WHITEPAPER

The road to a self-healing network

VIA AIOps

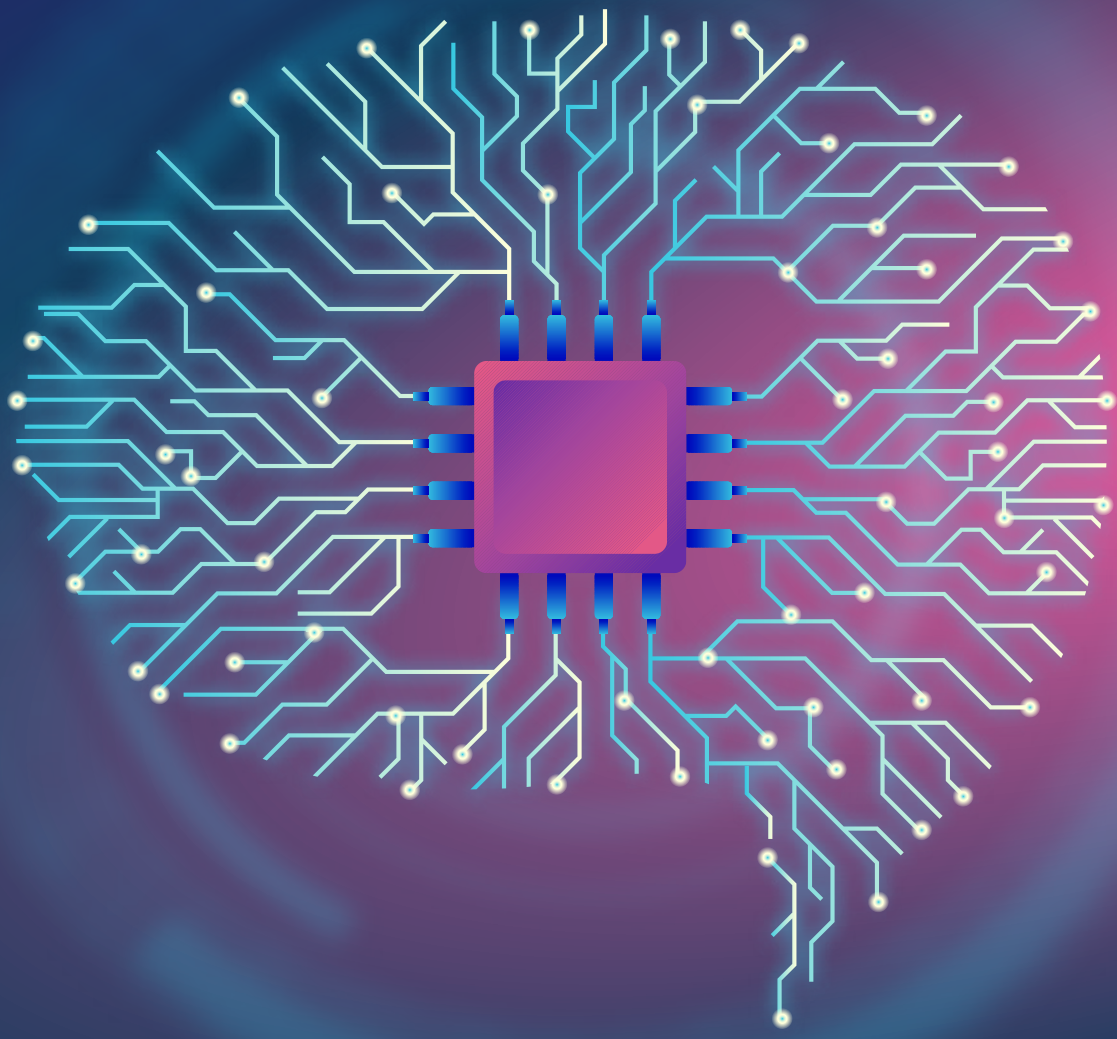

www.vitria.com

Telcos have begun the journey towards the self-healing network, but progress seen in research just conducted is slow and fragmented. One of the most striking findings has been the wide range of intelligence that will be needed – with the amount of diagnostic/predictive decisioning being significantly higher than other network automations. Indeed, the self-healing network might be the most concentrated area of intelligence requirements across a telco. This paper therefore starts by defining the self-healing network and then moves to discuss the intelligence and data required in order to fulfil telco requirements.

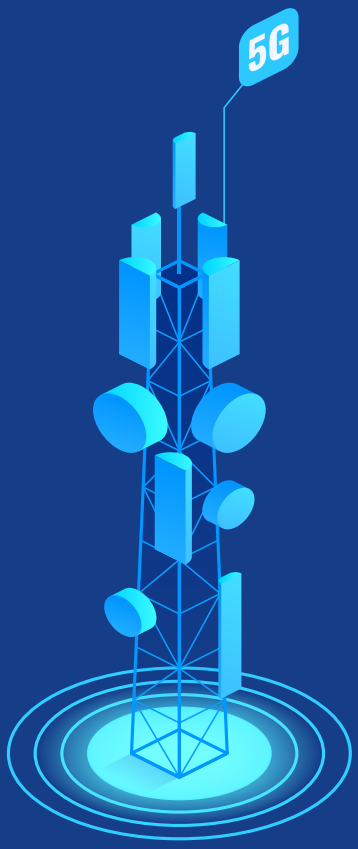

What is a self-healing network?

This paper starts by looking at the definition of a "self-healing" because, although it is part of the Self-X terminologies within the TM Forum's autonomous network project, it is often used to describe different functionalities within the network.

As one research participant discussed, "the sad reality is that many problems can be fixed by just turning off the offending hardware/software and turning it on again." This was certainly a feature of the first self-healing capabilities in SON REL 10, which focused on problems such as software glitches. The next iteration used self-healing to describe issues such as traffic re-routing in the event of a fiber cut in protected network configurations using SONET.

More recently, self-healing has been focussed around the cloud where services failover automatically to standby hardware and backup links. Indeed, the advent of virtualized networks allows the term to expand to quite simple activities (e.g. scheduling the nightly rebooting of a network function) and more complicated closed-loop activities adapting in real-time to faults or new demands on the network from services.

Figure 1 sets out to describe a future vision for self-healing and the various areas of activity that will be seen. In the columns, we have:

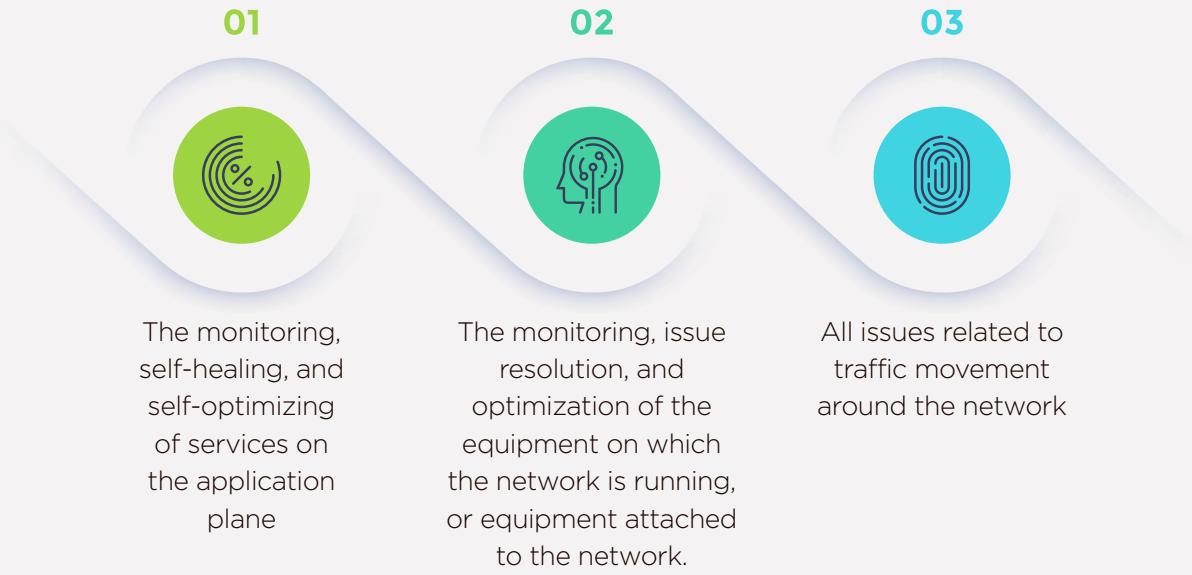

Figure 1: Definition of self-healing

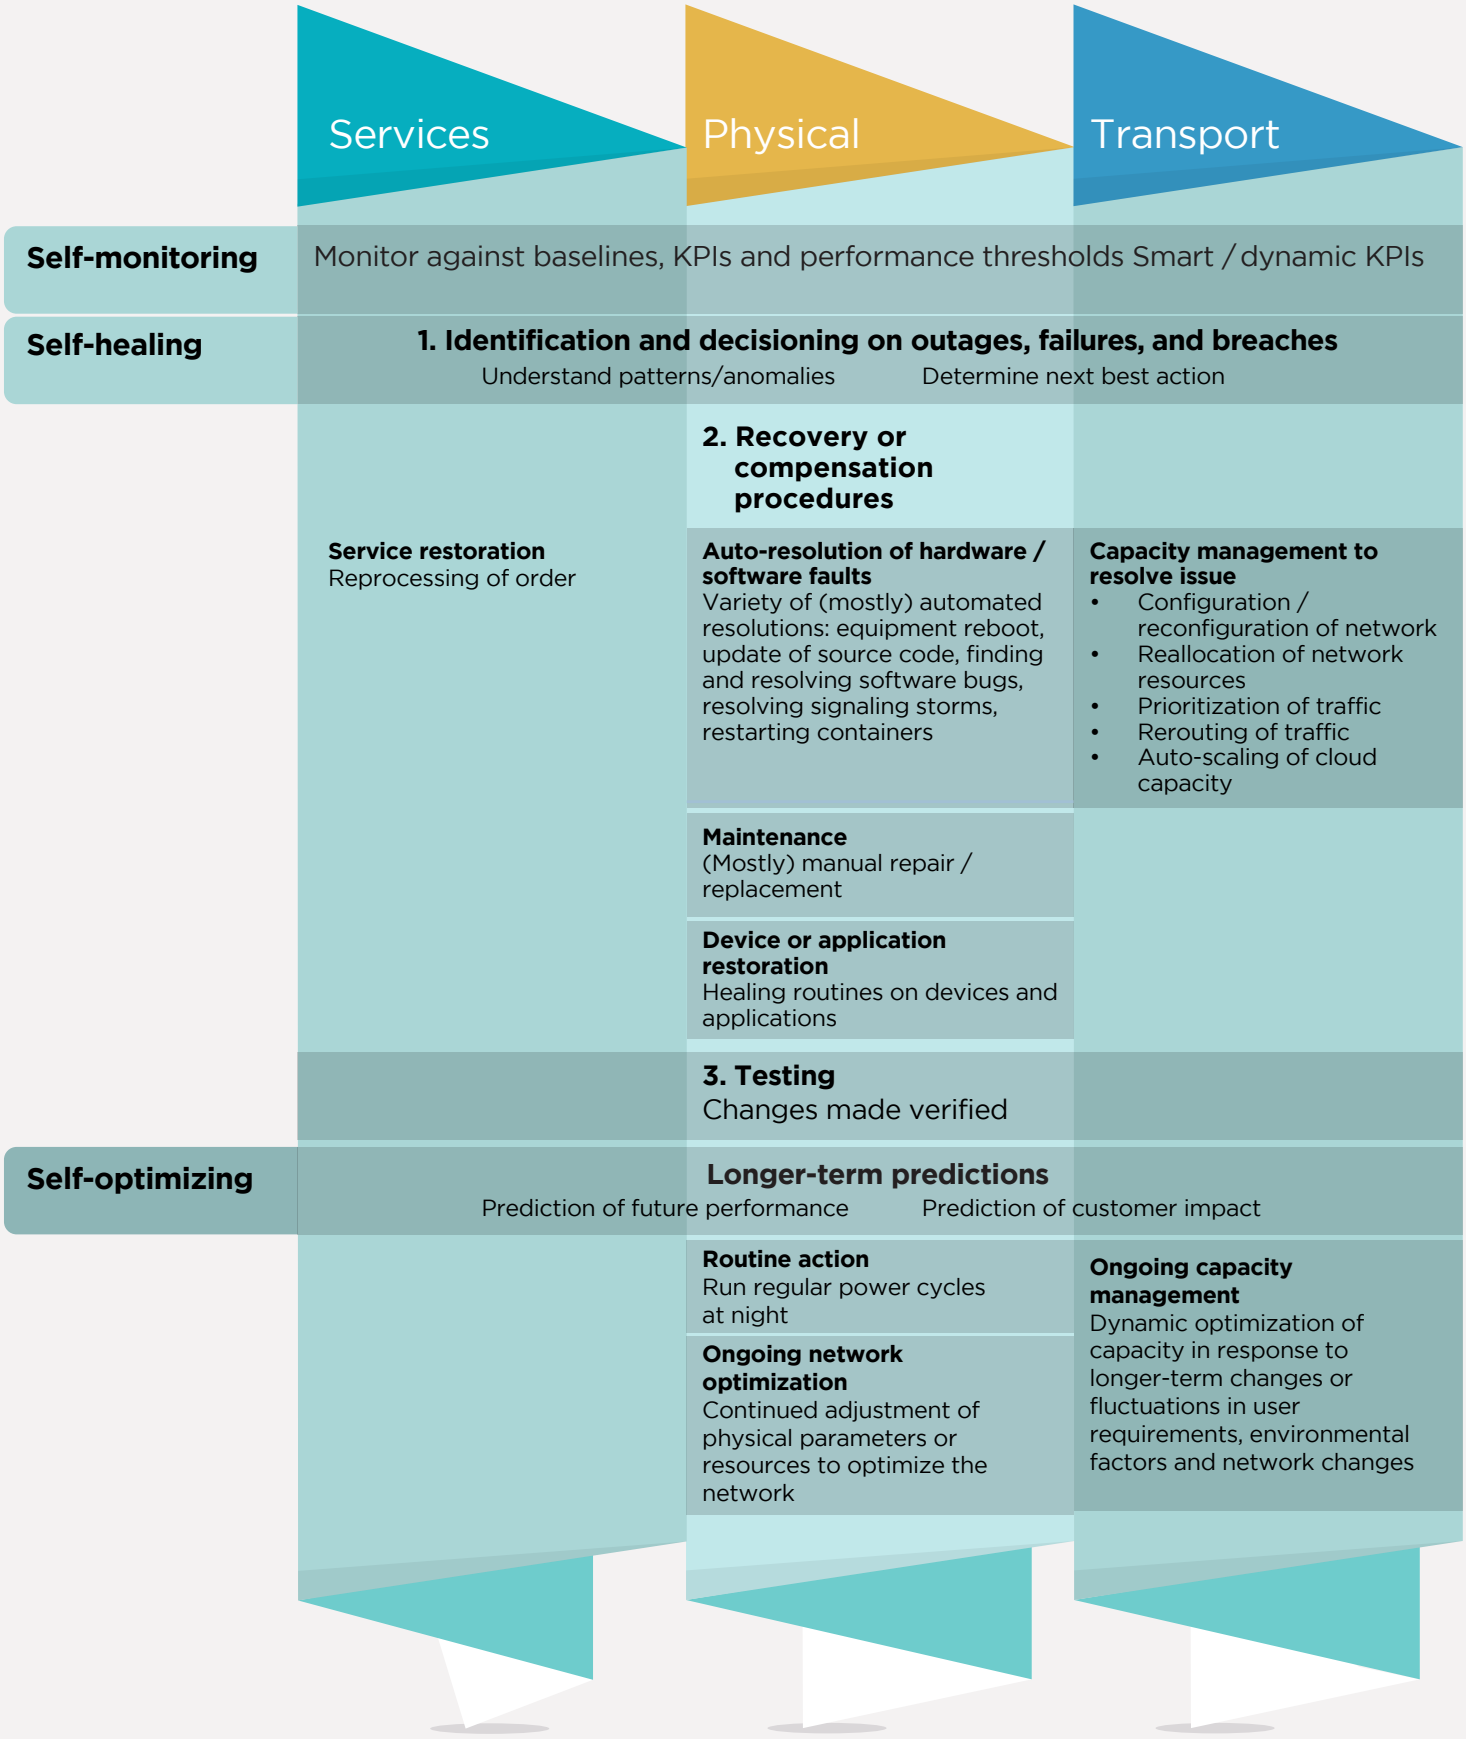

Source: Charlotte Patrick Research

Self-monitoring aims to continuously monitor the performance of key applications and traffic flows against various network requirements, baselines, and performance thresholds. This requires ingesting more data sources (e.g. traditional probing and other assurance including active assurance, customer experience insights or customer usage) from across the network with more granularity and frequency than previously.

Self-healing is the automatic detection and diagnosis of full/partial outages, failures, and breaches on the network. Recovery/resolution will either be automated and closed loop or require manual resolutions. Figure 8 below illustrates a typical self-healing process in more detail, but in short, the process monitors for faults and triggers a fault detection alarm when a potential issue is seen. This initiates a fault detection process, decisioning on whether there is a potential for self-healing or manual recovery, a diagnostic phase where the severity level of the fault is determined. At this stage, a recovery (e.g. repair or rebooting of equipment) or compensation (e.g. load balancing) procedure is initiated.

Self-optimizing is the automated adaption and improvement of efficiency and performance on the network, reacting to alerts from the self-monitoring systems or adjusting to continually meet intent, KPIs and other SLAs. It differs from the row above because it is not focused on responding to events and often deals with more underlying and reoccurring issues.

Drivers and barriers to the self-healing network

There are a range of significant drivers that will push forward the self-healing network. However, when discussing barriers to the deployment of self-healing networks, there were at least three mentions of elephants! All these mentions are related to the need to consume small pieces of the "self-healing" elephant to achieve goals.

Figure 2: Drivers and barriers to the self-healing network

Self-healing:

- Part of the move to autonomous networks
- Simplifies network operations
- Minimize network downtime
- Improve network efficiency

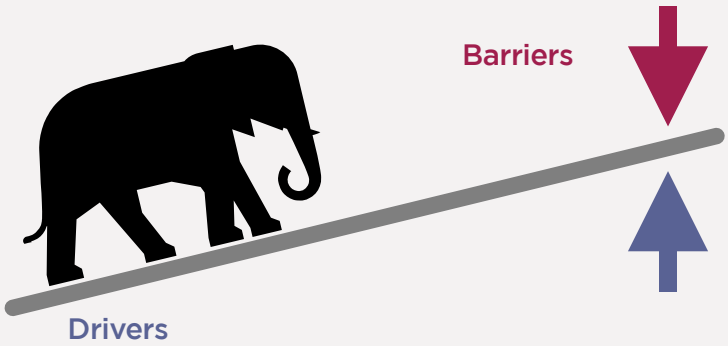

Source: Charlotte Patrick Research

Drivers

Part of the move to autonomous networks: Telcos hope that adding AI and automation will move the needle on general goals such as customer experience, reduced operational costs, network uptime, and support for increasing network complexity. This complexity comes from a range of potential new revenue streams for the telco:

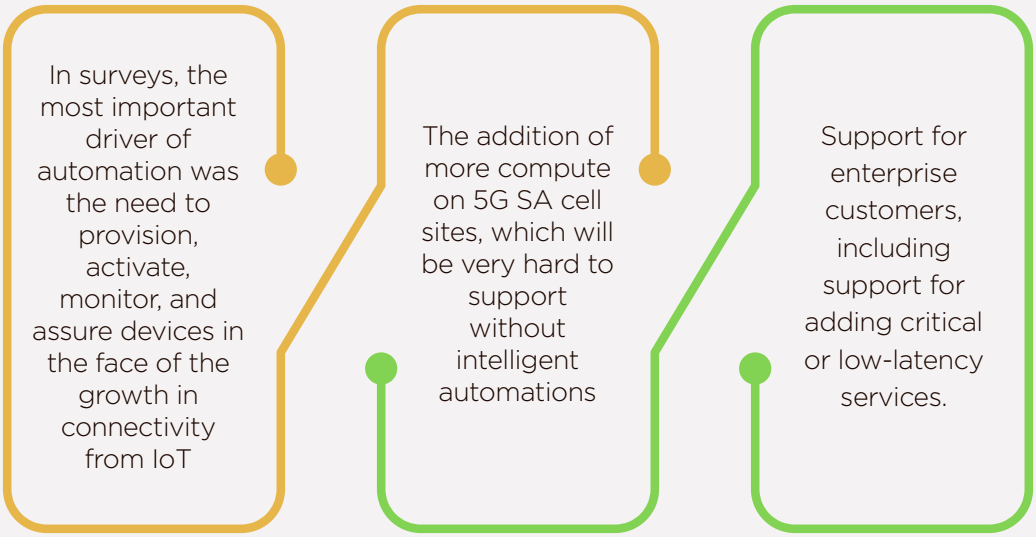

Simplifies network operations: Improved network self-monitoring should eliminate a significant amount of the mundane and repetitive manual interventions undertaken by network administrators. This would enable the (much reduced) NOC/SOC teams to focus on high-value tasks and innovation and help them avoid drowning in alerts.

Minimize network downtime: Self-healing aims to provide faster remediation for network failures/quality issues and improve business-critical application availability – increasingly without human intervention.

Improve network efficiency: Self-optimizing networks aim to enhance network efficiency, reliability, and quality of service while minimizing operational costs and manual configuration efforts.

Barriers

Interviewees discussed a wide variety of technical developments to enable self-healing, as shown in Figure 3 – including quality data, model development and process improvement.

Figure 3: Barriers to self-healing

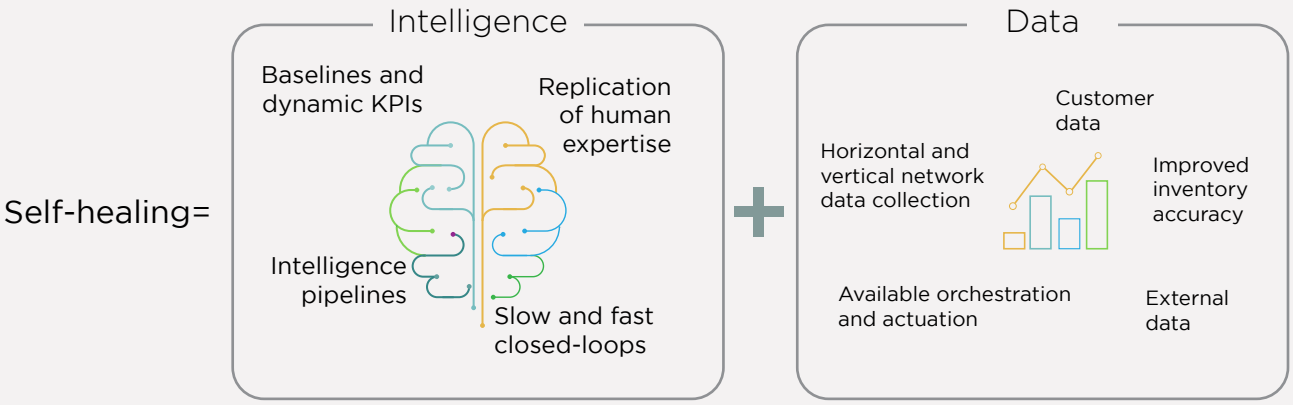

Source: Charlotte Patrick Research

Intelligence

Baselines and dynamic KPIs: Understanding how the network behaves using baselines and dynamic KPIs (which move in line with expected network behavior and real-time network conditions, technologies, and user demands) allows for optimization and fault detection.

Replication of human expertise: One of the main barriers to more advanced self-healing is that it requires human expertise, which is challenging to replicate. Models that will be deployed will replicate human-like thinking. Examples include:

Machine reasoning, which uses rule-based reasoning, deductive inference, pattern recognition, and probabilistic models to perform root cause analysis, predict maintenance requirements and provide insight for optimization.

Reinforcement learning, which learns from feedback in the form of rewards or penalties, adjusting its actions to achieve better outcomes over time.

Intelligence pipelines: ML agents will need to work together with "captured intelligence" and "local observation" to solve the more complex problems in self-healing. The models in a pipeline might solve one specific problem and pass on the answer to another model or act together with other models to solve a problem in parallel; a model might also be responsible for managing the underlying data set needed by other models or intelligently execute the prescribed results from upstream models. The phrase "captured intelligence" is used to describe solutions such as databases that hold the gathered wisdom on a topic. Graph databases not only hold a list of the inventory on the network but also model the relationships between the inventory, thereby holding knowledge about the dependencies between items in the list. Time series databases model evolving networks by tracking changes over time. An AI agent in a particular network domain may provide "local observation" or intelligently execute prescribed actions coming from other upstream models in the pipeline. These agents may include a Gen AI reasoning engine or more simple ML models and collect data when an issue occurs, collect data to fulfil a specific request or configure network elements on request (writing APIs or code as needed).

Slow and fast closed loops: Self-healing requires a large number of:

Fast closed loops for rapid response to immediate issues (e.g. dynamic demand allocation)

Slow closed loops, where data trends are looked at over a longer period and responses are more gradual and less reactive. Often used for longer-term optimization (e.g. capacity optimization)

Data

A **recent TM Forum survey** found the three main challenges facing the successful deployment of AIOps for assurance: inconsistent and fragmented data, a lack of data analytics expertise, and a lack of explainability in models. Self-healing requires a breadth of data not seen in any other process in the network. Data types include:

Horizontal and vertical network data collection: A full range of passive and active assurance data that stretches horizontally and vertically as required. Fault management may require frequent collection, and diagnostic routines will need the right data set at a granular level. The TM Forum survey found that multi-vendor environments were often the biggest barrier to collecting these data sets – with 50% of telcos initiating efforts to ensure a comprehensive assurance system for 5G.

Improved inventory accuracy: Creating visibility into provisioned resources and network topology is increasingly difficult in a multi-vendor, multi-domain, and multi-technology environment such as 5G. Inventory must be up-to-date where there are short-lived cloud network functions (CNFs), and vendors with self-healing often have to create shadow inventories to enable accurate root-cause analysis and the orchestration of necessary changes.

Customer data: The augmentation of network insight with customer data, both demographic and usage data, to enable decisioning.

External data: Ingestion of third-party data, such as weather conditions, information from supplier websites or web tools provided by the electricity companies, brings additional dimensions to diagnostics.

Available orchestration and actuation

Top-level and domain orchestrators and other actuators in the network need to be available to complete the self-healing prescriptions. This is typically still a 3-4-year journey and will underpin the deployment of self-healing decisions.

In Conclusion

Self-healing looks set to be one of the last automations in the telco network to reach maturity. Indeed, fully replacing human decision making in some of the areas of more complex decisions is going to be extremely difficult. It is therefore expected that it will take at least 8 years, if not more, to complete the self-healing network journey.