Going deeper with automation

How to tame complexity in wireless networks

Monica Paolini, Senza Fili

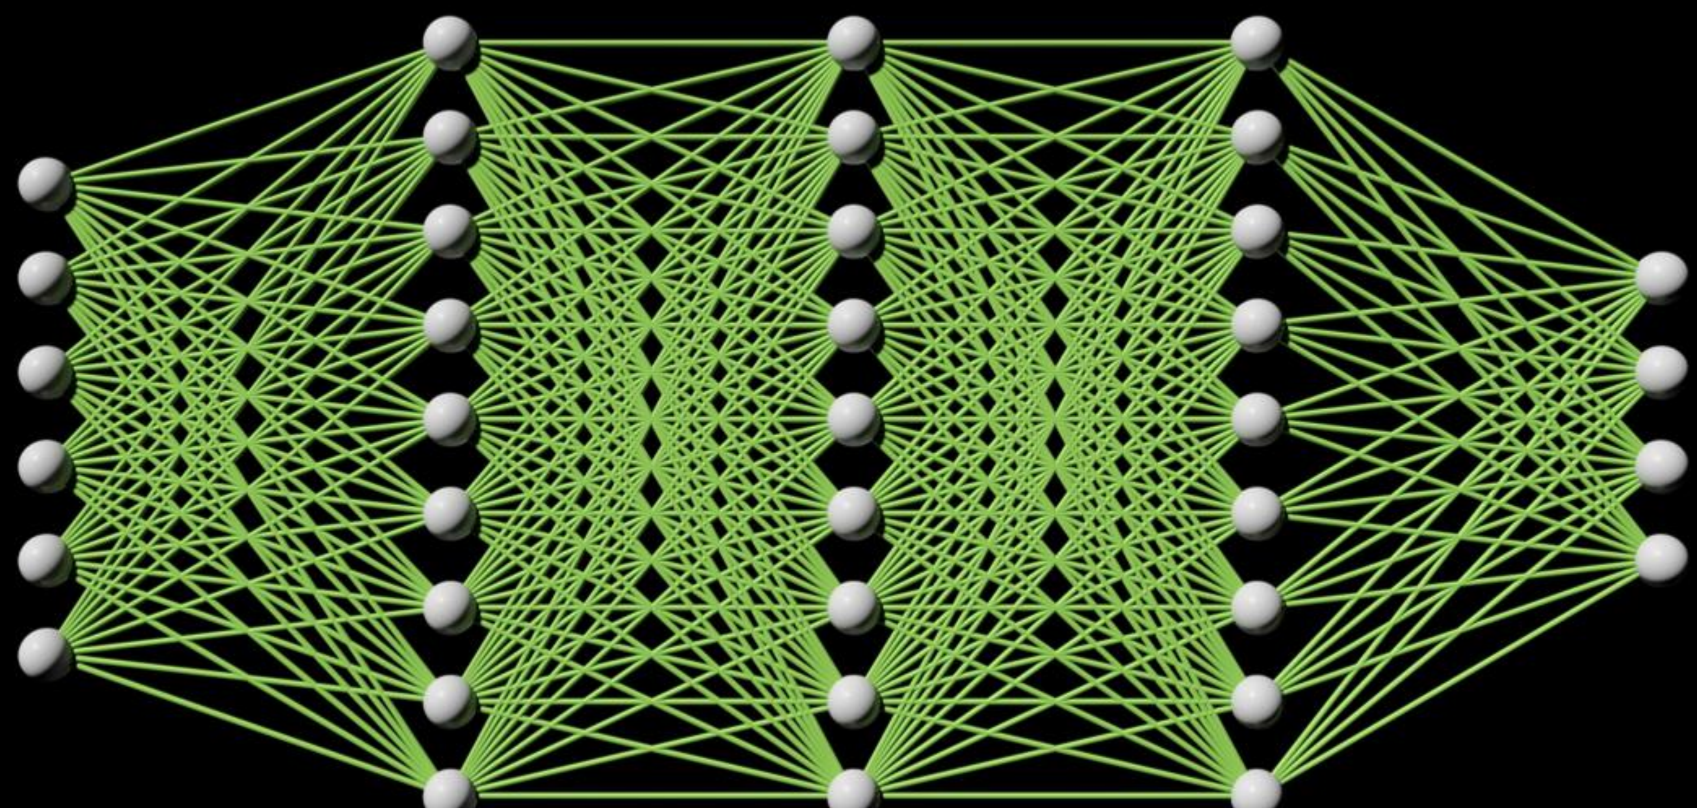

Report sponsors

HUAWEI

ciena
Experience. Outcomes.

BRIGHT MINDS SHARP SOLUTIONS | LUXOFT

TUPL

Additional conversations with

BT

In collaboration with
FierceWireless

Table of contents

Corpus Tetragonum, Secretum de thesauro experimentorum, Giovanni da Fontana, 15th century. A Renaissance device to automate memory and possibly text generation. Similarly to the wireless internet today, the invention of the printing press made content available in a new form and in a new scale. Overwhelmed and excited about the new opportunities, some creative minds designed and built machines to generate texts or memorize and store knowledge in ways that loosely anticipated today's more computationally heavy approach to the same issues. They had the right intuition, but not the technology and processing power to succeed, and only illustrations remain of these early attempts at automation.

Report

The view from above: Automation spreads as complexity grows

In wireless, there is a widely felt sense that with 5G we are on the cusp of a major transition, which some have compared to the industrial revolution or the introduction of electricity. A major transformation in how we communicate is afoot, and we are in a critical phase. But this process is much wider than 5G, although 5G may cover under its umbrella some of the technology tools needed to complete the transition. The initial catalyzers were the joint appearance and intertwined development of wireless communications and the internet. They have created an entirely new way for humans to communicate with each other and control the environment that no longer needs – for data, content or even some objects – to be instantiated in a physical object (book, piece of paper, banknotes). IP addresses guide us through what increasingly has become our reality.

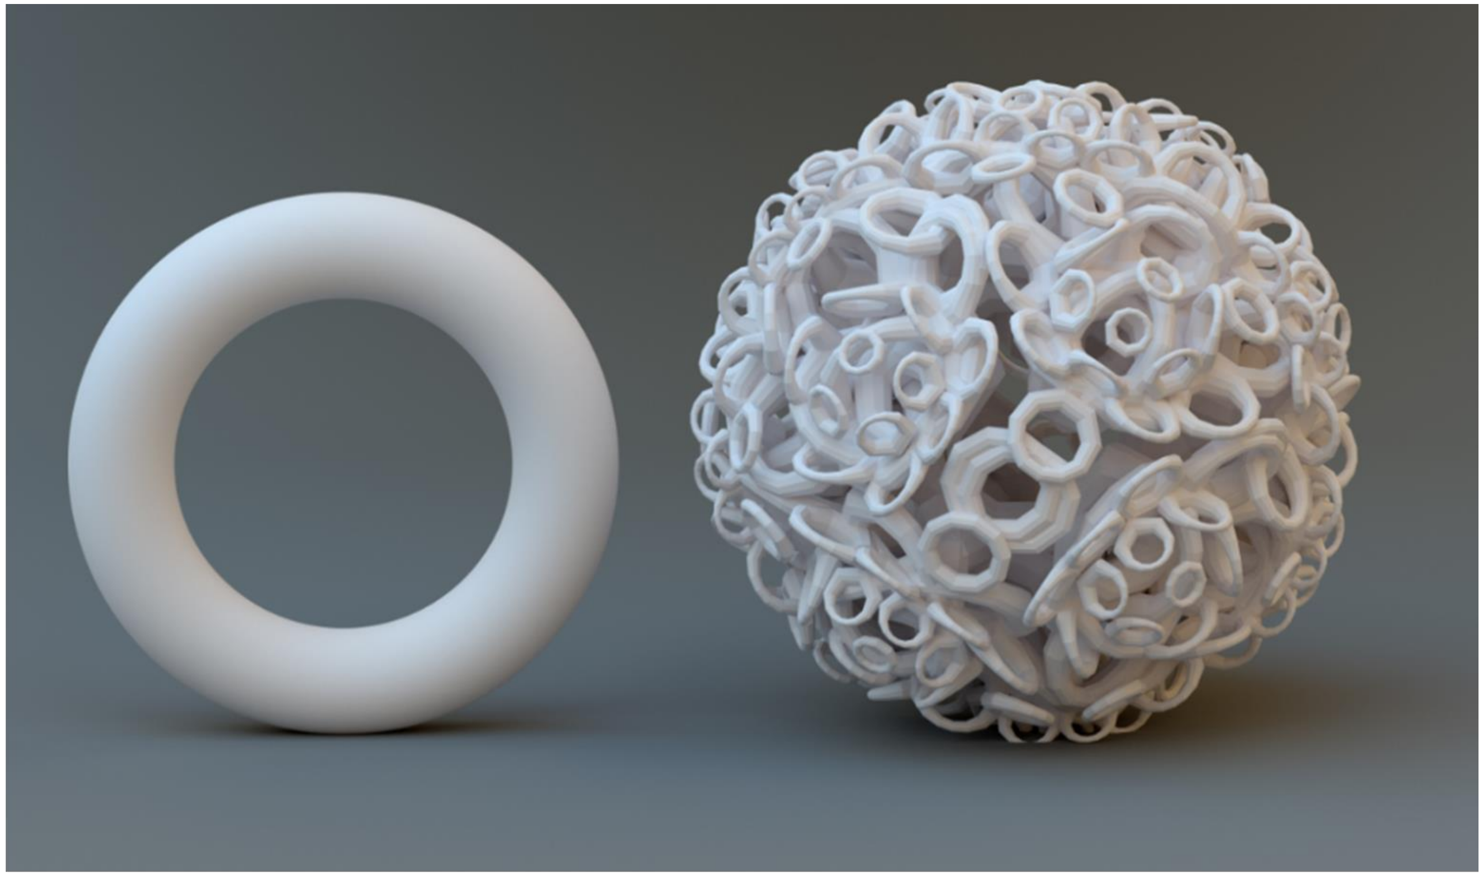

Growth in complexity

This process started decades ago. But now the wireless and wireline infrastructure that makes this transition possible is under strain and has to change at its foundation to cope with the demands and expectations that both the industry and the users have about connectivity. It is no longer sufficient to simply add more of the same – more cell sites in the RAN, more servers in the core. The need to scale in a massive way requires a qualitative change in how networks are built and run.

This is where automation comes in. The new automation solutions and platforms are no longer restricted to tackling simple processes; they allow operators to manage the increased complexity of wireless networks, deepen their understanding of how their networks work, better serve their subscribers, and take full advantage of the new technologies. The new breed of automation is much more than avoiding boring tasks and saving money. It is a core enabler that spreads across wireless networks end-to-end and deeply changes operations at both the technical and cultural levels.

This report explores the role of automation in shaping the evolution of wireless networks as we move toward 5G, massive IoT adoption, densification, and virtualization to meet a growing demand for a more diverse set of use cases.

1: Past / Future

Why has automation become such a hot topic now? What prevented operators from introducing it earlier? In truth, operators have been using automation for a long time – not as widely as one might expect, when compared to other verticals, but in a restricted manner, targeting narrow, highly repetitive processes. As we will see later in the report, operators are set to use automation in the future not only more widely, but in fundamentally different ways to fulfill a wider purpose.

Two critical enablers come from the broader ecosystem beyond telecommunications:

- Processing power has increased and become cheaper, making it affordable to deploy computationally intensive technologies or solutions that require crunching of very large data sets.
- More experience in using huge data sets from multiple sources – big data – and doing so, increasingly, in real time. This has given wireless operators the confidence and the tools to use big-data sources from their networks. The data has always been there, but operators still use only a small portion of it to derive main KPIs.

These changes prepare the ground for operators to expand automation. But it is the confluence of many drivers, and the synergies among them, with the move from what we call atomistic to pervasive networks that makes a compelling case for automation in wireless networks.

Pervasive networks emerge from the combination of virtualization, IoT, densification, and 5G to gradually replace today's atomistic networks. The table on the next page lists some of the directions in which wireless networks are evolving. Change is taking place across wireless networks – in different parts, at different depths, for different functions. The overall result is an increase in flexibility and efficiency. Higher complexity is the price to pay. As networks become less homogeneous, more layered, more dynamic and more capable of organizing themselves as needed, operating a mobile network requires keeping track of and optimize a far larger number of network elements and the interactions among them. In this environment, automation is necessary to efficiently optimize the network and allocate network resources appropriately.

For instance, in a multi-layered network, the combination of macro and small cells requires tools such as SON to coordinate transmission and manage traffic in a given area in an automated way. Done manually, this requires too much human effort to be cost effective, and most operators would settle for a less dynamic and less effective RAN management that would decrease performance and make resource utilization less effective – thus contributing a undesirable rise in per-bit costs.

End-to-end optimization is one of the major challenges in pervasive networks. In legacy, atomistic networks – including most of today's 4G networks – basic KPIs such as dropped calls and latency are used to optimize network performance, but this does not necessarily lead to the best QoE or network efficiency. In pervasive networks, maximizing QoE becomes the main target, but this is more difficult to do, because many elements contribute to a good QoE. Latency, for instance, becomes much more prominent, and the interaction of throughput and latency is key to assessing QoE, more so than just the two of them separately.

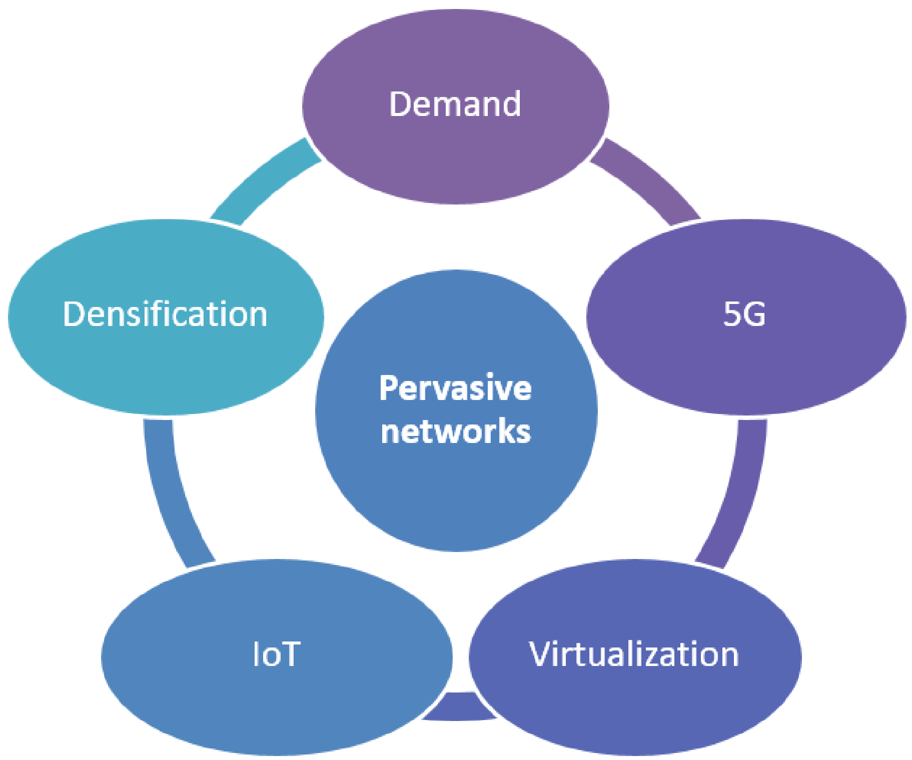

	Atomic networks	Pervasive networks
Network model	Network-centric: subscriber adapts to the network (e.g., goes to the window to make a phone call).	User-centric: network adapts to subscriber demand (e.g., ultra-dense wireless infrastructure in stadiums).
UE-RAN connection	One-to-one connection from the UE to the cell. Handoffs required for the UE to move association from one cell to the next.	UEs can connect to multiple antennas, use multiple bands. Flexible connectivity: device-to-device, dual connectivity, Wi-Fi. Multiple devices (including non-SIM and IoT devices) on the same plan.
User and control planes	User and data planes allocated to each access channel (e.g., sector).	Control plane can manage traffic for multiple access channels, so some access channels do not require a separate control plane (e.g., LTE in unlicensed bands, LWA, mmW). Short-range mobility can be managed without handoffs.
Densification targets	Coverage in the wide area, capacity in high-traffic areas, with most of the RAN infrastructure in outdoor locations and large venues (e.g., stadiums).	Vertical capacity increase and coverage extension driven by location-specific traffic or service requirements (e.g., service tied to a venue; IoT service).
Layers	Single macro-layer, with limited small-cell hotspot deployments, and with Wi-Fi offload.	Multi-layer networks, with extensive indoor and outdoor coverage with small cells, DAS or femto cells.

	Atomic networks	Pervasive networks
Spectrum	Cellular frequencies below 3 GHz.	Wider range of higher-frequency bands (3.5 GHz, 5 GHz, mmW), including unlicensed spectrum.
Core/RAN	Separate location and equipment, with RAN equipment located at the edge and core equipment in centralized locations.	Boundaries less strict, with CRAN and RAN virtualization, and some core functionality moving to the edge (e.g., MEC, CORD). Location of function (distributed versus centralized architectures) as a strategic decision.
Testing, monitoring, optimization	Testing and monitoring based on network KPIs and historical data. Limited optimization functionality.	QoE metrics based on the performance of UEs are tied to network KPIs to test, monitor and optimize networks in real time.
Capacity	Capacity per RAN element.	Capacity density (e.g., per sq km) and latency.
Traffic management	Maximize throughput. Capacity determines service availability.	Real-time traffic management, at the application or service level. Network slicing used to extract more value from network resources.
RAN equipment location	Telecom assets (e.g., macro-cellular towers, building rooftops), mostly in outdoor locations.	RAN equipment gets closer to subscribers and devices – closer to the ground and indoors.
Network ownership and control	Operator owns the network. Limited network sharing.	Multi-operator, neutral-host models, with some network elements (e.g., backhaul) shared among operators.

2: Simplicity / Complexity

Simple	Complex
Less scope for optimization	More dimensions to optimize
Performance easy to measure	Difficult to quantify performance
Maximize throughput	Maximize QoE
Stand-alone KPIs	Interaction of factors
Lower cost	More expensive, more cost effective
Well-understood cost structure	Financial models inadequate

When it comes to automation, complexity is a game changer. As wireless networks get more complex, automation does not scale with it linearly. It is not enough to do more of the same: automation has to do more, differently. The inputs are different, the processes to be automated are different, the methodology is different, and the goals are different. The next sections will discuss these differences, but here we present some of the key changes introduced by complexity that affect the role and purpose of automation.

In a simple network, the scope for optimization is limited. 2G and 3G networks are operated to maximize number of calls and throughput. They are largely static networks, so when they are fine tuned, there is no need to go back and make any changes unless there is a problem, a major change in demand, or an infrastructure change (e.g., a new cell site added). Automating the fine-tuning process can still be valuable, but the impact of automation is limited because it is a relatively straightforward task for an expert.

In a complex and dynamic wireless network, the scope for optimization widens. Multiple dimensions and variables can be jointly fine tuned to achieve the desired network performance – which is not necessarily the highest throughput, or the lowest latency, but more likely the combination of metrics that gives the best QoE, as defined by the operator. And each operator is likely to have its own way to quantify QoE and maximize it.

The wider scope for optimization expands the role of automation. Not only can automation cover more processes, it can (or eventually should) cover the interactions among them. Because the interactions can vary through time and network condition, automation can no longer be a prescriptive, rule-based mechanism. It has to constantly learn and adapt in order to operate in an ever-changing environment.

One example of these interactions is the impact of latency on throughput in TCP traffic. Most video traffic uses TCP, and video accounts for 73% of wireless traffic, according to Cisco VNI. TCP ensures good signal transmission by requiring the receipt of an acknowledgement from the receiver (the UE, for a downstream flow) before continuing to send packets. In a network with high latency, TCP messages in the downlink and uplink take too long and cause delays in sending video content, which degrades the user experience. From a network point of view, however, the delay caused by TCP acknowledgments in the uplink looks like lower throughput in the RAN downlink. The RAN may work flawlessly, but it is performing below capacity because the high latency slows down the acknowledgement process and, hence, the transmission.

In this case, maximizing throughput without regard to latency will not fix the problem. Lowering latency without changes in air interface or architecture is typically not an option. So TCP optimization has to take into account the goal – e.g., good video QoE according to the operator's definition – the latency, and the throughput. Clearly this cannot be accomplished manually or with a simple script, because of the high number of possible combinations of the variables involved. In cases like this, automation is not an alternative to a tedious and repetitive process, but the necessary basis to an effective TCP optimization.

Of course this does not come for free. There is a cost associated with the introduction of network-wide automation, but we expect the increased efficiency in using network resources to justify the investment because the overall infrastructure becomes more cost effective. As we see later, financial models such as the ROI and TCO are not well suited to capture these financial benefits: along with new wireless architectures we need to develop new models to quantify more accurately the financial implications.

3: Scripts / Learning

The main challenge with introducing automation in complex, pervasive wireless networks is that we do not know how to do so. We are still trying to figure out how best to operate these networks and extract the best performance from them. If we had this knowledge, introducing automation would be straightforward. Instead, we need automation as part of both the learning process (i.e., to learn to automate) and the operating process (i.e., to automate).

Much of the automation in today's networks has a narrow scope and is aimed at processes that can be automated because they are well understood. This information can be encapsulated in a script or similar tool created using the expertise collected by vendors or operators over the years, and is used repeatedly. Unless there is a network upgrade or a change in requirements, there is no need to change the script. It relies on a small data set that is easy to access, and follows a deterministic algorithm, that uses few resources. The main goal is to speed up the task, save employees' time, and cut costs.

While this type of automation is a positive step forward – but surprisingly still limited in adoption despite the obvious advantages – it falls short of what is needed in complex networks. Automating a specific task that can be isolated from the rest has only a marginal impact, and may even introduce some static behavior that could have a negative impact on a dynamic system.

In pervasive networks, automation creates a framework in which some tasks can be done with minor human assistance or none at all. And of greater value, automation here means network data can be used to learn about, and optimize, network performance and health. Although automation in a pervasive network can be applied equally to narrow and broad tasks, its reach is deeper and wider because it does not run on top of tasks but becomes an integral part of them. In the TCP optimization example in the previous section, automation is not used to simply direct how to use TCP but provides a way to relate TCP to real-time network conditions.

The learning aspect is the major change from today's automation – the most challenging to introduce, but also the most rewarding. In networks that are not just complex but also inherently dynamic, learning is not a one-off phase, to be completed at the beginning. Learning is a continuing effort, both to refine what has been learned to date, and to keep up with the changes happening in the network.

Expertise and experience from the vendor's and operator's staff is the starting point for setting up an appropriate learning model and supervising the process. Equally important is the contribution from internal (network) and external (third-party) data (see next sections) as the basis for learning, operating and optimizing.

This is where manual intervention is no longer sufficient, because extracting the relevant information, correlations and causal relations from a large body of data is difficult for humans – and downright impossible to do in real time. Humans are very good at high-level analysis, intuition and judgment, but our skills do not scale well as the volume of data increases, which is what happens in pervasive networks.

Tools and techniques from AI, machine learning and deep learning need an automation framework to process network data from multiple sources that changes continuously. While they do have high processing requirements especially if they operate in real time, their contribution may be critical in extracting the benefits of 5G and virtualization technologies, in developing new use cases, and in unlocking opportunities for new services.

Scripts	Learning
Surface	Depth
Narrow scope	Variable scope
Simple problem, known solution	Complex problem, unknown solution
Based on experience and expertise	Experience and expertise, plus data
Narrow data sources	Multiple data sources, external/internal
Static	Dynamic, continuously improving
Algorithms	AI, machine learning, deep learning
Low processing requirements	High processing requirements
Slow	Fast
Avoid tedium	Create opportunities, optimize

4: Historical / Real time

Dynamic networks change through time, and to use network resources efficiently, operators have to manage them in real time. They need to take into account the current network conditions, fluctuations of demand, performance issues, and anything else that may be relevant – e.g., traffic conditions, weather, major events.

Does this mean the historical data on which current automated processes are based is no longer relevant? Not at all. Historical data is the repository of essential information – the memory of the wireless network – and is the foundation for learning. Historical data frames the context that AI and ML need to structure the learning process, and helps operators select the best tools and techniques to meet their goals.

Many discussions of big data appear to assume that, as long as there is a massive amount of data, we can extract what we need from it by feeding it to the AI or ML tool of choice. This creates a misleading impression that muscling through data is a sure-bet process that will make the expertise collected through decades of experience irrelevant, along with the humans who collected it.

Quite the opposite. Learning how to operate and optimize dynamic and complex networks is not going to be simple or straightforward. It will require time, and we should anticipate setbacks as we make progress.

As we move to capture the dynamics of wireless networks by including real-time data from multiple sources, we face two seemingly opposite challenges.

On one hand, we want the flexibility to incorporate all the available data that might be relevant. The challenge here is to be able to port data to the automation/learning platform – data that at its source might be unstructured, or structured for different purposes.

On the other hand, we want to be as selective as possible, because learning is more effective with smaller volumes of higher-quality data, even if we have access to high-speed processing power. To do that, we need to grab the relevant data and exclude the rest. Establishing relevance is tricky. We cannot rely exclusively on experience, because then we risk introducing a bias against data sources we have not used before, possibly limiting what we can learn.

The relevance of history. Fabergé egg with clock, Fabergé Museum, St. Petersburg

Treading between these two goals will require fine-tuning, but will be essential to an efficient, cost-effective implementation of automation and learning that is neither bloated nor simplistic.

One of the crucial decisions in moving to real-time network management is deciding what counts as real time – selecting the appropriate time granularity (how frequently you collect data samples) and window size (for how long you collect them).

There is no definition of real time that applies across the network. What real time is depends on the task and its goals. For instance, real-time optimization of beam forming and massive MIMO will require a higher temporal resolution than real-time policy management.

But for each given task, using data with too high a granularity (i.e., too much detail) and too big a window (i.e., looking too far into the past) leads to an unnecessarily data-heavy approach, and it may be slower and less efficient. At the other end, less-detailed data may not be sufficient for the automation/learning model to converge to a satisfactory solution (i.e., not learning all there is to learn and hence not optimizing at the desired level).

Finding the right definition of real time goes beyond what is best suited for a specific task (e.g., massive MIMO or policy), and involves both capabilities and cost tradeoffs. At the margins, increasing time granularity for a task may no longer yield enough benefit to justify the increased cost and effort. However, as real-time solutions become more widely available and deployed, we can expect that, in general, the definition of real time will move toward higher granularity and, in some cases, shorter window sizes.

Even though it will require effort from operators, the real-time learning process that automation both requires and enables is where the opportunities lie. Historical data gives us the basis to define the challenges and what we need to learn. Real-time data opens the way for better optimization, new service creation and cost efficiencies.

5: Fix / Predict

Why stop at real time? As time granularity increases, we get close to the edge of the cliff, and we can jump off into predictive mode. With historical data, automation helps us manage what we expect from what we know already. With real-time data, automation allows us to learn, and to fine tune the network based on what we just learned. Predictive data gives us some control over what is about to happen, even though we may not fully understand it.

With prediction, the goal of automation expands from running, optimizing and fixing the network, to preventing problems and optimizing the network ahead of time. This is an improvement over real-time operation and, in many ways, a natural extension of it. If an operator understands the dynamics of the network in real time, it can also learn to anticipate what may happen next, by comparing historical data (e.g., what triggers a problem) with real-time data (e.g., are these triggers present now?).

Not everything can be predicted with a sufficiently high level of confidence, but there is a potential value, even if some predictions turn out not to be accurate. The cost of preventing a problem is generally lower than the cost of fixing it after it happens, especially if we take into account lost revenues and impact on subscriber QoE.

Let's take congestion as an example. If a cell site is congested, remedial actions can be taken, but only after subscribers notice – and the remedies may not take effect instantly. If the traffic load pattern or unexpected events in an area suggest impending congestion, the operator may shift traffic management to a protective mode – for instance, by compressing video or changing traffic priorities – that will reduce the impact of congestion on the RAN. The operator can use the same predictive capabilities to revert back to the default traffic management mode when congestion is gone.

While the operator cannot eliminate the increase in demand, it can diffuse the impact of congestion and manage it to minimize the impact on QoE. This is likely to reduce the cost of congestion if it results – as we would expect – in fewer customer support calls at the time the issue occurs, and in lower churn in the long term.

Blue Skillet, Seattle

6: Vertical / Horizontal

Automation today addresses a need to improve efficiency or reduce costs that is mostly at the task level. If automating a task brings a financial or performance benefit, the operator or the vendor will overlay an automation solution. Frequently, this is a solution-specific and vendor-specific tool that directly and exclusively interact with the solution – but mainly to automate a specific task, rather than to extend automation across domains. The tool sits comfortably within the confines of existing organizational silos, within both vendor and operator.

As automation evolves to meet the needs of pervasive networks, it is no longer tied to a solution, equipment element or task, but becomes an approach and/or a platform that is deployed horizontally across the network. The organizational, operational and financial investment needed to implement it requires a broader approach, even though we expect the adoption to be gradual across the network.

As a horizontal effort that spans the network end-to-end, automation needs access to multiple internal and external data sources. It becomes a much more powerful and disruptive agent of change – one that breaks traditional boundaries inside vendors and operators alike, and forces business units to collaborate more closely.

Initial efforts may target a single or a few automation use cases (e.g., TCP optimization, backhaul management, or customer care). This enables the operator to become familiar with and confident of the new approach. The initial efforts can subsequently be expanded to other use cases, because the data sources, the learning tools, and, indeed, the learning can, to a good extent, be shared.

The congestion example mentioned earlier illustrates how automation in learning about and managing RAN traffic may subsequently be used in automating customer care. Knowing where there is congestion and what is being done to address it can help inform the customer care response to congestion – for instance, by proactively providing information to customer support staff, or by taking action before subscribers call in (e.g., by sending them alerts about upcoming congestion or offering incentives to reduce usage). When adding backhaul automation, customer care can gain wider visibility into the network and can use this knowledge, for instance, to identify the RAN or the backhaul as the cause of QoE degradation.

If RAN and backhaul optimization were done manually or through task-specific solutions, the scope for sharing the benefits of automation would remain local and not spread to other parts of the network. An end-to-end automation platform makes the learning and automation process transparent across the network. It also makes it traceable. And that information itself can be used to deepen the learning process in the same area of the network or in another one.

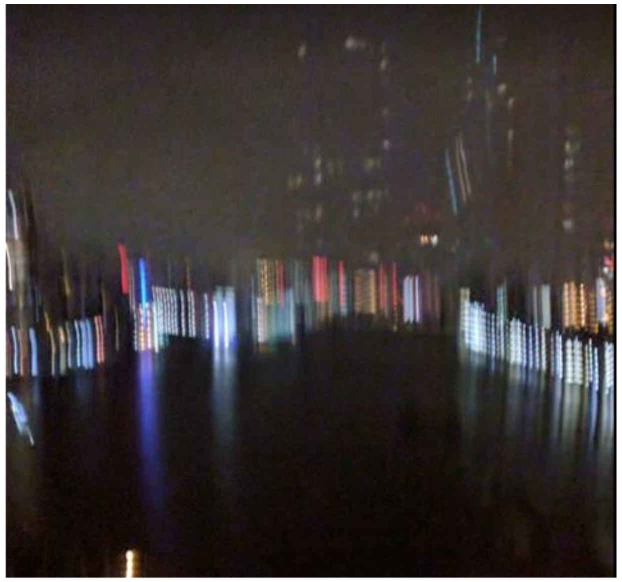

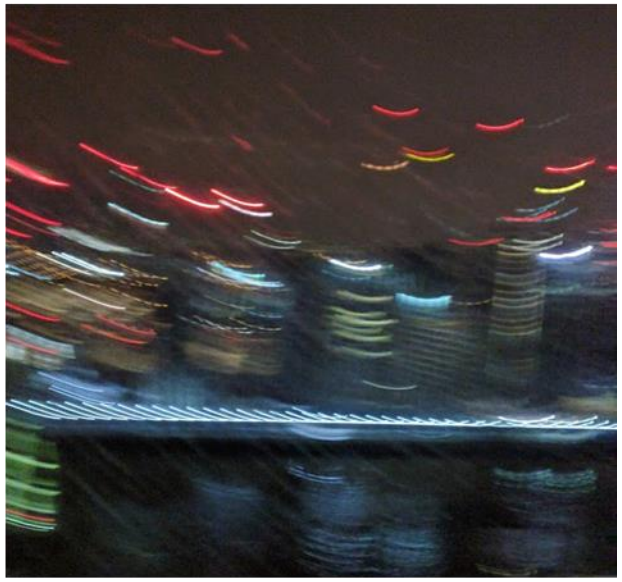

London, two views at night

7: Legacy / Virtualized

The most powerful driver to automation is virtualization – more powerful than 5G, in terms of both timing and impact. 5G will require automation, too, but more because of the increased volume and complexity of traffic. The complexity spawned by virtualization, in contrast, is concentrated at the level of network architecture where automation can be of the most benefit.

By separating function from hardware – allowing flexibility in where functions are physically located – and by allocating resources in real time, virtualized networks change continuously. Network management has to adapt to this dynamic environment. Tuning, resource allocation, monitoring, optimization and fault resolution can no longer be done manually or only with minimal automation – nor can they be done occasionally at set times, as they can in legacy networks. As the network continues to change, they need to operate continuously, and this massively raises the processing requirements for network management – from testing all the way to optimization.

In a static, legacy network, it is easier to track performance, identify root causes and fix problems. Operators largely know what can go wrong and how to fix it. This is not to say that managing 2G, 3G or 4G networks is easy; even in these non-virtualized networks, the step-up in complexity from one generation to the next has enlarged the burden on network management. But it is something that operators have learned to handle with limited automation.

In a virtualized network, the number of moving parts – and hence the interactions among them – is much higher. There are more sources of problems, more difficult to track down. The same issue may appear at different times or different places; the root cause may be different and more difficult to identify.

In this environment, much of network management has to be automated simply to keep it sustainable within current staff and funding levels.

This is not bad news, however. Virtualization makes the adoption of automation more urgent, and this allows operators that embrace virtualization to benefit from automation's wider benefits – the cost savings, the more efficient allocation of network resources, and the optimization of QoE.

Legacy networks	Virtualized networks
Known set of problems	More things can go wrong
Planning and testing restricted to initial deployment phase	Testing widens in scope, in parallel with monitoring
Historical performance	More difficult to keep track of performance
Root-cause analysis leads to specific network elements	Root causes are complex to identify, need to deal with interactions
Fix problems	Optimize resource allocation
Automation limited to specific tasks	Wider scope for automation

8: Technology / Culture

Technology may be the easy part of automation. Even though automation requires substantial effort, and its impact on the network is deep and wide, and the roadmap and requirements are still works in progress, there is wide commitment from operators to move toward automation, and from vendors to provide solutions that address operators' requirements. Adoption of automation in other sectors has proven beneficial, and it is reasonable to expect a similar outcome in wireless.

The bigger challenges to automation are cultural. Automation introduces a cultural change that is going to be initially disruptive, even though in the long run it will be positive. Technological change can be rolled out according to a plan. Cultural change could meet a resistance that may hamper or slow down the technological change.

Mention of automation – and even more so of AI – often generates fears that humans may no longer be relevant or needed, at least not at the current levels or in the current roles. But we are far from a world in which robots have the capabilities to take control of economic activities or even social activities, and we are unlikely to get there in the foreseeable future – if for no other reason than we humans have not been able to build machines that exhibit the kind of intelligence humans are effortlessly capable of. But what automation, AI and ML promise to do is to alleviate the burden of running the increasingly complex, high-performing and sophisticated networks we are planning to deploy, and make them financially and operationally sustainable at the current levels. They are tools that make wider technological advances possible, rather than a threat to human relevance or intelligence.

For a successful adoption of automation, it is crucial to articulate and manage the expected cultural changes within the ecosystem players – operators for sure, but vendors as well – and to address the concerns of both staff and society at large.

Probably the most commonly voiced concern is that, with automation, operators will be able to shed some of their workforce, because some jobs will be done by a combination of servers, drones and robots. While this is true to an extent, the increase in complexity is such that it would not be possible to run future networks with today's levels of staff. And without automation, the escalation in staff costs could prevent or slow down the pace of technological evolution.

Firing staff may not be the main worry. Instead, what operators will need to do is retrain their employees to manage automated processes. This will give employees the opportunity to avoid repetitive tasks and use their expertise more extensively and make their jobs more rewarding.

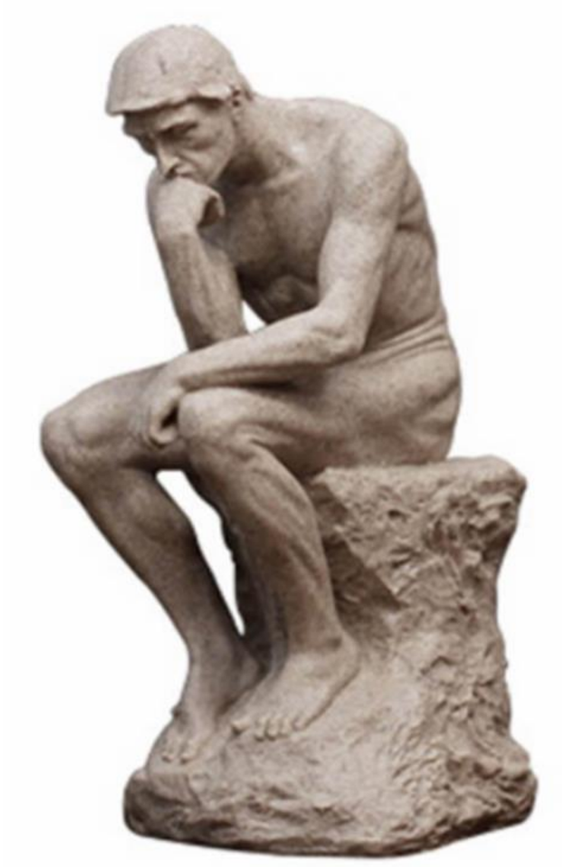

Rodin,
The Thinker

9: Cost / Savings

It is not difficult to make a business case for the task-specific type of automation used today. It requires an upfront investment that is not very high and that can be recouped by reducing staff costs and time-to-market.

A similar business case can be made for the new automation approaches that traverse the entire wireless network. Because of their wider scope, the cost is going to be higher, but the benefits will also be greater. But it is more complicated to assess the full extent of the financial benefits of an approach to automation that traverses the entire wireless network, compared to a task-specific automation solution.

As we discussed in the previous section, although automation will reduce staff costs and time-to-market when compared to the case without automation, an even stronger case can be made: without automation, managing highly complex networks will become unsustainable, both financially and operationally. In this context, automation is primarily an enabler and secondarily a cost-saving solution.

For specific uses of automation – those that operators are likely to adopt initially – the business case for task-specific automation can be used to assess the short-term cost savings.

However, this leaves out the higher investment needed to make the deeper end-to-end changes that move to automated networks from automated tasks, including indirect costs, such as those linked to staff re-training. But also, and more importantly, it leaves out the long-tern financial benefits from automating the network and using automation to optimize performance and QoE and to allocate resources more efficiently.

These benefits are more difficult to capture. They escape traditional ROI and TCO models that track costs and revenues on specific network elements or functions, but they are likely to be the largest benefits, especially in the longer term (i.e., beyond the first 2–3 years from initial adoption, assuming a gradual adoption that has the effect of delaying the end-to-end benefits). The same limitations of ROI and TCO models make it equally difficult to assess the long-term benefits of virtualization.

To assess more accurately the impact of a network-wide change – whether automation or virtualization – that disrupts the demarcation lines separating network areas, it is necessary to move past the ROI and TCO models, and evaluate the network-wide impact.

For instance, in the example of congestion management, better allocation of network resources may indirectly lower customer care costs and churn as well as making the RAN more efficient. But in addition, a more efficient RAN will reduce the need for capacity upgrades. This does not affect the initial investment, but it reduces the long-term investment needed to meet a growing capacity requirement.

10: Choice / Necessity

To date, automation has been a matter of choice. All operators have some tasks automated, but the degree of automation, its depth and the type of tasks that are automated varies greatly. In some cases, automation is developed internally by the operator; in other cases, a solution is provided by a vendor.

Driven by pervasive networks, network automation, as we have seen, is no longer a nice-to-have option. It becomes a necessity for dealing with the greater complexity, especially as networks become virtualized and 5G approaches.

However, operators have a wide range of choices as they chart their automation strategy:

- **Start.** Some operators are more aggressive than others in implementing automation. The more innovative operators at the forefront of virtualization are also in the vanguard of network-wide automation. Other operators may be committed to expanding automation but not ready to do it yet or still do not see it as urgent.
- **Pace.** Adoption is going to be gradual. Most operators will start with some network areas or some tasks where they have a strong need for automation or they believe it will be most beneficial. Some may work on the learning component later or prefer, at least initially, to keep a narrow focus. A few operators, though, may adopt an end-to-end learning/automation platform at the outset.
- **Area.** Similar to what is happening with virtualization, operators will start with both trials and initial commercial deployments in targeted areas. SON is one of the early automation efforts; others may follow in optimizing other aspects of the RAN, in customer care, or in network slicing.
- **Depth.** The endgame is to get as close as possible to zero-touch automation, in which tasks are completely automated and humans only supervise them. Again, movement in that direction will be uneven, and some tasks will continue to need a degree of human intervention or supervision. As we move to real-time operations, zero-touch becomes almost a requirement, with human intervention restricted to oversight of learning and to problematic cases.

Interviews

Ciena

As a leader in packet and optical networking, Ciena is deeply committed to shaping next-gen networks by promoting virtualization, openness, collaboration and automation. Automation is a key element in supporting the scalable, programmable, adaptable and dynamic networks that will be the foundation of 5G.

Ciena offers a wide portfolio of software automation solutions that, along with its networking equipment, enable operators to transition from today's static networks to more flexible and dynamic virtualized networks

Liquid Spectrum is based on an agile, flexible and dynamic software-defined network architecture. It combines programmable hardware with software applications to optimize the use of network resources – in terms of efficiency, capacity, reach, availability and automation.

With Liquid Spectrum, operators can deliver bandwidth and services on demand. They can use the software platform to allocate resources and move capacity where needed, when needed. They can create new services quickly and extract higher revenues from their existing assets.

Cloud-based applications within Liquid Spectrum give operators additional tools to manage their network's performance and asset allocation:

- Performance Meter measures performance SNR in real time at the single-channel, path or network level, to monitor and optimize performance, plan for new services, or so historical trend analysis.
- Bandwidth Optimizer allows operators to avoid the manual selection of wavelength to set capacity; instead, it recommends the best configuration for allocating the required capacity. Using service policies, Bandwidth Optimizer can also respond quickly to the demand for new services.
- Liquid Restoration uses Performance Meter and Bandwidth Optimizer to increase service availability.
- Wave-Line Synchronizer accelerates service provisioning by automating the synchronization of modems and photonic lines.

The Blue Planet Intelligent Automation Platform is an open software suite designed to give service providers access to network performance and use analytics. This insight helps providers enable end-to-end automation, optimize services and operations, provide on-demand services, and reduce costs. The product includes the following functionality:

- Multi-domain service orchestration (MDSO) for end-to-end service automation across physical or virtual domains
- NFV orchestration for instantiation, management and chaining of VNFs in centralized, distributed, and hybrid topologies
- Analytics to collect, process and store network data, and for use with machine learning to identify trends, predict usage patterns, and anticipate network and service disruption
- Network health predictor (NHP), to assess network health and conditions historically and in real time, helping operators work proactively to address issues before they arise
- Manage, control and plan (MCP), a life-cycle management software tool to automate operation of Ciena's equipment, increase agility and flexibility in the network, and speed up service creation and delivery

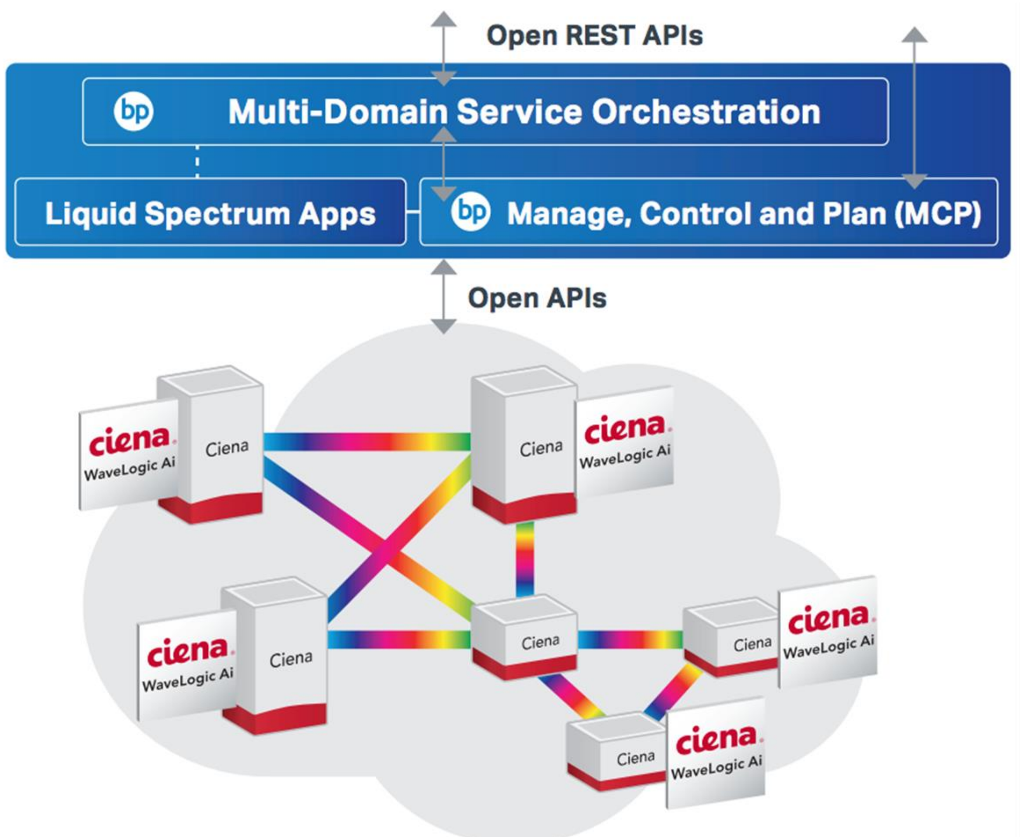

Example of a Liquid Spectrum implementation, using a fully integrated Ciena Solution *Source: Ciena*

Ciena

From automated to autonomic networks

A conversation with Kaniz Mahdi, Vice President, Advanced Architectures, CTO Group, Ciena

Automation in wireless networks is still surprisingly limited in both scope and depth, but operators have been trying to automate some functions over the years. This type of automation has been driven by a top-down, static approach that translates the knowledge of engineers into scripts that execute repetitive actions. As networks become more complex and powerful, though, this type of automation is no longer sufficient.

With virtualization and 5G, automation becomes more crucial to ensuring even basic operability. And automation itself changes in fundamental ways to enable wireless networks to eventually self-monitor, self-configure, self-optimize, and self-heal, with engineers overseeing and monitoring the process but not driving it.

I had the pleasure of talking with Kaniz Mahdi, Vice President, Advanced Architectures, CTO Group, at Ciena, about the transition from static automation to new, autonomic networks.

Monica Paolini: Kaniz, you have been working on automation for a long time. What is your current role at Ciena?

Kaniz Mahdi: I head the advanced architecture team at Ciena. One of the key mandates of my team is to identify areas of new growth, to help Ciena spread its wings beyond what would be normally categorized as a natural evolution of the current portfolio.

As far as automation is concerned, we already have a very broad suite of tools in the tool kit. Automation has been in Ciena's DNA for many years, ranging from automated configuration of our Packet Networking portfolio – with SAOS, our Service-Aware Operating System – to the recently launched programmatic software control of our optical gear with Liquid Spectrum. Then, of course, our Blue Planet Multi-Domain Service Orchestration (MDSO) management platform is focused primarily around life cycle management of physical and virtual network functions.

On the industry front, we are proud to be founding members of ONOS and CORD, which are industry flagship projects on software-defined networking and, to my knowledge, one of the very first implementations of automated networking control.

We are also members of ONAP, the Open Network Automation Platform, which is on the path to becoming the de facto industry reference on open network automation.

In addition, we are leading various research and industry initiatives, the most recent being the Edge Automation working group that we are co-chairing with our partners in AT&T, under the auspices of the IEEE 5G initiative.

Monica: What about ETSI MEC?

Kaniz: ETSI MEC is one of first few initiatives in this space. ETSI MEC started out with RAN APIs for programmatic interface exposure from classic radio access network (RAN) solutions. They have recently evolved toward SDN NFV.

IEEE Edge Automation WG targets disaggregated RAN, and its integration with automation platforms – for example, ONAP – for automated control of spectrum resources. It's much more involved on the control side than ETSI MEC, which is mostly focusing on the application enablement of the RAN systems.

Monica: Automation is nothing really new. We have been trying to automate the networks for a long time. Why is there such an interest in automation right now? What's new?

Kaniz: That's a very good question. Process automation has been around for quite a while. When I was in engineering school, in fact, I did a project on PLCs, automated control of electromechanical systems – and that's 23 years ago now. Then you have automation of OSS/BSS workflows, which has been in our industry for quite a while now.

What's new, Monica, is basically the self-automating aspect of it. The automated systems of today are mostly very simple automation processes driven by static policies. They are fine for the systems of today but have limitations and restrictions on what you can and cannot do.

One of the first uses of self-optimizing or self-governing systems came about in radio systems, where we have SON, Self-Organizing features built into the RAN, and it has been around for three or four years.

But when it comes to NFV and SDN, and bringing all the piece parts together, static policy-based automation does not help you much. You would have to look more deeply into automation techniques that rely on self-learning systems or self-driving systems.

Artificial intelligence, machine learning, and data analysis have been available for a long time, but the use of analytics, machine learning, and artificial intelligence in telecommunications has just recently shown some promise.

On one hand, analytics and machine learning are coming to fruition, while on the other hand, a need is created by the virtualization of network functions.

You marry the two, and then the combination is what we call autonomics. In the most simplistic terms, it's a collected control of multiple, possibly heterogeneous, discrete close loop control (e.g. self-driving systems) which are governed by common business goals. You need that when you're bringing different autonomous systems together.

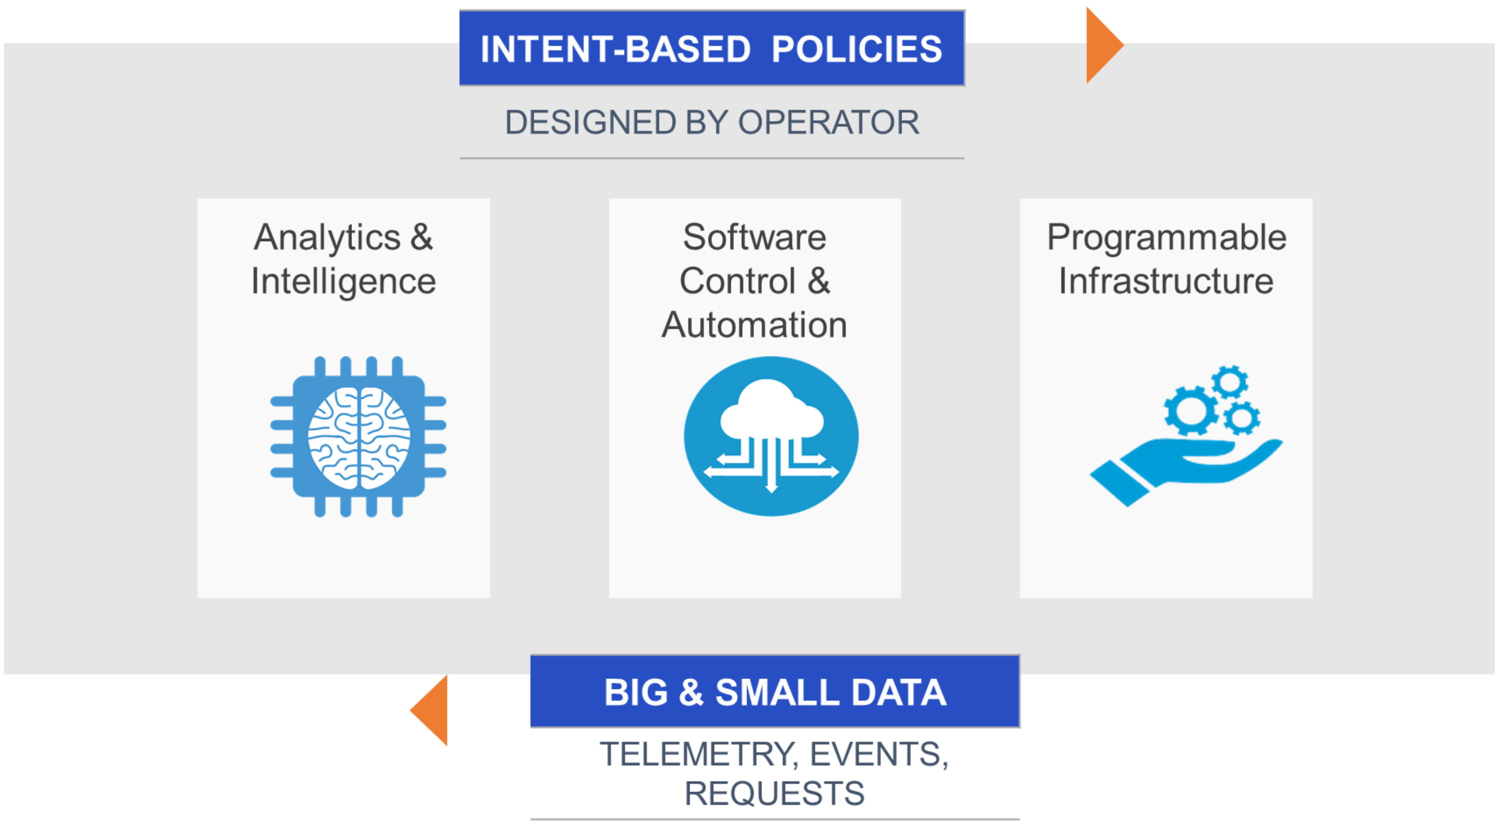

The role of automation in managing wireless networks *Source: Ciena*

Monica: Where are we now in terms of moving from automation to autonomics?

Kaniz: SDN was a game changer. It enabled separation of control plane and data plane, which is essential to automating systems using externalized programmatic controls.

NFV has also helped us move a few inches. I'm hesitating here a bit, because straight virtualization does not really help you too much with externalization of control.

We have made some progress, but lots needs to be done, specifically around the telcos' brownfield systems, which have been running and operational for ages, in many cases.

The current wave of virtualizing network functions has taught us some lessons. One is that – unless you rethink the applications, the way they are written, and the way they expose control – data is essential to making any progress toward a fully automated system. That's where the real challenge lies at this point.

Monica: You mentioned AI, machine learning, and analytics. What is the relationship between autonomics and those tools?

Kaniz: The relationship is intimate, of course, because you can do autonomics with static policies, but then you end up in a limited, restricted, closed-loop system, which is not able to react or adapt in real time, and not very intelligent.

Machine intelligence, which includes artificial intelligence, and all facets of machine learning – supervised, unsupervised, reinforcement learning – all of these play a critical role in autonomic network control, with the key reason being that today's automation systems rely on static policies. And although you have big-data analysis already embedded in these automated systems, what you have is data collected over a significantly long period – hours, days, weeks, and months. You have folks sitting in the back end doing data analysis, data scientists enrolled in generating policies, updating policies, a typical human-in-the-loop kind of control.

This is fine when it comes to designing, configuring, and controlling the simple systems of today. However, when it comes to deep nesting of disparate control loops, most of these being real-time control loops, it becomes a real challenge. It's almost impossible to predict and prepare for the needs of these control loops ahead of time, even when looking at the data collected over days, weeks, and months.

What you need is a system which learns by itself, with self-learning policies and self-learning control, which is monitoring itself and then adapting to the changes exerted from within and from outside. That is only possible with machine learning techniques.

The short response to your question is yes, they are integral to each other. However, although AI and machine learning have made quite an impact in other industries – as evident with the fact that they're around us everywhere – using machine intelligence for automating network control is still nascent in research.

Ciena has been engaged with academia on this for a while. We are quickly learning that, even in academia, a lot of effort needs to take place before we get on a solid footing, where we can say, "Yes, I can do a reinforcement-learning-based automated control of an optical system," for example.

I would like to use this opportunity as a call to the community. We're seeking collaborators from the industry who are willing to come help us disrupt ourselves and take on what seems to be an impossible task. I think that with help from the like-minded in our industry, we should be able to make it a reality sometime very soon.

Monica: You raise an important issue. This is a challenge to the whole industry; there is no single type of player or individual player that can solve it on their own. It really requires players to work together to find solutions, to get the different parts of the network and the different players to join together.

Kaniz: Exactly.

Monica: How does that research work relate to what you hear from the operators themselves?

Kaniz: The operators are doing the same. Some of the work we have been engaged in with academia is a partnership along with the operators. They're also learning alongside academia, alongside the vendors, and alongside the startups and the data scientists.

You said it quite well. It's going to be a collaboration, a learning exercise for all of us, a coming together of sorts, from the IoT industry, from the cloud industry, from the telcos. The biological sciences, as well, because autonomics takes inspiration from how the brain works and how it controls different parts of our body.

It's a collective effort, based on the pulse that we have at this point from around industry, academia, and the vendor community. Almost everyone is learning at this point and trying to figure out a way into collectively joining forces.

You have pockets of collaboration. You have the IEEE 5G initiative, where we are trying to see how we can do some AI-based automation for the edge. You have the ONAP community, which is looking into how to integrate the real-time control of virtual network functions – for example, the radio systems, the optical systems, the packet cores – bringing them into a real-time control loop, which is controlled, in turn, by a system like an ONAP or ECOMP or MANO framework. You have the universities, which are looking into modeling telecommunications systems. That will help to train the algorithms that are needed to drive some of these controls.

Monica: Let's talk about 5G. Obviously, automating a network isn't technically tied to 5G, but in 5G, you do need automation to effectively scale. At the same time, 5G might accelerate the automation of wireless networks. How do you see the relationship between the two?

Kaniz: If you look at the history of how packet systems have evolved from 3G to 4G, and now getting into 5G, the architects, myself included, were pressed to get things out the door most of the time when we were designing these things. For example, the packet gateways, the security gateways, the mobility management system, and how you create a connection between two wireless endpoints, it's all basically bolted on top of the packet network that was designed originally for the internet.

We ended up building monolithic gateways, which were bolted onto a fully operational packet system. That resulted in quite a lot of complexity, which was OK, because these systems were self-sufficient. For example, a packet gateway is, in itself, a self-controlling entity; it controls its own forwarding plane.

The same goes for the radio networks. You have a packet processing function and a control function, and, just recently, the SONs, self-organizing networks.

These wireless packet systems were enabling simple communications, which were mainly point-to-point, human-to-human, or even simple machine-to-machine communications. Later on, video was added on, and virtual reality. But even then, there should not be any problem serving these use cases with existing wireless architectures.

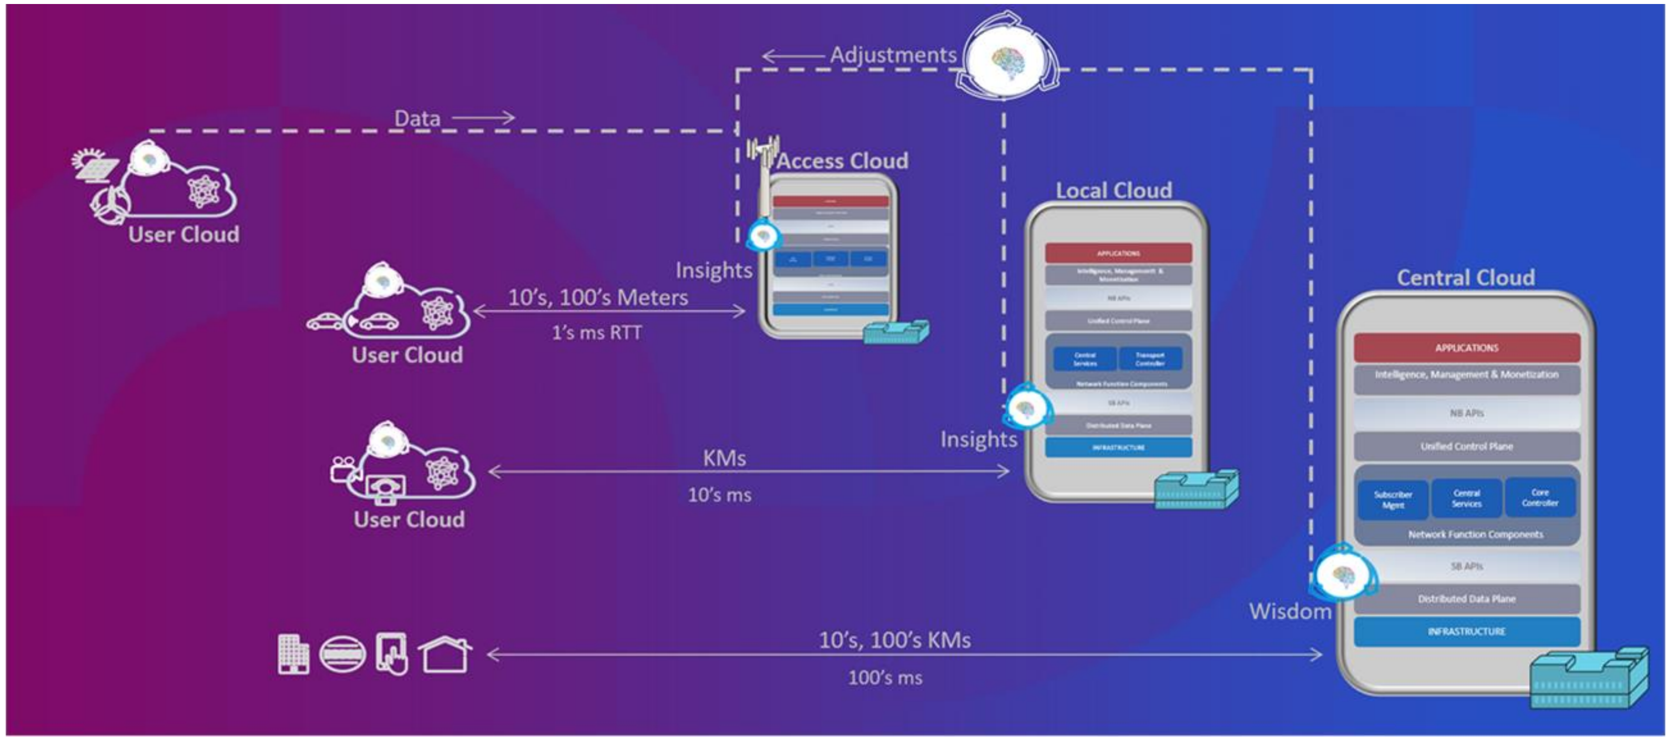

Autonomics underpins the 5G promise

Source: Ciena

However, when you look at the full user landscape that is being promised with 5G, things get extremely complicated. You have basically three different extremes that 5G is expected to serve.

One extreme is the massive bandwidth required by use cases like augmented reality, virtual reality, and real-time gaming.

The second extreme is the massive volumes of highly dense machine-to-machine communications. We're talking about thousands, if not millions, of devices in a very small geographic area.

The third extreme is massive performance. One of my friends, Misha Dohler from King's College, London, talks about internet of skills, which is basically delivering skills remotely, for instance over the Atlantic. You're talking about very high, stringent requirements on the latency and performance.

If, as is the expectation for the 5G system, you have one system delivering these three different, extreme variants of SLAs, and you expect it to dynamically serve these three different variants at any given time, then the monolithic systems of today fail to deliver. That's where things get very, very complicated.

People thought that 5G benefits from SDN and NFV, so it should not be so bad. However, straight implementation of SDN and NFV on the monoliths that we have today – i.e. virtualize them, use automated workflows for life-cycle management, and distribute the control plane and data plane using SDN – even that's not enough.

What is needed is a full separation of control plane and data plane, a full separation of hardware and software, coming to a point that you have almost fungible systems. You have building blocks that can be stitched into dynamic compositions that cater to a specific user flow.

To enable such dynamic compositions, what you need is an extremely intelligent autonomic control system that is able to stitch these flows on a dynamic basis. These flows are expected to cut across disparate networks.

Especially if you look at uRLLC, the ultra-reliable low-latency communications, in the internet of skills example that I just gave you, you have confluence of IoT, cloud, and telecommunications systems. There is a need to join these systems together through some form of a reference framework.

Nothing exists at this point, so there is a need for folks from all of these three or four different market segments to put their heads together to create this reference framework that allows autonomic control of user flows stitched together on a dynamic basis.

Monica: That brings to mind edge computing and network slicing. They should help manage all those different use cases and improve performance as well. What's the relation there with autonomics?

Kaniz: Network slicing, in my opinion, has been overly hyped. By analogy, you can make a huge cake and cut it into different sizes and shapes and different forms, and assign it to the users or the folks who have a liking or a need for one type of slice versus the other. That is network slicing.

That technique will allow multi-tenancy on systems of today. However, it is not able to deliver the dynamic user flows needed for the future systems – enabling, for example, those three extreme use cases I described earlier, around massive broadband, massive volumes, and massive performance.

What you need there is dynamic compositions, and autonomics is the linchpin to dynamic compositions. Again, you can compose things using static policies, but those will be template-driven, static compositions. You will not be able to dynamically serve up user flows that are serving one specific use case, in a particular given time, based on a particular context.

That's what you need for the future systems, and that's what Ciena is striving towards.

As for edge, again, the term "edge" can mean different things to different people. You have the multi-access edge, which is sort of an evolution of ETSI MEC. Then CORD and other communities are building some similar solutions. The telecommunications providers and vendors are marching ahead with the multi-access edge. You have the OpenFog guys, who are looking at it from the cloud perspective. Then you have Amazon and Google, who are also releasing edge products. Autonomics is a common denominator to all of these.

In the end, when you are serving the uRLLC type of use cases – the future autonomous cars, remote surgery, or remote delivery of skills – for all of these, you have to stitch all of the disparate edge systems together.

The way that will be done, in my opinion, would be through a use of autonomics that allows these disparate systems to be governed under their local policies and local jurisdictions. But then you have a coming together of these systems under a common, overarching policy.

That's where autonomics really helps you. It allows you to stitch different jurisdictions together, which allows for proper separation of concerns. You have the local policies governing the local controller, with east/west and north/south peering that allows you to federate across these disparate control systems.

Monica: What do you think are the priorities, in terms of moving to networks that are more automated over the next few years?

Kaniz: I'm not sure I'm ready to say what the priorities are. There are a lot of challenges we have to solve. Once we solve those problems, then we'll be able to say, "These are the things I do first, and then I go to the next step."

As a very first step, we need to solve three key problems to tame the complexity. Telecommunications networks are already heavily complex. Most of the complexity has been built over time, because of all these monoliths which are sitting and controlling themselves in silos. Then on top of that is all the promise of 5G. It's not just 5G; for instance, it's also anything and everything that has to do with self-driving things all around us.

First, you look at autonomous cars, autonomous drones – autonomous everything. How do I cater to this new breed of use cases with telecommunications system software that has been running there for ages?

Key to enabling a useful automation is an ability for an external system to control an entity. The software sitting in the telco system is just not ready, not prepared to be controlled by an external system. That is the hindrance number one that has to be removed.

Today, what you have is element management systems, which are a protector that sits atop the vendor's software and allows you to punch tiny holes, so you're able to have some degree of observability, and then some degree of influence – not control – on the control that is sitting inside the application.

That lid has to be lifted up. We need to remove the EMS, and we need to allow a fluid exchange of observability and control, bi-directionally between the controller and the controlee. That is item number one that needs to be taken care of when enabling telco systems to be fully automated.

Another issue that I see today is centralized network topologies. Systems such as the communications systems, and the CDNs are designed on tree structures. They are very much centralized. You have a strict hierarchy, which is sort of a tree or a fan.

Those types of topologies really start to break down when you think of the distributed systems that will be required of the future networking gear we'll be serving – the autonomous cars of the future, or the smart cities, or automated farms.

Traffic distribution is the second challenge that we need to study, and Ciena is deeply engaged in those studies, along with our partnering service providers.

How will these swarms of devices that are about to be activated in the near future shape the traffic? How will they impact the algorithms that dictate the network topologies?

The third issue is the control hierarchy. It's easy to draw these nested control loops on paper and make it very complicated. I have an extremely complex system that will solve world hunger. However, it's a control theory problem.

How do I create the framework that allows me to operate these control loops without interfering with each other? How do they interact with each other? How do we establish the control hierarchies?

These are three basic problems that Ciena is engaged in solving. We are taking help from our vendor and academia partners, and from the service provider community, to solve these problems.

Hopefully, once these are solved, we will be in good shape to talk about prioritizing the steps one should take to start implementing these systems that will help us move forward toward a fully automated system of tomorrow.

About Ciena

Ciena is a network strategy and technology company with a unique ability to translate best-in-class technology into value through a high-touch consultative business model. We believe in driving an industry ecosystem and offering the greatest degree of flexibility in how our customers consume technology to deliver the most rewarding experiences and business outcomes. We develop and apply technologies across packet and optical networking and distributed software automation that drive openness, virtualization and network automation for customers worldwide to meet surging bandwidth demands and facilitate the shift to the cloud.

About Kaniz Mahdi

Kaniz Mahdi is Vice President of Advanced Architectures at Ciena, and a highly regarded telecom visionary. In this role, Kaniz is responsible for Ciena's global technology strategy and target architecture for 5G era Networks underpinned with Open Architectures, Network Function Virtualization (NFV), Software Defined Networking (SDN), Machine Intelligence and Autonomics. Kaniz spent the last few years as VP/Head of Architecture at Ericsson passionately shaping technology landscape for the multi-faceted transformation of telecommunication industry with Cloud, SDN, and 5G. Prior to joining Ericsson, she headed Communications Services Standards Research at Huawei Technologies, and held various roles in System Architecture and Product Design at Nortel Networks. Kaniz has a stellar record of continuously pushing the envelope on new technologies with more than 60 patents and publications; she is lead inventor of key technologies essential to current Voice over LTE systems and played an instrumental role in adoption of Web RTC. She holds a Bachelors of Engineering degree in Electrical Engineering and a Master of Science in Telecommunications.

Tupl

Tupl was founded in 2014 to unlock the potential that artificial intelligence and automation in wireless have for increasing efficiency and creating new business opportunities. Tupl sees this as a revolutionary transformation in telecoms which could have an impact comparable to that of the industrial revolution.

Today, operators do not effectively use most of the performance and operations data from their wireless networks. Tupl wants to change this and to do this it has created a platform that transfers performance and operations data into a digital format that enables automation and optimization in wireless networks.

Tupl created TuplOS, an AI engine, as a tool for developing automation utilities and creating a digital knowledge database to harness complex processes.

TuplOS system includes utilities to collect data from multiple sources, to create models and to simplify training processes, with a goal to make it easier to automate complex processes.

Automation of manual and repetitive tasks can reduce costs and help operators use existing network resources more efficiently.

In addition to digitizing performance data and automating processes, Tupl's solution continuously learns from the data it captures, to further refine the automation and adapt the processes to changing requirements of increasingly dynamic and complex wireless networks.

Tupl claims that its solution can deliver a 90% reduction in manual labor, a 100-fold increase in speed, and a fourfold increase in accuracy compared to existing, manual engineering processes.

Tupl applications within TuplOS include automated technical customer care (Automated Customer Complaint Resolution [ACCR]), Proactive Care, automatic network troubleshooting (Network Advisor), network construction prediction, and operations workflow management (UNIFIER).

Tupl ACCR brings together network and customer data from multiple sources to find the root cause of customer problems, then provides recommendations to customer service personnel and network engineers. Heuristics and machine learning work jointly to identify the cause of a subscriber complaint.

According to the company, 80% of complaints can be resolved in a closed loop, without human intervention. For the 20% that require human supervision, the time it takes to the resolve the issue is cut in half, because engineers can drill down into the data or escalate the resolution process more effectively.

Tupl Network Advisor uses AI and ML to do root cause analysis and troubleshooting. It automates the data analysis and helps engineers refine their tools for tuning network performance.

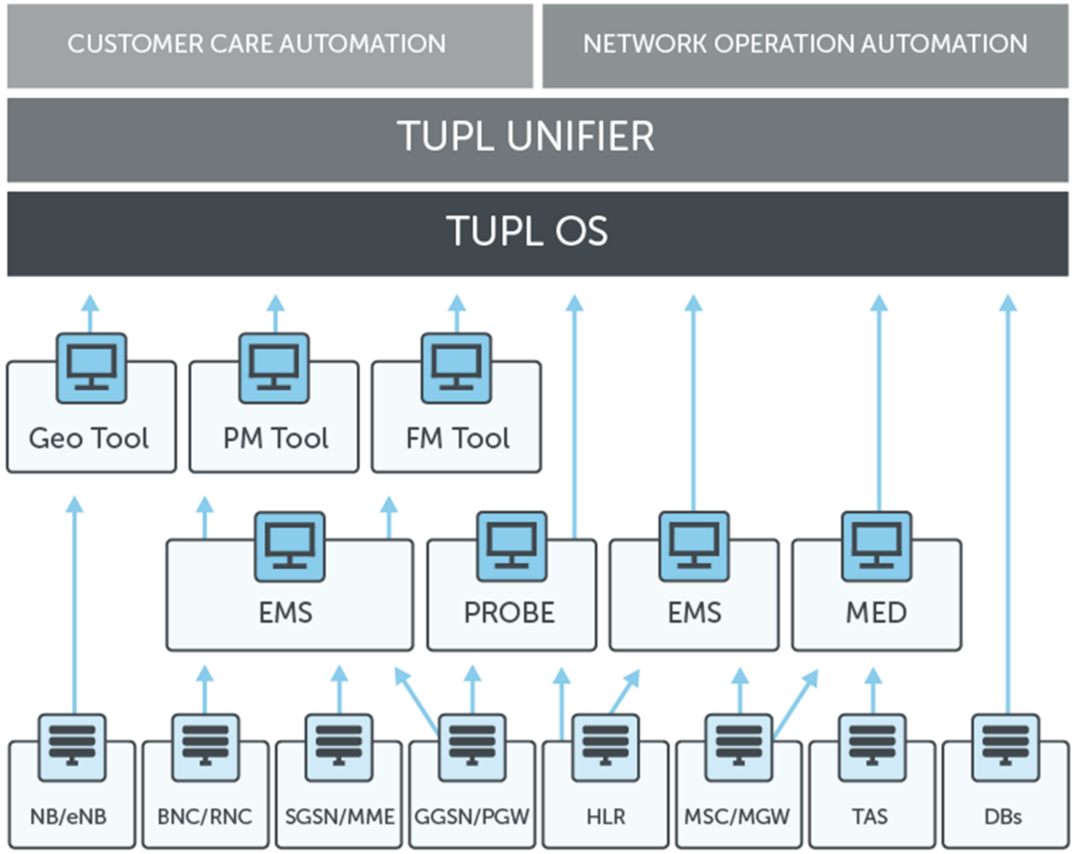

Tupl solution *Source: Tupl*

Tupl UNIFIER is a platform that gives the operator a high-level and detailed view of its network's health and performance. UNIFIER brings together data from multiple sources at multiple granularity levels, including:

- Network status data: network performance KPIs, alerts and anomalies; configuration parameter values, changes and discrepancies; faults, trouble-ticket and work-order events
- Customer data: geo-location records, call performance, service performance.

Tupl

Learn to automate, or automate to learn?

A conversation with Petri Hautakangas, CEO, Tupl and Pablo Tapia, CTO, Tupl

What happens when an operator starts to automate its processes? How does the internal culture change, what are the financial benefits, and how does it affect the services the operator offers? I posed these questions to Petri Hautakangas, the CEO of Tupl, and Pablo Tapia, the CTO of Tupl.

Monica Paolini: Tupl is a new company. Could you tell us what you do?

Petri Hautakangas: We are a relatively young company. We were founded in 2014, with a good group of industry experts. We saw a real need for automation coming up, especially in the most aggressive tier 1 operators.

There weren't any really good solutions out there. That was the premise for founding the company: to start developing a solution that can face the automation tasks when it comes to complex but repetitive manual work that happens day in, day out in network operations.

Monica: How does automation relate to other sister efforts, such as AI, machine learning, or deep learning?

Pablo Tapia: Automation has been going on in the networks for a while. Operators have been trying to deploy solutions to be more efficient – for example, in frequency planning, or cell planning. Some tools are automated in the networks using artificial intelligence, which is a way to put people's thoughts into code.

The real breakthroughs are techniques in machine learning that enable this automation to go one step beyond what was possible before. There are algorithms that don't necessarily require coding. These automation solutions can be created faster, and they can go deeper than real-life coding. They can capture things that are in gray areas.

Machine learning and deep learning are taking over in many industries, not only telecoms – for instance, in automotive or agriculture. It's only natural that it would come to telecom, and it's starting to produce notable results.

Monica: Is there anything special in the wireless industry that's motivating the adoption of these technologies?

Pablo: In the wireless industry, the complexities of the networks are increasing dramatically, not only because of the technologies that are deployed, but also because of the size of the networks, and the number of components deployed, the small cells, and more devices with IoT.

At the same time, customers are becoming increasingly demanding on quality. Mobile network operators need to become much more efficient at operating these systems. The only possible way to do it is through automation.

Monica: Automation is not simple. How can mobile operators get the benefits without too much disruption?

Pablo: This is a long-term journey. It's not a big solution that you buy, deploy, and then you transform your operations from one day to the other. It's a gradual path where, with very little investment, operators can start to get into this transformation, and they can start to see benefits already from day one.

There are a lot of things that can be automated. There are customers that have big targets for automation in over 80% of their processes. This is a journey of automating, little by little, all these processes that are manual today.

In order to do that, you have to have a vision of how you're going to orchestrate all these automation solutions, so that you do not end up creating a new problem that replaces the one you're trying to solve.

Monica: Before we talk about the technology, can you tell us about the cultural changes that go with it? Technology is one part of it, but automating a network means you're going to lose some control. It means that people will be doing different tasks than they do now.

Petri: First off, we don't believe in losing control. It has to be a transparent system. It has to be run by engineers. We try to take very repetitive tasks that are not what the engineers should be doing anyway. They should rise above those tasks. That's one part of the system that one deploys for any use case.

It should start from looking at the existing way of doing things, but doing them automatically, and then adapting them to the operator's existing environment. Then you don't have that kind of cultural problem, at least in terms of changing the way things are done right now. It's just that tasks are done automatically, but the engineers continue to be in control.

The next step could be that you start looking into adding new value, new functionality, that's not possible with just manual processes.

There are, of course, some obstacles also, technical ones and departmental ones, because to create value, you may need to break down some silos that exist in pretty much any operator at the moment.

Pablo: There is an exciting perspective of how this can change the way people are operating the networks today. As Petri mentioned, there is a lot of manual work that goes on day to day.

Control needs to rise to a higher-level role, where engineers, instead of doing the actual job, are more like designers, and they move from being a telecom engineer to a sort of software engineer who handles all these models, the decisions, and the knowledge base that is being created.

It's an interesting change of paradigm. It doesn't have to happen overnight, but it will be interesting to see how it develops.

Monica: How do you envision the adoption path unfolding? Are the operators going to start automating end to end, or will they start with specific functions?

Petri: We are engaging with extremely aggressive operators in Europe, here in North America and in Japan.

Every operator is slightly different, but the most aggressive operators have decided that they will need to reach a high level of automation, otherwise they will not be able to cope when the complexity increases in the network.

Automation tasks start from individual functions. That's what we see. You have small tasks to be automated, and then you have also rather big ones – for example, optimization, or solving technical customer care issues in an automatic way.

Monica: What are the functions that operators want to automate first?

Pablo: It depends on the operator; every operator has its own pain point. It depends also on the phase the operator is in. Some operators are growing, so they need an augmentation of their workforce. Things like designing the network faster or troubleshooting problems quicker are relevant to them.

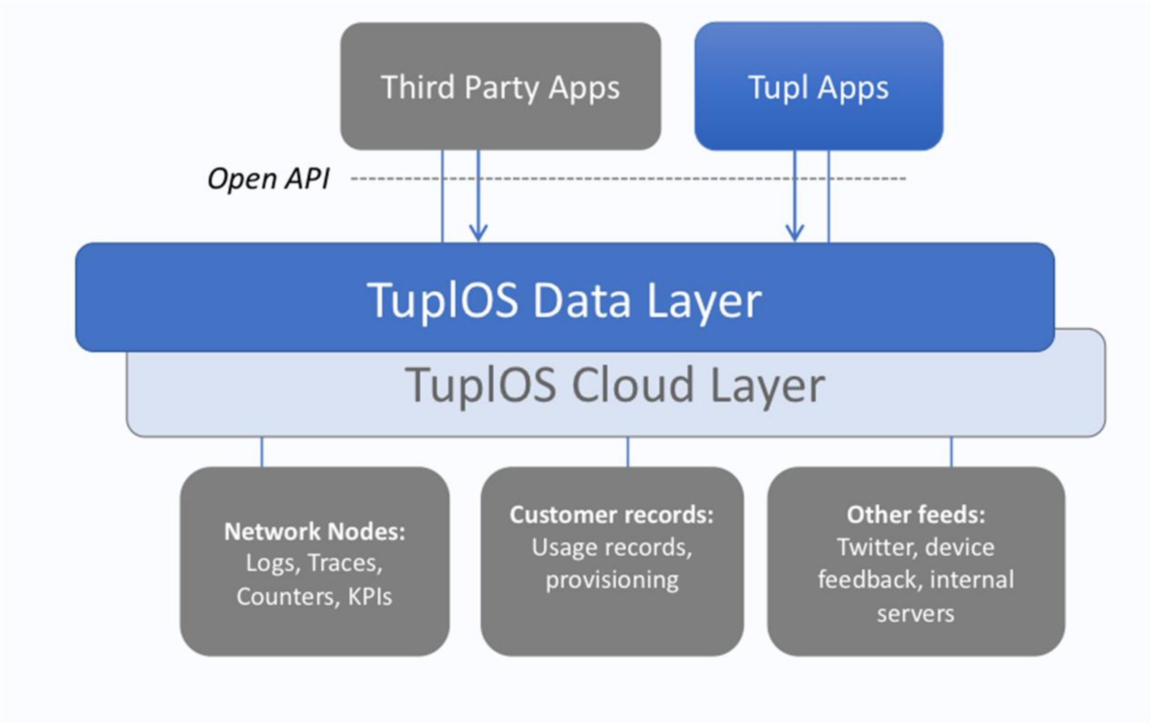

TuplOS AI Engine for automating network operations *Source: Tupl*

Other operators are focused on customer experience, and they want to see their customer care processes improve. Right now, if you look at typical ratings, operators don't get very good reviews from a customer care perspective. There is a lot of room for improvement.

Depending on the operator, depending also on the geography, there is a focus on one point or the other. What is important to understand is that there are many, many things to do. Ultimately, all these things fall into buckets that all operators will end up having to use.

Monica: Petri, you've worked in Europe, Japan, and the US. Have you seen any differences?

Petri: Yes. A surprise, in a way, is that Japan has created an almost countrywide mentality in support of automation. It's a very strong, holistic requirement that makes Japan one of the most aggressive countries in automation. And that flows into the operators, definitely.

The US is very business oriented and customer oriented. Operators want to be efficient and serve their customers the best possible way. If you have too many manual tasks, for example, it takes too much time to respond to customer issues.

If you are able to automate that, just think about how quickly an operator can help the subscriber. There're always some issues, but if you are able to respond quickly to the customers' problems, you are in a different league than other operators.

Monica: Pablo, you mentioned that this is not an expensive addition to the network. Can you expand on that?

Pablo: Because of the technologies that can be leveraged for these solutions, you do not need to develop massive systems. What is important is to focus on the value provided to the customer.

In many instances, it's a very clear-cut business case for the customer, because it provides a significant amount of savings, and you are able to quantify it. You can demonstrate the ROI in a few months. It's quite an easy investment, from the operators' point of view.

Monica: At the same time, you need to go through the learning process at the beginning, that results in a long-term change. Is there a long-term component in the ROI?

Pablo: It depends on the solution you're using. We've put a lot of purpose into making those aspects very simple in our solutions – into how to create and train those models, and into how to make sure maintaining the automation solutions does not impact the bottom line.

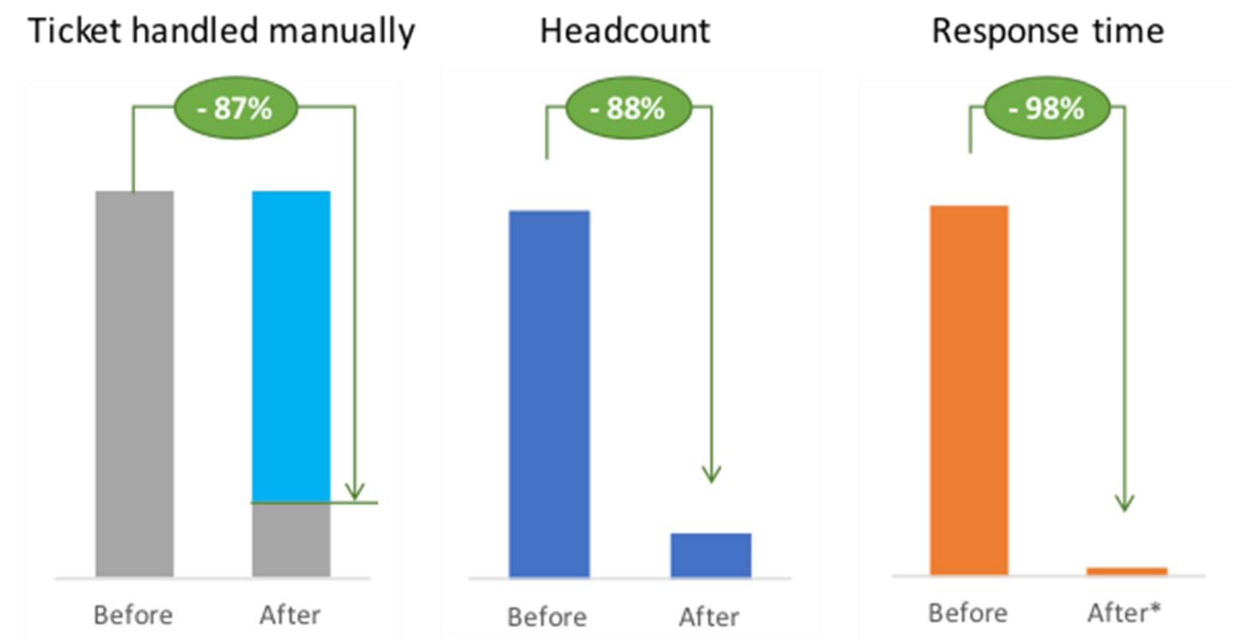

Increased efficiency with the Tupl AI Engine *Source: Tupl*

Petri: If anything, the ROI can be extremely fast. These automation applications are rather complex, but if there is a good underlying system to quickly create the automation model by using the engineers' workflows in a particular customer, you see the benefits really, really fast. We can make a very interesting case on the ROI.

Pablo: One interesting aspect is that these are cumulative savings. It's a change of mentality. When one looks at this problem, you're not trying to solve everything, or to get a big chunk of your opex with the first step.

When you look at the final goal, you see cumulative cost savings that, many times, come from hidden costs and, at the same time, empowerment for doing new things.

Monica: What is Tupl's focus in network automation?

Petri: Our slogan is "operations made simple." At the moment, we deal in automating the technical customer care, even looking into doing that in a proactive manner, and we are automating network engineering and optimization tasks.

We have a second product called Network Advisor, which just recently won a Fierce Innovation Award in the network optimization category.

In the network operations center, we do task automation, alarm handling, and so forth. There's lots to be done in a more intelligent way in that area, also.

Those are the three buckets we do.

I can mention briefly what is technical customer care automation. About 20% to 30% of the overall customer care issues are related to something technical.

The technical issues end up landing with the engineering and operations people. They spend a huge amount of time going issue by issue, looking at six different tools, and trying to, in their head, correlate what's the root cause, and how to fix it and how to communicate back to the customers. It takes a long time. It also consumes a lot of the engineers' time. It's not necessarily consistent, because you have a group of 100 or 200 engineers doing it.

We set out to automate this. This is a good example, because in a particular customer case, we started correlating a certain amount of feeds, a task that the engineers typically did. The results were immediately valuable, using the engineers' expertise to figure out the right route, and what to communicate. We've trained that model with the engineers.

Then you start to add more information into the system, so you get more granularity and more accuracy to really pinpoint the issue. That's how you add new value to an existing process for which it was impossible, in a way, for the engineers to have all the data in hand. The savings are significant, and the customers are really happy. Our clients are able to resolve customer issues 100 times faster by automating up to 90% of the engineering tasks that had been done manually in the past. Plus, they are getting results that are four times more accurate, and 100% consistent.

Monica: The system keeps learning. As you say, at the beginning you have the initial learning phase, but then as you go on, it perfects itself, so you should be seeing better results, right?

Petri: Yes, but perfecting is maybe not the right word, because the situation also changes. Let's say you build a perfect system. It will be a static environment. We all know there's nothing static about mobile networks.

New technologies and changes in the network are introduced. The network has to adapt in a scalable manner, in a maintainable manner, and the only way to have this is by AI learning.

Pablo: It's not like it's going to be always getting better, though. There's a point at which it's good enough, because even engineers don't agree on the results. What is important is to have the ability to provide this feedback mechanism to continue adapting and re-teaching the system.

There is also an interesting question of how you go beyond the first use case. There are two ways to go.

One is to continue going deeper into that use case. For example, in the case that Petri mentioned, the customer care resolution, now we are going into the second phase. Instead of reacting to a customer complaint, we are going to proactively identify and resolve the issues customers complained about.

Petri: I call it the holy grail of customer care.

Pablo: At the same time, once you have incorporated all these different data feeds and data correlations, you can build other use cases on top of what we call our TupIOS platform. That's an AI correlation engine that allows you to build more value, and then realize this dream of creating multiple automation solutions that solve the operator's challenges.

Petri: It's a journey.

Monica: Once you understand what's going on, and what the causes are, you might be able to preempt them. You might see the problem coming, and also, you might generate new services that address the concerns users have.

Petri: Exactly. That's reactive customer care moving towards proactive customer care. You validate your model and the underlying algorithm in the reactive mode. And then you start to have the pro-activeness, so even before there is a physical call, you are able to use the same process to take action proactively.

Pablo: Not everybody calls to complain.

Monica: Subscribers may churn without calling. You want to catch them before that. There is a huge amount of value if you can do that. Operators do not know today why subscribers churn. This is a way to address that, because the operator can go back and see what happened before a subscriber churned.

Once an operator adopts your solution for customer care, it has a platform that can work in other domains as well. The operator is learning a lot about its network, which goes beyond customer care, because, as you said, it's a technical problem.

How easy is it for them to then move to new areas of automation with your solution?

Petri: We have always had a vision of a multitude of use cases developed on our AI engine. From the beginning, we've designed our use cases with consideration for this expansion.

Once you have integrated the first use case, which depends on the first pain point that the operator has, then adding the next one is a simple effort from an integration point of view. You just focus on the value creation through models, training, and so on.

To give you an example, with customers, I always start with the customer care solution. Now we are going to a second use case, which is the network optimization and troubleshooting, which we call Network Advisor. It uses many of the same feeds, but it solves a different problem, which is the problem of the RF engineers in the field.

Monica: What's next at Tupl?

Petri: We've set good foundations, especially with the most aggressive operators. We now see the next wave of operators requesting the same solutions. It's continuing with the journey.

We already have requests for IoT, but we see it almost as a service that has to be managed separately, although it is in the same network.

These things take additional paths. Those need to be properly managed and automated as much as possible, because again, networks are not getting any simpler anytime soon. We plan to support all these future paths going forward.

What we have designed is completely applicable to any network operation. We already see the complexities increasing in verticals such as utilities. It's pretty easy for them, to be honest, for applications such as electricity metering. It's a point-to-point connection, and every now and then, a couple of bits come.

Now, with smart meters and so forth, you have many other networks where the same process applies. To be honest, you can't find a more complex place to be than in mobile networks. We've proven ourselves by solving even the most complex tasks in wireless networks by automation.

About Tupl

TUPL

Founded in 2014 by telecom, big data and AI veterans, Tupl is transforming customer access and experience in the telecom industry through improving operations with leading wireless operators across the US, Canada, Japan, Mexico, and Europe. Its AI Engine, TuplOS, utilizes machine learning and several other utilities to enable faster innovation cycles for network and customer care operations. Tupl is headquartered in the US in Bellevue, Washington with presence in Spain, Mexico and Japan, and is continuing its rapid global expansion in 2018. To learn more about Tupl and request a demo, visit www.tupl.com.

About Petri Hautakangas

Petri Hautakangas is the CEO at Tupl. Petri is well known in the telecom industry, having worked in senior positions in Japan, USA, China, Mexico, Brazil and Finland. In Japan, his most recent position was Nokia CTO for Japan & APAC markets. During an earlier stint in Japan, Petri was instrumental in growing Nokia's network business six-fold, reaching well over 1BUSD in sales. Petri also worked in the Seattle area as CTO from Nokia to T-Mobile, orchestrating new mobile network solutions that millions of T-Mobile users are now enjoying. Prior to Japan and USA, Petri worked in Brazil to help build business from near zero to the biggest Nokia operation in Latin America.

About Pablo Tapia

Pablo Tapia is the CTO and Founder at Tupl. Pablo founded Tupl after a 15 year career in wireless, with multiple roles across various areas and disciplines. Prior to Tupl, Pablo was a Member of the Engineering Staff at T-Mobile USA, where he led major technology projects, including Self-Organizing Networks (SON), Quality of Service (QoS) techniques, as well as spearheading the evolution to HSPA+ and LTE. Before T-Mobile, Pablo led a development team at Optimi, a software startup developing optimization solutions, and held various positions at Nokia Networks R&D. He is an author and contributor to several books on Wireless technologies. Pablo holds an MSc EE from the University of Malaga, Spain.

Luxoft

Headquartered in Switzerland and with more than 13,000 employees worldwide, Luxoft is a global IT provider of technology solutions that serves various lines of business within communications, media, financial services, automotive, retail, healthcare & life sciences sectors.

Their approach to media and communications is rooted in governing the ever-increasing complexity of industry trends, with customer data, video content, personalization, customization and automation as the drivers behind chatbot, 5G, blockchain, machine learning, and SDN and NFV technologies.

Luxoft provides solutions with an industry focus to improve customer experience, network transformation, adoption of service enablement and increasing operational efficiencies.

They aim to supply service providers with intuitive tools and solutions in automation and virtualization, AI and machine learning, predictive and prescriptive analytics, OTT Solutions, WAN infrastructure and Ethernet switching, as well as network optimization.

Luxoft's Software Defined Lab (SDL) is designed to help service providers virtualize their networks while lowering costs, reducing risk, and shortening time-to-market for new service launch. According to Luxoft, the simplification and automation of the testing process reduces the setup and testing time from weeks to hours, and improve the

accuracy of testing across multiple vendors, hardware elements and Network Function Virtualization Infrastructure (NFVI) solutions.

The virtualized testing environment supports multiple functions within service provider organizations:

- Proof-of-concept and network trials
- Product evaluation and acceptance
- Network, service chain and solutions testing
- Hardware design support for SDN and NFV
- Platform and infrastructure validation

Luxoft's NEM services include:

- Regression
- Engineering validation
- Full release testing suites
- Multiple deployment configuration simulations
- Rapid fault reproduction and fix validation
- Stress testing for infrastructure, network, dependent services, bandwidth and failure scenarios

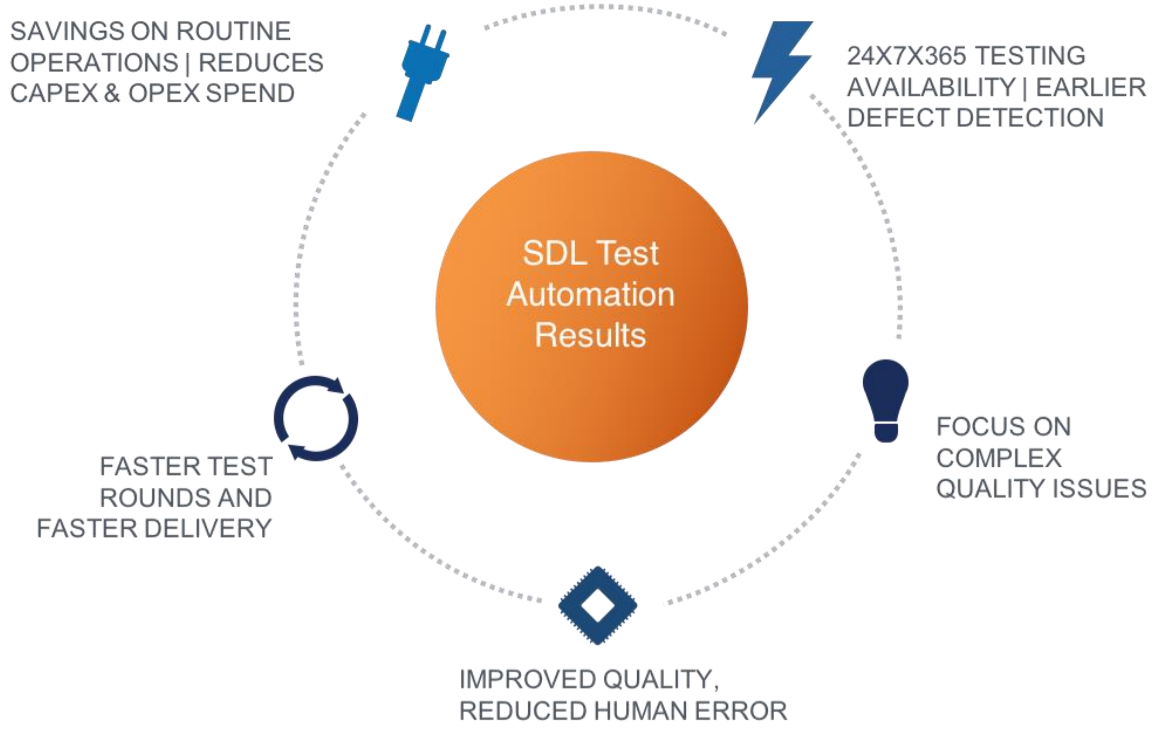

Benefits of SDL

Source: Luxoft

The company also provides communication solutions for their adjacent verticals. In agriculture, for instance, Luxoft has worked with an agricultural equipment manufacturer to develop mobile applications helping farmers to optimize, in real time, tasks such as tilling, planting, spraying and harvesting. Luxoft developed cloud technology to collect and analyze third-party data (e.g., weather) and historical data, these apps connect wirelessly via a tablet device to the farmer's equipment, providing real-time decision support recommendations for adjusting settings during the operation in order to maximize crop yield.

Luxoft

Test automation in virtualized networks

A conversation with James Hopson, Product Manager of Telecom Solutions, Luxoft

Testing is an area where automation is crucial as we move to virtualized, complex and dynamic networks, where more elements and vendors have to interoperate, and operators have to ensure everything works as expected. I have the pleasure of talking to James Hopson, the Product Manager of Telecom Solutions at Luxoft about how they help operators meet the challenges of testing in virtualized networks.

Monica Paolini: James, what do you do at Luxoft in the field of automation?

James Hopson: Luxoft has a test automation product we built for NFV. It started with an SDN product and a test automation product in the legacy space, combined those together, and developed that out for NFV test solutions. We have hundreds of people who work in test automation within the NFV space using our testing tool Software Defined Lab (SDL), and we utilize it to help our customers develop their test environment.

We also work with third-party products that we've integrated to that tool to provide test traffic generation.

Monica: What pain points are you are trying to address with automation?

James: We found that a lot of our customers were spending too much time testing. There are a lot of tools on the market that can handle the automation of test cases, but the manual environment setup of the NFV space and networking, the configuration of both systems and the various testing tools were taking most of their time. Sometimes customers would spend 90% of their time in testing on this one phase.

So, we focused on developing a set of tools that could automate this process (including setup and reconfiguration), while ensuring it would work with any of our customers pre-existing equipment or software.

Monica: What limits automation today, and what will change with virtualization?

James: Until today, the automation has been mainly about writing scripts to execute tests against a fixed test bed. The setup of the test bed was typically manual. The setup of the equipment that you were testing was typically manual, as well, or preconfigured by the vendor.

In the NFV space, the whole environment has become fluid and flexible. The vendor component that we're trying to test has to be configured, and it usually can be configured in different ways in the NFV environment.

The testing tools in the NFV environment have a lot of configuration options, a lot of capabilities, so most of our testing is trying to determine what is the best option for those various components, as well as working out the best orchestration within a fluid environment.

The challenge is the limitless number and types of configurations possible, coupled with the speed at which new hardware hits the market and software iterations are released.

Monica: There is acceleration in the field of automation. You can automate more, and automate faster, and the scope is wider. Automation requires some work initially, because you need to change the way you do things. But it's also an investment for an operator. What are the drivers toward automation?

James: The main driver is cutting down development time down in building the test environment. Generally, our customers have already invested in the test case automation, going from test case definition to generating the appropriate traffic, and making the appropriate measurements from a test system.

What they have trouble with is automating that network setup. We see customers who will spend six to eight weeks to set up a test bed for a two-day test. Within the NFV space, we can now make the test bed itself virtualized – build it in the NFV space and set it up in an automated way that is just as fast as running the tests themselves.

Because this is a new area, customers don't have any automation for that aspect. They have automation for small test cases, but not for testing the configuration of a physical or virtual network or a combination of both. Consistency is needed to carry out tests across multiple environments and configurations to reduce time and cost-to-market and to remove human error, but there are limited internal skills and resources to manage this.

Monica: Do operators see time savings as a way to cut costs, or to improve performance, or to reduce downtime? What are the immediate drivers for automation?

James: There're a couple of drivers. One is that, because there are so many NFV components coming to market, they need to evaluate more products and more combinations than in the past. They have the same number of people to do testing, but now they have more test load on their hands to deal with.

Also, because the environment changes so often – with security updates, new versions of infrastructures, new versions of NFV – they have to retest their existing components much more frequently than in the past. These teams are struggling with how to get faster.

It's not really faster tests on the same product. It's more testing, more iterations of those tests, with the same number of people in the same lab, the same environment: "How do I accelerate that? If I were just trying to make each test shorter, then yeah, I would be able to reduce staff."

What they're finding is a backlog of additional testing they need to clear out. And by reducing the setup time, it helps them cycle test projects quickly, so they can get on to their next project and increase the pipeline flow.

Monica: And because it's a dynamic environment, you need to keep testing the system, because it keeps changing as well.

James: That's exactly right. Not just the environment, but the systems you're bringing in.

It used to be that you had a physical appliance with 10 or 15 functions built into it. You would take a few months to test that, integrate it into your network. You would turn those functions on and off, but it was a fixed appliance that had a fixed set of stacks and interfaces. You knew how that was going to work for several years, as long as that appliance was in place in your network.

Now, with NFV, we have VNFs to come into play. Having configurable and variable flavors of each component enables scalability and more efficient use of network resources, but creates multiple configurations, all of which need to be tested.

Now that single piece of equipment in your network has become 10 times two or three vendors, maybe 30 different pieces of software you need to test and integrate, plus the combinations of those components, to get to a capability in the NFV space that is equivalent to what you had with

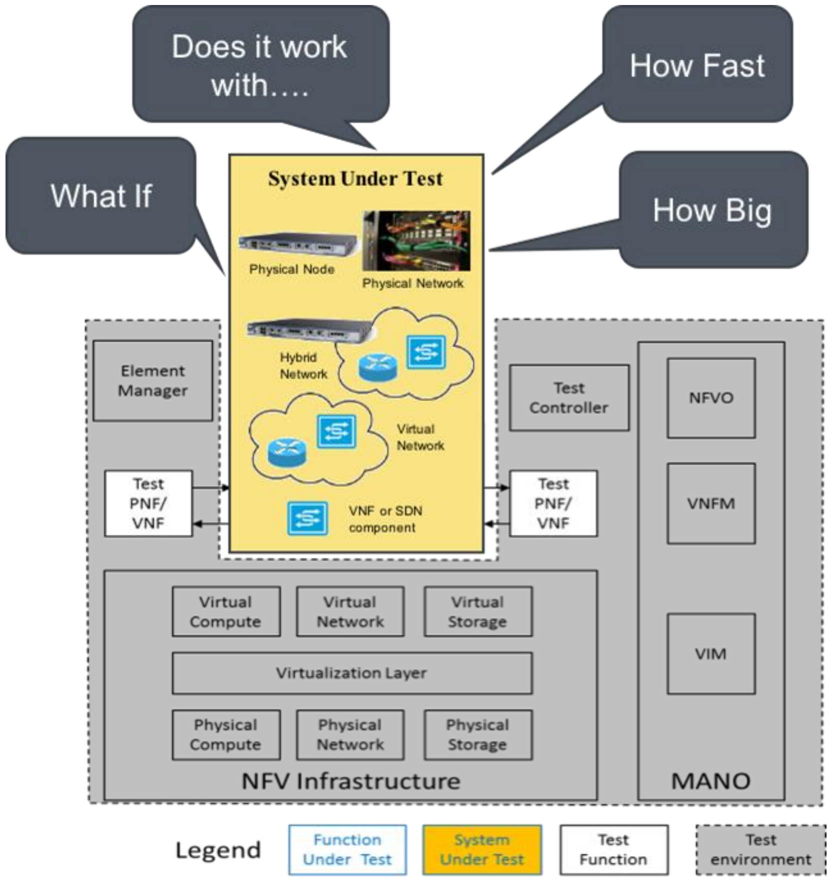

Dimensioning in Luxoft's SDL *Source: Luxoft*

the fixed appliance.

A lot of teams struggle with the fact that there are too many variables. But as soon as you get the automation in place, you figure out how to test component A with component B, and then you automate that. I can switch out the vendors. I can switch out the combinations very quickly without having to go back and spend that six-week setup time.

Monica: The amount of testing you have to do doesn't grow linearly with complexity. It grows at a faster rate than the number of test targets, because you have all the interactions among test targets to test separately.

James: Right. For a given function, we're running about the same number of tasks. It's the fact that we're running more functions, testing more vendors and more combinations of functions than we used to.

Monica: Often, there is a worry that with automation, people will lose their jobs. But it doesn't seem like we are going in that direction, because there's so much more work to do than in the past. Automation will help you do a larger amount of work with the same number of people, but because there's more work to, it seems unlikely that operators will need fewer people – at least initially.

James: That's exactly what we're seeing. The acceleration in the number of components coming into the network, the agility of the network, the number of changes that we have to process, that's where the real pressure is. Automation is just helping the staff to keep up with that, rather than displacing positions.

I guess if you had one company testing one product, one VNF, it would be more efficient and be able to reduce some staff. But that's not what the operators are seeing. They're seeing more VNFs, more components coming in, more vendors, and the need to either build up those teams or build up their efficiency, and the efficiency's really the best way to do it.

Monica: Before, you mentioned that there are going to be more vendors, so there is an evolution

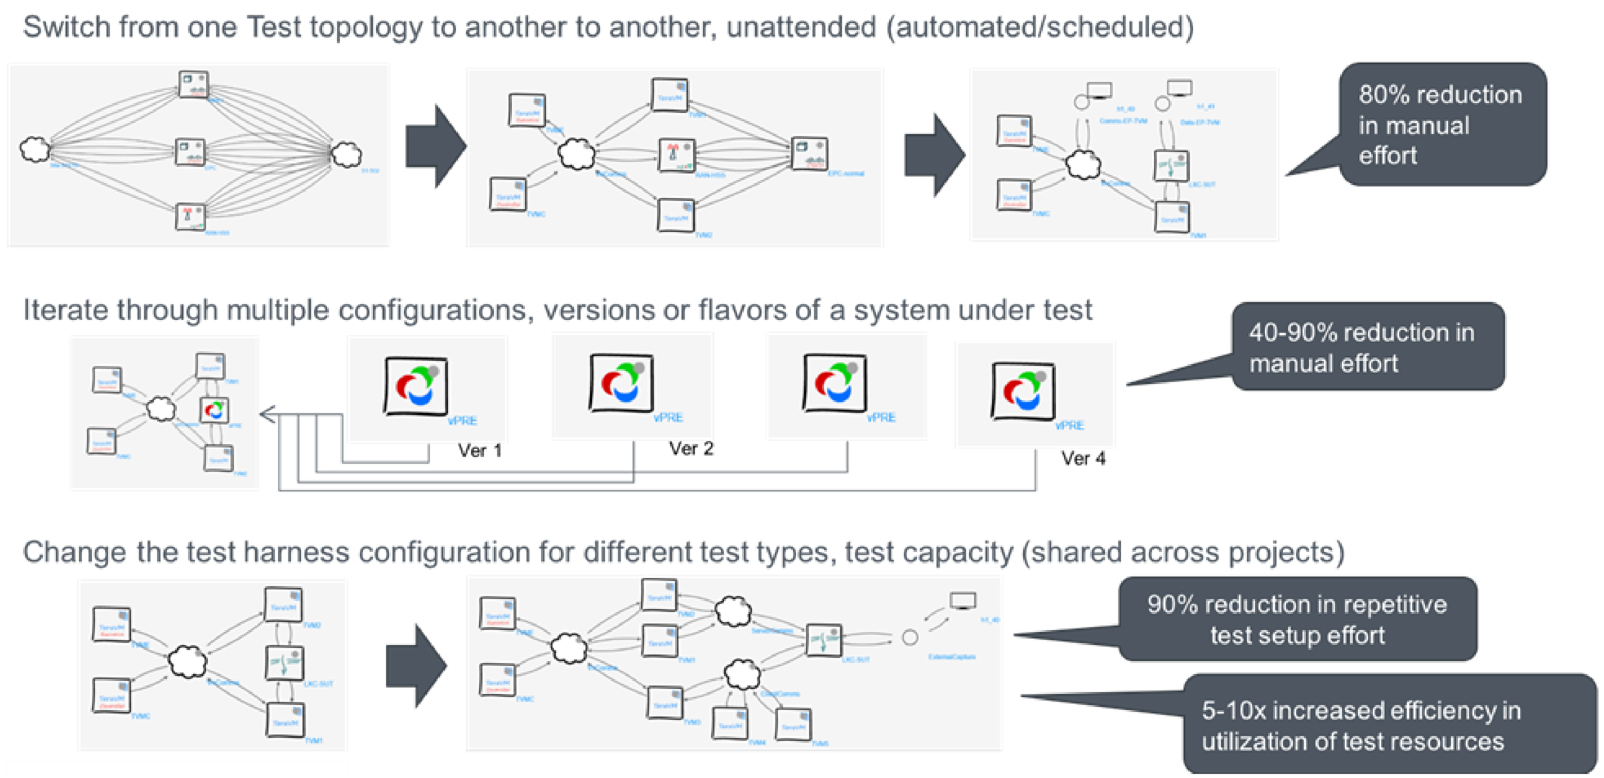

Reduction in setup and testing time

Source: Luxoft

in the ecosystem of the individual operator. How do you see that changing? What are the drivers? How can an operator cope with that?

James: When you were dealing with appliances, you wanted to deal with a small number of suppliers. You wanted to deal with a piece of equipment you knew about. There was a lot of capital investment, so there was a lot of long-term commitment.

As we move into the software space, if your vendor doesn't meet your needs, changing out, let's say, one firewall for another with our SDL in play, it's simply a matter of some testing, with some configuration changes, as long as the functionality of that component is the same.

The ability in the software space to make decisions on smaller components, to make decisions faster, to have a shorter mandatory lifetime in the network gives the networks and the operators more flexibility in the services they can offer. It can cut their costs, because now they can go to smaller vendors or more focused technologies.

To some extent, it provides the customer with more choices as well, because now instead of having only one or two types of, let's say a network edge device, the operator may select and configure them from a palette of choices, as long as they've been certified to work in that network.

Monica: With more vendors, operators have a bigger responsibility in ensuring that all the different elements work together. Is that in itself changing the way they approach testing?

James: Certainly, it changes the way things have to be done. The traditional six-month product acceptance process, and then six-month engineering rollout process, is really too cumbersome when you're trying to roll out these smaller VNFs or trying to turn them around in a month to get those out to your end customers.

The expectation to have more agility in modern networks is continuing to intensify. It's a follow-on to the iPhone generation, where if I want a function, I can go out in a matter of seconds to shop the app network and find various applications that meet my needs, pick the one I want, and if I don't like it, I can download another one within a few minutes.

That's the expectation and freedom our mobile devices have given us. And we see, certainly on the customer-premises side, trying to push for that same capability from operators. Our customers want to be able to come in and say, "I don't like the way that firewall's working. I don't like the way that mail proxy's working. Let's try a different one. Here's one I found on the internet. Can you put that in for me?" This eliminates weeks of ramp-up and testing to determine if it's a viable option.

That agility, that flexibility in our customers, that quick decision, quick connectivity is leading from the mobile side of the network into the access.

Monica: With automation, operators lose the direct control they are accustomed to. If they decide to tune up the network, they may no longer be able to do it manually. In this perspective, automation is more than a technological change. It requires a cultural change as well.

At Luxoft you provide the technology, but when you talk to operators, do you see the challenges that the cultural change brings?

James: Yeah, I actually think that's the biggest challenge, taking an organization that has been in a test and engineering environment for some number of decades. They've gone through several evolutions of the hardware appliance side, from proprietary hardware into parts hardware, and even into virtual machines running on parts hardware.

But this next step, this step into individual functions being pluggable into the network, is a real cultural change that requires a top-to-bottom rethinking about how that whole process works.

What does the pipeline look like from the author of the software all the way through to the customer? How do I shorten that pipeline? How do I ensure integrity? And, especially, given as flexible as some of these networks have become, how do I know that that piece of software, which worked well in network A, is going to work well in network B? These are some of the challenges our customers face.

We believe the way to address these challenges is through the ability to turn up and run the whole suite of validation tests against a function anywhere in the network and anywhere within the life cycle, while teaching the operators how to structure their organizations to continually make this happen.

Monica: Do you see any differences across different geographies?

James: Not as much as you would think. We see that the big operators in all the major regions are starting to make large investments in the NFV space. They are beginning to roll out services. Their teams are all in the same stage of a cultural revolution internally about how they're going to deal with this pipeline.

A lot of times, operators think of their deployment teams as the primary key to deploying a new service. But in this case, it involves the entire supply chain validation – the Lab, Deployment, as well as Operations. They're all involved in figuring out how to deploy this new technology.

But we see that occurring about the same amount in all regions. The big operators are making that move. The smaller operators are mostly followers. Some of them are being progressive and jumping into this feet first, and ready to move forward. Others are holding back a bit, trying to see what the best tools and techniques are going to be, who the best service providers are going to be.

We've seen a lot of the smaller operators look at outsourcing engineering, or outsourcing some of the testing capabilities, just because it's so daunting to them to try to take up this role.

Monica: Do you need real-time testing in a network that is fundamentally dynamic and, hence, changes all the time?

James: We certainly see the need for continuous testing throughout the life cycle, all the way out to customer deployment, especially with the flexibility we're introducing into the network. The customer may say, "I want you to configure these three components, these three functions for me," and the operator may never have put those three together before.

What they don't want to see happen is having to spin all the way back to their procurement side and start the process over in order to find a different vendor to test out their requirements.

But think about it – if these three components requested by the customer could be tested to validate end-to-end traffic and end-to-end compatibility in a matter of a few minutes, then be able to pass that over to their customer – it's an attractive capability.

That's why it's going to be important to build the test tools along with the NFV functions themselves and carry them through that life cycle.

Another thing we is that our customers also want to use this as part of their SLA. If one of their end customers calls in and says, "I'm not seeing the traffic flow at the rate I expected or the rate I've ordered," they can run a traffic simulation, spin that up, and watch that traffic flow through the component. They can find out where that bottleneck is, and see if something's configured wrong, or see if, perhaps, it's running fine up to the network edge, and it's actually the customer's equipment that's causing the issue.

Being able to simulate that traffic in the live network at real time is a valuable capability that is needed in the NFV space. We can turn up those test generators right inside the live network.

Monica: In that case, the testing becomes closer to monitoring?

James: Yes and no. Monitoring is something you would do passively. The testing that I'm talking about would be more active, with various traffic types generated at whatever edge of that chain of services the customer needs us to test.

If the customer has an issue, we can say, "Let me handle this now. I'm going to block your network for five minutes while I run a test." Or maybe, "I'll do it in off hours to validate that that traffic is actually flowing properly."

You could certainly also do it at a lower rate. If you're seeing a protocol problem, where they're saying, "All my YouTube videos are coming through, but my Google videos aren't," or "my Amazon videos aren't," then you would try maybe a protocol test, without consuming lots of bandwidth.

Monica: What's special about what Luxoft does?

James: We are vendor agnostic, plain and simple; our solution works with various software and hardware vendors. We've built standards-based test suites for product evaluation across multiple vendor products/versions.

We have many environments where we'll have two or three different, essentially competing NFV infrastructures under evaluation. We can write one test case and test the same function in all three environments simultaneously.

The same thing with the test tools. A lot of tools generate traffic and can simulate network protocols. We don't care which tool it is in the network. We'll talk to any of those; the customer has already made an investment in licensing and in appliances to generate traffic. We plug into those and automate those without having to be tied to one vendor.

We haven't seen many companies willing to jump into the market completely independently, like we did, but we needed to do so ourselves, internally. We have many different tools, customers, and platforms inside of Luxoft that we use for our projects, and we needed that tool ourselves, so we built SDL.

Our customers find this is very attractive. Most of them haven't completed their decision process. They still have options they're looking at. They still have infrastructure variables to contend with, or they're rolling out new versions of infrastructure and new versions of tools, and they want to compare old versus new.

In fact, probably half of our market right now is evaluating off the appliance-based function with our test tools, and then comparing that to the virtual equivalent to make sure they're seeing an apples-to-apples cut-over as they move their customers out of a physical environment.

Monica: There is a lot of value to being independent, because traditionally, a lot of the testing would be done with the vendor. If you don't have that many vendors, it's easier to do that. As the number of vendors increases, so does the value of an independent testing solution.

James: That's right. On the testing side, it's independence from the infrastructure and independence from the test traffic generator, because we find that some traffic generators are better in one domain than in another. Some of them are very specialized and lead in certain areas of the network functionality.

An operator may have had to invest in two, three, or four different traffic generation tools to meet the needs of his network. Likewise, on the infrastructure, they may have some of their infrastructure open sourced. They may have some they've built in house.

About 30% of our customers have a proprietary version of some aspect of their infrastructure in house. Because we're not tied to any particular vendor, we can go in and make a couple of tweaks to our communication modules, and we're off to working with their proprietary environment.

That flexibility on both sides allows us to come into an existing environment and fill the gap, which is that automation and simplification of that setup and repetitive testing, without having to disrupt all the tools and replace them.

Monica: As we move forward, what is our next big challenge? What specific challenges are you trying to address at Luxoft?

James: I see in the future that the culture change in the testing environment and the tools around that are still going to be a challenge. It's going to be, probably, three to five years before the culture and tool upgrade in the NFV environment are complete.

We've talked today about a lot of areas to grow into, on the acceptance testing side, on the engineering side, even out to the customer – giving the customer the ability, for example, to say, "I think my Amazon videos are nor playing at the normal rate," or "My Skype calls are having trouble. Let me run a quick SLA test using virtualized tools that will spin up in the operator, or maybe even on my customer premises, as well. Run that test. Give me the results in real time so I can see whether my network is performing properly or not." I see that as the future.

I also see us continuing to move toward virtualized services on the customer premises. In other words, a CPE box that itself is a mini-NFV infrastructure, so now we're essentially deploying remote applications, getting closer and closer to that iPhone model of "Let me see. I think I want a different security gateway. Let me download this one and apply it to my customer premises equipment." We see that as a trend, as well.

These abilities to validate the service change, the end-to-end traffic flow, the bandwidth of that capability are all going to be tools that we will need.

About Luxoft

Luxoft (NYSE:LXFT) is a global IT service provider of innovative technology solutions that delivers measurable business outcomes to multinational companies. Its offerings encompass strategic consulting, custom software development services, and digital solution engineering. Luxoft enables companies to compete by leveraging its multi-industry expertise in the financial services, automotive, communications, and healthcare and life sciences sectors. Its managed delivery model is underpinned by a highly-educated workforce, allowing the company to continuously innovate upwards on the technology stack to meet evolving digital challenges.

About James Hopson

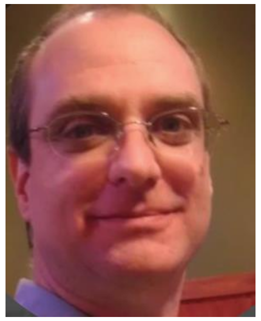

James Hopson has been a pillar within telecommunications for more than 20 years as a Product Manager. His area of expertise ranges from Network Operations, Solutions Architect, and Virtualization; most recently shifting his focus to specialize in NFV, Testing and Analytics. He is currently responsible for the Twister test automation framework and the Software Defined Lab solution, Luxoft's automated test and modeling environment for complex VN based network deployments.

Huawei

"Automation is necessary for future network" is Huawei's response to the increasing complexity in wireless networks that is driving costs up without a corresponding increase in revenues. HetNet topologies and network slicing are two examples of this trend toward greater complexity – a trend that also multiplies the opportunities to expand subscriber services. To manage complexity and costs, the wireless industry needs to optimize the use of resources more effectively.

To do so, Huawei promotes the introduction of Wireless Intelligence across the network, to help operators find ways to improve performance and automate network operation in real time.

A massive increase in data traffic, coupled with an equally impressive increase in computing power, has created an environment in which Wireless Intelligence can thrive, building upon AI, machine learning, deep learning, and reinforcement learning. Moving ahead of the rule-based expert systems of the near past, we have now entered a machine learning phase that uses network data more extensively and efficiently. Huawei expects a deep learning phase and a cognitive phase to follow, centered around the creation of logic.

Wireless Intelligence adds the following in wireless networks:

- Unified data management
- Analytics
- Automation and closed-loop operations in planning, deployment, monitoring, and service assurance
- The ability to evolve through the use of cloud-based applications and open APIs

Increased network complexity is a driver for Wireless Intelligence *Source: Huawei*

Rather than being confined to narrow domains, the impact of Wireless Intelligence will spread horizontally end-to-end across wireless networks, as a comprehensive tool box that operators can use to address challenges wherever they arise.

Huawei's trials and initial commercial deployments give insight into the contribution that Wireless Intelligence makes. For example:

- **LampSite topology optimization.** Topology and RF planning in diverse scenarios, taking into account user distribution analysis and traffic distribution, to increase capacity and reduce interference
- **Adaptive KPI anomaly detection.** Using historical data and fixed thresholds based on expert experience to detect KPI issues and define a more sensitive adaptive threshold using machine learning predictions, resulting in a 90%-plus reduction in fault alarms
- **Motorway traffic recognition and optimization.** Use of traffic patterns to estimate traffic demand and optimize RAN configuration (e.g., antenna tilt) to reduce congestion and call drops
- **Massive MIMO pattern optimization.** Use of machine learning modeling to define and set antenna parameters
- **Bottom user throughput optimization.** Use of machine learning for root cause analysis to understand and resolve subscriber issues, such as dropped calls

Huawei

Wireless Intelligence enables automated networks

A conversation with Kevin Xu, Head of Marketing Operations, Huawei Wireless Solutions

The networks of the future will be automated, but how will we get there? Today's networks are only marginally automated, and network optimization still mostly consists of painstaking manual processes. What is the path that will enable us to leverage artificial intelligence, machine learning, and analytics, and evolve to the automated networks that 5G requires?

I talked with Kevin Xu about the role of what Huawei calls Wireless Intelligence in setting the path and enabling network automation. Kevin Xu is the Head of Marketing Operations at Huawei Wireless Solutions, which is part of R&D at Huawei.

Monica Paolini: Kevin, thanks for sharing your thoughts with us today. I know you have worked in the wireless domain for years, and now your focus is on network automation. What is the high-level approach to automation at Huawei?

Kevin: This class of technology doesn't have a standard vocabulary yet.

First, we consider artificial intelligence a GPT, a general-purpose technology that can be applied to all the suitable cases or scenarios.

Second, we do believe that this kind of technology will create real value for the telecom industry. That's why we have used artificial intelligence to develop research and many solutions.

Third, on a higher level, the entire wireless industry is just beginning to realize how to approach, and how to deploy, artificial intelligence technology.

That's why we think that at this early stage for the telecom industry, we need to focus on the real value that artificial intelligence creates for operators.

Monica: Traditionally, we have a lot of AI in other fields as well. It's not a telecom technology. What is AI's role in the wireless industry?

Kevin: I want to start with the data. When we talk about artificial intelligence technology, it all comes from big data and data analysis. It is a precondition for the outcome. In wireless networks, a tremendous amount of data is generated every day, every second.

To use this technology, operators start with a good advantage, because all of them have plenty of data and enough computing resources, and they can find a lot of value through this technology.

Today, wireless networks create a lot of operations and maintenance challenges, due to the networks' complexity.

From 2G to 5G, there are many bands, many frequencies, and many scenarios: the macro-cell site, the small-cell site, and many other things. A technology such AI is well suited to wireless operators, who can adopt this technology to find a smart, easy way to optimize the network performance, as well as to simplify the maintenance.

Monica: As you say, it's a perfect case, because we have so much data, but also so much to learn, because we still don't know how to optimize the network. In this context, what is the role of automation?

Kevin: We use several concepts to describe what we call Wireless Intelligence. From our side, the first stage towards the automation of the entire network is some basic intelligence technology.

Automation is a bigger concept. We also started it as a basic-level technology, with a goal to make the entire network automatically operated or maintained. Network automation needs to start with some basic functions.

Monica: Mobile operators traditionally – and I would say surprisingly – have always done many things manually, in a just–turn-the-knob operation mode.

Automation is going to require a big change for operators. How big a change is it really, and what needs to happen? What needs to be done differently?

Kevin: First, I need to clarify one thing. Working with partners like Huawei, operators already use a lot of automated technology to maintain and manage their networks. A lot of tools also apply for network optimization, and they are very helpful.

The changes ahead of us are driven by new use cases. If you look at the network today, it's getting more and more complicated, because we want to deploy new layers, new technology, new spectrum.

Every day, we ask ourselves the same questions. Can we make the wireless indoor coverage better? Can we deal with the uneven traffic in the entire network? When we maintain a network, we also make a very good utilization issue for the entire asset.

We must get more revenue, so we should think about using this asset to get more revenue. Efficiency, also, should be considered. To achieve better results, we are thinking we should meet the new requirements. That's why we think about Wireless Intelligence technology.

We expect three changes to happen.

First, it's about performance and network optimization. With some combination of parameters manually set by human beings, we might meet some limitation, such as latency. The

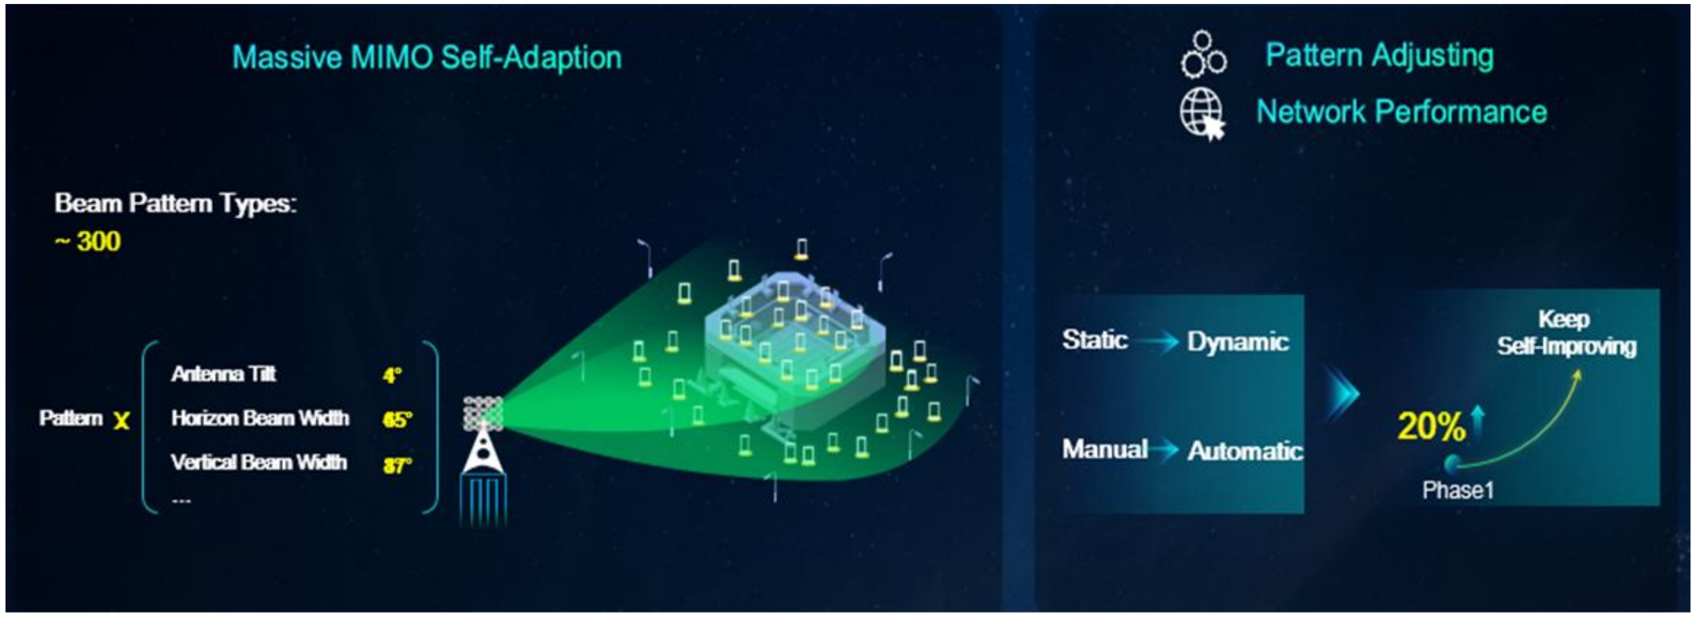

Wireless Intelligence tackles massive MIMO *Source: Huawei*

new machine learning algorithms will find the best way to solve a problem.

The second change might be in operations and maintenance, O&M. In this case, we might have some work we can deliver to the machine learning or intelligence technology, and free up people for other tasks.

The third change is about new business opportunities. Sometimes we define a scenario based on the coverage. In a rural area, an urban area – those are based on coverage.

Through Wireless Intelligence, maybe we may find other types of scenarios that open new business opportunities.

Monica: We will see huge changes both on the operational side and the revenue side. The better you manage network complexity, the more revenues you may get from your network assets.

In discussing any topic in telecoms today, we inevitably end up talking about 5G. But there is much work that can be done, should be done and will be done, before 5G. What is that work, in terms of both automation and Wireless Intelligence?

Kevin: In Huawei, we have an approach for using some technology for what we call 4.5G, which will prepare 4.5G to evolve to 5G. In the current stage, a lot of technology uses 4.5G. Massive MIMO is an example of a 5G technology that already is being deployed ahead of 5G as part of 4.5G.

For this kind of installation, some Wireless Intelligence technology also can be applied. For example, there're some parameters in this massive MIMO network. You need to consider a combination of different parameters to achieve the best performance.

This is deployed today. You have to solve these issues today, not with full 5G. For products such as massive MIMO, we have Wireless Intelligence solutions to help people find the best combination of parameters for all scenarios.

In urban areas, high-rise buildings, rural areas, high-density areas – automatically, Wireless Intelligence will enable operators to find the best combination of parameters. There are some cases where Wireless Intelligence is already in use, but there is also a much wider scope for future thinking.

Monica: A lot of the preparation for 5G can be done ahead of time. How dependent is the success of 5G on automation, Wireless Intelligence, and all technologies that rely on big data?

Kevin: First, 5G enables more possibilities for the industry as it adopts Wireless Intelligence. Our company's vision is to be the leading ICT solution provider, and we are committed to realizing this vision of a Wireless Intelligence connected world.

One major feature of 5G for operators is going to be eMBB, as well as enabling the wireless industry to be digitalized, and to connect more devices and processes, such as intelligent manufacturing, remote medication, and health treatment.

Point number one, 5G can enable a lot of new industries which already have intelligent solutions, but the wireless connection provides a better stage.

Point number two, 5G networks enable a lot of new features, like the bandwidth. The high throughput,

Customer value from Wireless Intelligence *Source: Huawei*

stability, and very low latency create possibilities for new applications. Previously, maybe GPS gave us location information. With some new features in 5G, maybe we can get a very accurate positioning through the whole network.

Monica: As you mentioned before, you've been working on this for some time now. It is all relatively new, but you have quite a lot of experience already. What have you learned as you've worked on this?

Kevin: To be honest, the telecom industry is not the leader in AI, because a lot of internet companies in other sectors started earlier.

So first, it's time for the entire industry to catch up with this golden age, to create real value.

Second, as I mentioned before, operators do have a unique advantage for using this technology, because they have great data. They update their data every day. All this input should be professionally analyzed. It's a very good asset.

Third, today we don't think we should be too excited about Wireless Intelligence technology. We need to consider, what's the real value? We should not follow some hot concept but find the real value for the operator.

Rather than hype the technology, we prefer to start with real values created for the operator. Step by step, we can approach new concepts.

Monica: You're working with a lot of operators. What are they telling you?

Kevin: In general, operators ask for this technology. They really care about this kind of technology. Recently, we've received more and more requests to cooperate on this topic.

They want to understand the general concept of Wireless Intelligence, and how to approach it. Today, there are many concepts, many studies, many voices in the wireless industry. They take different approaches, but they all point towards this real value for the network.

Monica: You've been working with operators on many case studies. One I find very interesting is the one on massive MIMO. Can you tell us about that?

Kevin: This is a very interesting and successful case. Massive MIMO technology is a new product concept which needs more optimization of parameters. With three or four parameters, plus the coverage scenario, you will have plenty of combinations. You can use them to configure your products. It'll result in different coverage, different performance.

In this massive MIMO case, you have 300 combinations as options for one site, one cell, to configure the parameters.

Doing it manually, you find several combinations based on your experience. You try one, and you try another one, to see the results, and you adjust a little bit. It takes maybe months, several months, to find the best combination of parameters for one cell, or one site.

Now we use a machine learning algorithm for this. The machine learning algorithm will set up one option. Let's assume it's not the best one; the optimization software will automatically try another one, get the feedback, then run a third one, and so forth.

This gets done very quickly – maybe only a few days, maybe less, if the optimal scenario is found early on. Machine learning has really reduced optimization time. Once the site is optimized, it will accommodate a huge amount of traffic, and this will create the value for operators by increasing revenue. It's a really exciting case.

Monica: It increases revenues, but also, they can get more out of the hardware that they already have in place. They can squeeze out better performance. It's a winning situation, both on the revenue side and on the cost side.

Kevin: Yes. Sometimes, people cannot find this best combination manually, but the machine learning can.

Monica: It's a tangible benefit. During all the trials you've done with operators, what are the key advantages of your solutions?

Kevin: There are several unique values from Wireless Intelligence.

First, it takes us beyond the limitations of expert experience. Previously, a lot of optimization was based on some engineer's experience, and on-site execution.

This new Wireless Intelligence technology is the best and most appropriate solution in some very difficult cases, where the complexity is too high for a human to tackle.

The second value is the efficiency improvement. Today, new elements of the network are deployed and used in new scenarios we find every day. With the resources we have today, how can we improve the network performance? Wireless Intelligence gives us a way to create a higher-efficiency outcome for the entire network, not just for any single one resource.

The third value is in the new opportunities. It's not artificial intelligence or Wireless Intelligence itself that creates new business opportunities. It's what people do with it. They're collecting data and analyzing data. This will be a very good input for the marketing and the promotion team. They can find more links from the sales to the network service. This is also an approach that works for the entire industry.

Monica: Let me ask you a final question. What's next at Huawei? What do you plan to work on over the next few years?

Kevin: In general, our vision already is to establish an intelligent, connected world. Huawei has several business units. I think you already heard about this On-Device AI, which was announced by our terminal side.

Wireless Intelligence is an important and interesting area – the entire wireless industry is focusing on this one; many other industries are too. This is one of our main research directions. For sure, we will keep an eye on it, and the research will continue for this one.

The second, as I've said, we are not in a rush to create a concept. We want to create real value. This is the major trigger point for Huawei Wireless. All our solutions and all our roadmaps are based on this point. We want to see what we can really do, and what value really can be created for our customers. This is how we do it.

The third direction in our vision for the future is 5G and 4.5G. The entire industry will be strengthened. The bandwidth, the performance in wireless networks will be much stronger and more powerful.

We do think all these things can be connected, and not only for the network element – maybe something else, too. The market requirements for Wireless Intelligence will grow fast, and we look forward to exploring more possibilities.

About Huawei

Huawei is a leading global information and communications technology (ICT) solutions provider. Driven by a commitment to sound operations, ongoing innovation, and open collaboration, we have established a competitive ICT portfolio of end-to-end solutions in telecom and enterprise networks, devices, and cloud technology and services. Our ICT solutions, products, and services are used in more than 170 countries and regions, serving over one-third of the world's population. With 180,000 employees, Huawei is committed to enabling the future information society, and building a Better Connected World.

About Xu Yan (Kevin)

Xu Yan (Kevin) is Head of Marketing Operations, Huawei Wireless Solution. Mr. Xu leads marketing strategy and operations, program development and execution, new product launches and regional business development across Huawei Wireless products and solutions. Previously, Mr. Xu worked for various Huawei regional offices, gaining extensive knowledge of global markets. Between 2012 and 2017, he held senior positions in Huawei's North Africa office, Middle East office, and Levant regional office where he served as Chief Wireless Engineer and Vice President of Products. During this time, he was responsible for product strategies and solutions design for significant projects in Morocco, Egypt, Jordan and Lebanon. Mr. Xu joined Huawei in 2006 as the Branding Director of Huawei Wireless Product Line. In this role he led brand management and operations across Huawei's Wireless products. Prior to joining Huawei, Mr. Xu worked for multiple telecommunication operators including China Mobile.

BT

Is zero-touch automation possible?

A conversation with Mansoor Hanif, Director of the Converged Network Research Lab, BT

In this conversation with Mansoor Hanif, Director of the Converged Network Research Lab at BT, we talked about how automation is evolving, from the scripts of a decade or more ago – which were limited in scope but did their job, helping to manage simple wireless networks. As we move to 5G and virtualized networks, automation is not just a requirement; it penetrates the wireless network core to RAN – so deeply that we may reach the ultimate goal of zero-touch automation.

Monica Paolini: Automation has been a hot topic lately. Mansoor, can you tell us about the history behind automation at BT and what type of efforts BT is making in this area?

Mansoor Hanif: For anybody who's been involved in any type of network operations over the last few decades, automation is nothing really new. Any operations engineer worth their salt has been doing some form of automation, whether using scripts or other methods, over the last 10 or 20 years.

The extra focus now is coming through the optimization of the techniques around automation, the ambition to get to zero-touch automation – which is something we haven't really tried to get to before or seen a need for – and the application of AI machine learning techniques.

Automation has been used in various areas of the fixed-network side of BT for many years. Our research team, my colleagues in the data analytics and AI machine learning team, have been working on various aspects of automation with the fixed network.

I came from the mobile side from the EE side, so I know from my own background in various networks that to a certain extent, automation has been applied to optimization and maintenance of mobile networks, as well.

Now we're trying to bring that together in a converged way, as far as possible – trying to bring together the learning from the fixed and mobile sides – and applying the latest techniques to get as close as we can to zero touch, even though I believe personally that zero touch is still a few years away.

Monica: What are the challenges with zero-touch automation?

Mansoor: Zero touch refers to completely removing humans from the loop, apart from setting the robot going or the automation going.

Networks are so complex, with so many variables that can occur in them, that at this stage it's quite dangerous to run a zero-touch automated process on anything significant that can impact the customer experience, simply because there's always going to be something unexpected.

There are areas of the network where you can completely automate the process. Not just the network, but also our finance systems, any type of systems across the IT estate. Those are usually what we call RPA, robotic process automation – very simple manual tasks that we're all aware of, the manual back-office tasks. Those are quite easy to do.

When you're talking about automating important functions across a network, it's possible to do that at a low level. But at any point in a very complicated loop with many interactions, you're likely to meet something totally unexpected.

At that stage, it's really important that, somewhere in that chain, there is an expert who understands the whole loop and can guide the system to the right conclusion. If you look at our mobile network today, for example, you've got roughly 19,000 sites in the UK. Many networks around the world, depending on the size of the country, are bigger; many are smaller.

We'd love to be able to say that each of those sites is identical, and that the landlord's identical, and that the access rates are identical, and what can go wrong in any site is identical, and the data we have on each site is equally accurate.

In reality those sites are hugely different, which means you need to define several thousand processes, or maybe even a dozen thousand processes, to be able to handle all the eventualities. That's why zero touch is still a little way off for the really complex network types of processes. But we can get fairly close to zero touch in some areas.

Monica: Which processes do you think can be fully automated, and which ones will need supervision?

Mansoor: Many areas can be fully automated. If you look at an optimization process – for example, a self-healing process or a dynamic capacity fulfillment process across a radio network, or even a fixed network – to get to zero touch would mean automation could handle 99.99% of the cases.

At the moment, I think that probably 80% of the cases can run through the automation loop with a really high confidence that there are no issues, and there are no risks involved. I'm just going to make up the numbers here, but it's the rough proportions that are important.

In the remaining 15–20% of the cases, the automation process itself may not be quite sure it's making the right decision. Therefore, those are the cases that should be put forward to an expert system, managed by an expert, to guide the algorithm.

That is also part of the learning process. Every time the human intervenes and says, "Hang on. In this case, this is the right course of action," and resets the process, that should all be memorized, used, put back into the machine learning algorithms to improve those processes.

The number of variations in those processes is such that it will take several years before we get closer to the 99% point. That's exactly what I mean when I say we can do a lot of automation but not completely zero touch at this stage.

Monica: As you say, a crucial part of the machine learning and AI is to learn from what you're seeing. You need artificial intelligence and machine learning to discover what you do not know already. If you knew already how to automate, then you wouldn't need AI and ML. As you learn, you are able to increase the level of automation that you can have in the network.

Mansoor: We've done a lot of that in some of our field operations. For example, as I said, my colleagues have been working on algorithms to optimize our resources, especially in the fixed network, which we've applied to our own network, and how to optimize the use of field engineers – for example, for Openreach.

There's a lot more we can do across the board there. The key thing is to be able to get to a point where the data we're feeding into the machine learning and AI algorithms is really useful and it's as close to real time as we can get.

We have two big challenges.

One is that, for example, the data that's filled in trouble tickets on the network is usually simple labels that don't give us enough information to make detailed decisions. The duty managers have to do a lot of filtering on the maintenance side to interpret those tickets.

If you try to automate those tickets today, you will come out with a huge number of variations, simply due to the way the tickets are labeled by the maintenance engineers. Two tickets could be exactly the same case, but they just use a different word. It's going to take a bit of time before we get everybody to understand the impact of entering data in a certain way.

One point is what level of detail we put in, and how that data should be presented to allow it to be used by machine learning or AI algorithms. The use of the algorithms, once we get the data inputs to a better quality, will increase as we get closer to real time.

The second challenge is that at the moment, the delay in processing is too long to make effective use of the data, for example, in scheduling your field engineers to the right site with the right tools, with priority over other faults impacting fewer customers.

We're working to get that closer to real time. There are two ways to do real-time analysis. It's by working either with crowdsourced data or with systems from third parties to be able to very quickly understand the impact of a fault on our customers and how to prioritize our field engineers to fix those faults.

At the same time, you can also use machine learning techniques to profile, during every part of a day, what is the typical impact on your customers of every single cell site. You can build a picture of potential impacts, and you could feed that into your algorithms. That's not real time, but it's using machine learning to simulate the real-time impact based on previous data.

We're looking at both approaches at the moment.

Monica: When you say crowdsourcing, it's basically getting data from the handsets of your subscribers to understand what they're seeing?

Mansoor: Yes. Obviously, all our mobile users use handsets or dongles, so to some extent, yes, coming from a handset. It doesn't mean necessarily from the handset itself. It can come from third-party service providers, application providers – whatever data our customers are comfortable sharing with any of our partners, we can make use of on an anonymized and aggregated basis.

We have our own technical data, but our customers are happy that, for example, Google or Facebook or Akamai or any third party such as a RootMetrics, P3 or OpenSignal provide us with that data, because it can complement our own data very effectively.

We're happy to work with all of those, but we need to make sure the algorithm is actually using all of those inputs, making them into a coherent whole, integrating them with our data, and then making sure that data is used correctly.

We're not fussy, but the biggest thing is that there are few third parties out there that can get a significant enough footprint of our overall user base, which in the UK is around 30 million or more, to make it worthwhile.

What we can't do is have 15 different sources of data, each one of them addressing a million users, and then try to mix and match. That would be really, really difficult to optimize. What we can do is work with three or four sources of data that cover the whole 30 million, then perhaps check that off with our own data and come to intelligent conclusions, and then optimize those with time.

Monica: That's going to require some time, because obviously you need to be able to properly understand the input you're getting from the third parties, which might have different ways to report. That's a lot of work.

Mansoor: That's a good point. At the moment, I can't mention any confidential details, but there are many third parties who are working with operators on the best way to do that. As you mentioned, at the moment, (a) not everybody's willing to share with other third parties and to integrate that, and (b) they may have different ways of representing that data, which makes it more difficult.

Those are all things we're actively working on through our own groups, but also with other operators through TIP, for example. We're in the TIP working groups on AI and machine learning to see where we can maximize the benefit they bring.

Monica: What are the parts of the network where automation has the most potential? You were talking, I think, mostly about the RAN. What about other areas? Are there any low-hanging fruits?

Mansoor: My own background is RAN. Obviously, that's what I'm focusing on, and I can see a lot of benefit there.

The RAN and the fixed network are similar in the biggest quality-of-experience benefits that we can get immediately, in remote areas in particular, by optimizing where we send field engineers – whether that's for fixed or mobile – and making sure that when they do go to a remote site, they go with the right equipment, they know what the issue is, and we send the right company.

I wouldn't say that's uniquely a radio issue, but also in many cases transmission. It can also be core, to a certain extent, but it's mainly transmission, transport, RAN, but also fixed network. It's the same issue: managing spare parts, managing tools, managing third-party subcontractors to make sure they're going to the right place at the right time with the right tools, and they have a good idea of the right problem to fix.

You'd be amazed how often, at present, we have to do several iterations before we get to the right loop. It's not specifically a RAN issue. It's more a field operations issue. That is the most useful in terms of what we can do today with optimizing, using AI and machine learning techniques in our network.

I'm looking specifically at the quality of experience aspect. We're also looking into many other aspects to help our customers in their daily travels. We'd rather start with the basic things that really, really matter to our customers more than anything else, which is quality of service, coverage, reliability.

Start with those. Then on top of that, you can build a whole layer of services that the customers might want.

Monica: Does automation bring any advantage when networks are working fine or there is no disruption?

Mansoor: What I described is getting the basics right when something does break down, how do you optimize what you fix first and make sure you fix it properly. That's the first step.

You can also use optimization to do preventive maintenance. A number of our partners are using statistics to look at preventive maintenance. That's certainly the next step forward.

Beyond that, there's self-healing. Especially as we go into 5G, this will be more prevalent with network slicing: you can use your data not only to optimize a field force but to reconfigure network slices dynamically, based on the information you have, to optimize a certain situation in real time.

That will be very, very relevant also for when you've got an issue on the network and you need to satisfy an SLA for a specific corporate customer.

Beyond actual faults, as we get more and more into dynamic networks – especially in the virtualized world of 5G – automation can help in adapting networks, flexing networks, moving functions around the network.

To get the most benefit of automation, we need to have an open discussion with our suppliers – this is already happening in the NFV space – to make sure the licensing models allow optimization over an area, the same as we did for capacity in the past, rather than only for a specific element or a specific function.

In the NFV world, we have the idyllic virtualized network spread out from the core to the aggregation nodes to the edge, when and if we need it. The future algorithms will define the best place to run those functions. We'll be automating not only the capacity of the network, but the capacity of the actual functions and the specific SLA that a customer needs from a specific function.

Again here, it should be very relevant in the corporate world or in the mission-critical world. We'll trigger the automation that moves a function from the core to the edge for a specific period and then pulls it back again when you don't need it.

It's going to get even more beneficial to have that automation in place. Before, it was just capacity, and maybe software capacity, because you couldn't move hardware around. As we go to a virtualized world, then you will have the capability of moving capacity and functions.

Now, in some cases you will need to over-dimension the hardware in certain areas to make it possible to move the software. That's one of the key things we're looking at in some of the TIP working groups – to understand how we can dimension hardware on a common platform in a cost-efficient way that gets the benefits of this flexibility provided by automation or the virtualized network.

Monica: That would apply to latency in conjunction with MEC, edge computing, or network slicing. You can manage not only vertically across functions, but also horizontally for different applications. It would require too many people to do it manually.

Mansoor: That's a good thing about automation. Automation isn't an enabler for any of those functions. In my personal view, it's the other way around: automation will allow us to scale to a point that it gets to be operationally cost effective.

In the past, we had the scripts. For the scale of operations we had, we could manage it manually. Today, if you're talking about moving capacity around, moving functions around, managing very different SLAs for different customers, and including different variables like latency, speed, throughput, whatever you're selling as a service, it's absolutely impossible to do that without automation at a scale that actually makes it worthwhile over a single platform.

That's why automation is so much of a key to the actual business model behind this.

Monica: At the same time, people are nervous about what's going to happen to the people, the cultural changes that automation introduces. Are we going to need fewer people, or are they going to do different things? Are they going to be made irrelevant? How do you see the cultural changes that innovation brings about?

Mansoor: This is a big subject. It's something that in BT we feel very strongly about, because we have 100,000 people who work for BT. Many of the people have been there for decades. We have such a big role in the UK economy. Beyond our own people, the knock-on effect of people who work for associated companies is so big.

We know that simply moving people or replacing them by robots is not our objective here. What we want to do is reskill people to the future. That's what we're going to be focusing on.

To a certain extent, we need to be more efficient. The nature of the work will change. But we are going to have to go through a massive cultural change and massive reskilling operation, and we're starting that already. We will be making sure that people understand the new techniques, and what their roles are in those.

Some of the old jobs, especially back office jobs, will be replaced to a certain extent. But they'll be replaced by jobs that, as our universe is expanding, are there to control those automated processes through the expert skills of the people who are managing them. But much more, the back-office jobs will change to customer-facing roles.

The customer is at the center of everything we do, and our main objective is to satisfy the customer. A lot of people who are perhaps working in the back office now can be reskilled to deal with customers, which is a great thing. From our contact centers for customer care to the retail stores to other types of customer service, we'll see a shift of roles.

I'm simplifying here, but I could see two directions of travel.

One is towards technical jobs, but towards controlling the autonomous processes or interfacing with layered skills. You would have a layer one, layer two, layer three skilled resource managing all the automation processes. People will be reskilled towards that.

The field resource would also continue to be important. They will be working in collaboration with automated processes to make sure the field resource is able to solve problems effectively and efficiently.

In the second direction, you'd have jobs geared towards customer interaction, which is always going to be key to us, because it's a huge part of customer satisfaction. It's having that interaction with real people, as well.

That's where I see the cultural changes happening. New players coming into this field could start with hugely automated systems from the beginning. Most big operators, on the other hand, have a social responsibility as well. We have a social contract. Therefore, we are absolutely committed to making those changes and reskilling our people.

The good news is that we also see that commitment from governments. The UK government, especially, is in line with that. In the budget for this year, for example, quite a lot of money's been set aside for putting new skills out there, in terms of AI in particular. There are, I think, 200 new PhDs that are going to be offered in UK universities. That area is quite well understood.

Monica: That's going to be important, because it's a long-term change. It's a good thing that it requires time, so everybody involved has the opportunity to adjust from a cultural point of view.

Mansoor: In the long term, people probably think of it as a long graph of gradual transition, with some tipping points perhaps. The reality is it's going to be a jerky process with ups and downs. There will be mishaps, and there'll be difficulties, and there'll be points where we're going to have to deal with issues. The reality is obviously a lot more complex than the graph.

About BT

BT's purpose is to use the power of communications to make a better world. It is one of the world's leading providers of communications services and solutions, serving customers in 180 countries. Its principal activities include the provision of networked IT services globally; local, national and international telecommunications services to its customers for use at home, at work and on the move; broadband, TV and internet products and services; and converged fixed-mobile products and services.BT consists of six customer-facing lines of business: Consumer, EE, Business and Public Sector, Global Services, Wholesale and Ventures, and Openreach. For the year ended 31 March 2016, BT Group's reported revenue was £19,042m with reported profit before taxation of £3,029m. British Telecommunications plc (BT) is a wholly-owned subsidiary of BT Group plc and encompasses virtually all businesses and assets of the BT Group. BT Group plc is listed on stock exchanges in London and New York.

About Mansoor Hanif

Mansoor joined EE in November 2011 and led the technical launch of the 1st 4G network in the UK and was also accountable for the integration of the legacy 2G and 3G Orange and T-mobile networks. Until 2016 he led the team who plan, design, rollout, optimise and operate all EE radio access networks, including Mobile Backhaul and Small Cells, and was accountable for the coverage aspects of EE's Emergency Services over LTE programme. He was also a board member of MBNL (the joint venture of EE with H3G) until 2016. During the acquisition of EE by BT, Mansoor led the EE network Integration team and is currently Director for Converged Networks and Innovation in BT R&D. He is a member of the BT Technology Steering Board, Scottish Innovation Programme, Intel CommSP Technical Advisory Board and alternate board member of the Telecom Infra Project (TIP).

Glossary

5G	Fifth generation
AI	Artificial intelligence
API	Application programming interface
BP	Back propagation
BSS	Basic Service Set
CA	Carrier aggregation
CDN	Content delivery network
CORD	Central Office Re-Architected as a Data Center
CPE	Customer premises equipment
CRAN	Cloud RAN
ECOMP	Enhanced Control, Orchestration, Management & Policy
eMBB	Enhanced mobile broadband
EMS	Element management system
ETSI	European Telecommunications Standards Institute
GPS	Global positioning system
GPT	General purpose technology
ICT	Information and communications technology
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
IT	Information technology
KPI	Key performance indicator
LTE	Long Term Evolution
LWA	LTE WLAN aggregation
MANO	Management and Orchestration
MEC	Multiple-access Edge Computing
MIMO	Multiple input, multiple output
mmW	Millimeter wave
NEM	Network Element Managers
NFV	Network Functions Virtualization
NFVI	NFV Infrastructure
O&M	Operations and maintenance
ONAP	Open Network Automation Platform
ONOS	Open Network Operating System
OSS	Operations support systems
PLC	Programmable logic controller
QoE	Quality of experience
RAN	Radio access network
RF	Radio frequency
ROI	Return on investment
RPA	Robotic process automation
SDN	Software-defined networking
SIM	Subscriber Identity Module
SLA	Service level agreement
SNR	Signal-to-noise ratio
SON	Self-organizing network
TCP	Transmission Control Protocol
TIP	Telecom Infra Project
UE	User equipment
VNF	Virtualized Network Function
WLAN	Wireless local area network

References

[1] Cisco, Cisco Global cloud index supplement: Cloud readiness regional details, 2017.
[2] GSMA, Unlocking commercial opportunities: From 4G evolution to 5G, 2016.
[3] International Telecommunication Union, IMT Vision – Framework and overall objectives of the future development of IMT for 2020 and beyond, Recommendation ITU-R M.2083-0 (09/15).
[4] McLuhan, Marshall, The Gutenberg Galaxy: The Making of Typographic Man, 1962
[5] Noction, Understanding the impact of network latency on service providers' business, 2012.
[6] Paolini, Monica, Getting the best QoE: Trends in traffic management and mobile core optimization, Senza Fili, 2016.
[7] Paolini, Monica, Massively densified networks: Why we need them and how we can build them, Senza Fili, 2016.
[8] Paolini, Monica, Mastering analytics: How to benefit from big data and network complexity, 2017.
[9] Paolini, Monica, Power at the edge. Processing and storage move from the central core to the network edge, Senza Fili, 2017.
[10] Paolini, Monica, The smart RAN: Trends in the optimization of spectrum and network resource utilization, Senza Fili, 2015.
[11] Sculley D. et al., Hidden Technical Debt in Machine Learning Systems", NIPS'15 Proceedings of the 28th International Conference on Neural Information Processing Systems, 2015.
[12] Smith Aron, Anderson, Monica, Automation in everyday life, Pew Research Center, 2017.
[13] Tapia, Pablo et al, Implementing Operational AI in Telecom Environments, Tupl, 2018.

About Senza Fili

Senza Fili provides advisory support on wireless technologies and services. At Senza Fili we have in-depth expertise in financial modeling, market forecasts and research, strategy, business plan support, and due diligence. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless, and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations. We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, use these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafiliconsulting.com, or contact us at info@senzafiliconsulting.com.

About Monica Paolini

Monica Paolini, PhD, founded Senza Fili in 2003. She is an expert in wireless technologies, and has helped clients worldwide to understand technology and customer requirements, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She frequently gives presentations at conferences, and she has written many reports and articles on wireless technologies and services. She has a PhD in cognitive science from the University of California, San Diego (US), an MBA from the University of Oxford (UK), and a BA/MA in philosophy from the University of Bologna (Italy). You can contact Monica at monica.paolini@senzafiliconsulting.com.