WHITEPAPER

New Frontier: The Self-Healing Network

Organizations are making slow but steady progress toward realizing the value of self-healing networks. One reason is that unlike other network automations, the self-healing network requires a wide range of intelligence to support diagnostic and predictive decisioning. In fact, the self-healing network represents the most concentrated area of intelligence for telcos and other network dependent businesses like financial services.

This paper starts with the current definition of the self-healing network and then unpacks the data and intelligence requirements to unlock the full potential of this advancement.

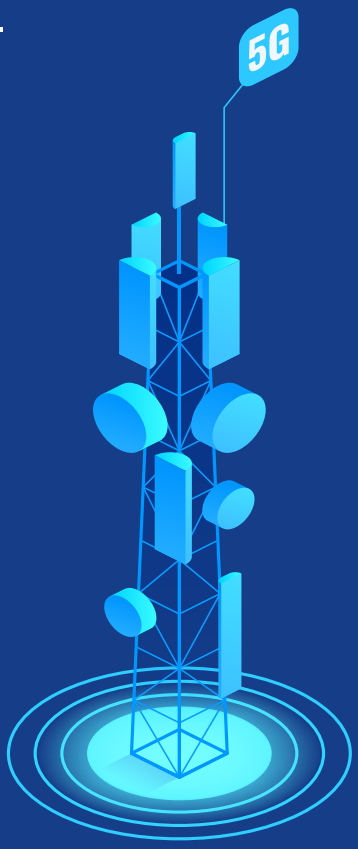

What is a self-healing network?

Self-healing networks are increasingly prevalent in modern network architectures. They leverage automation and artificial intelligence (AI) to constantly monitor network performance, identify potential issues and even take corrective action. Self-healing networks gather data on aspects of the network like traffic flow, device health and use AI and machine learning algorithms to analyze the collected data looking for patterns that predict potential problems. Finally, if an issue is detected the system is designed to automatically take corrective action.

Advancements in network management tools and protocols offer more visibility into network operations. As networks became more complex and critical to business functions, the need for automated responses to failures continues to grow.

A Look Back

Around 2019 the TM Forum, a global industry association serving telecommunication service providers, suppliers, and consultants, published the "Autonomous Networks Manifesto" which positioned self-healing networks as a 'key element for achieving network resilience and efficiency' and described some self-healing capabilities like turning devices on and off and re-routing traffic in the event of a fiber cut in protected networks. Most recently, self-healing is discussed in the context of cloud where services automatically failover to standby hardware and backup links. The advent of virtualized networks expanded self-healing to include other very simple activities (e.g. scheduling the nightly rebooting of a network function) and the introduction of more complicated closed-loop activities like adapting in real-time to faults or to new service demands on the network.

Today, self-healing is defined in three dimensions:

Self-monitoring: monitoring against baselines, KPIs and performance thresholds

Self-healing: includes decisioning, recovery and compensatory procedures example: rerouting traffic) and testing to verify changes.

Self-optimizing: predicting impact on customers and predicting future performance.

Three Stages of Maturity for the Self-Healing Network

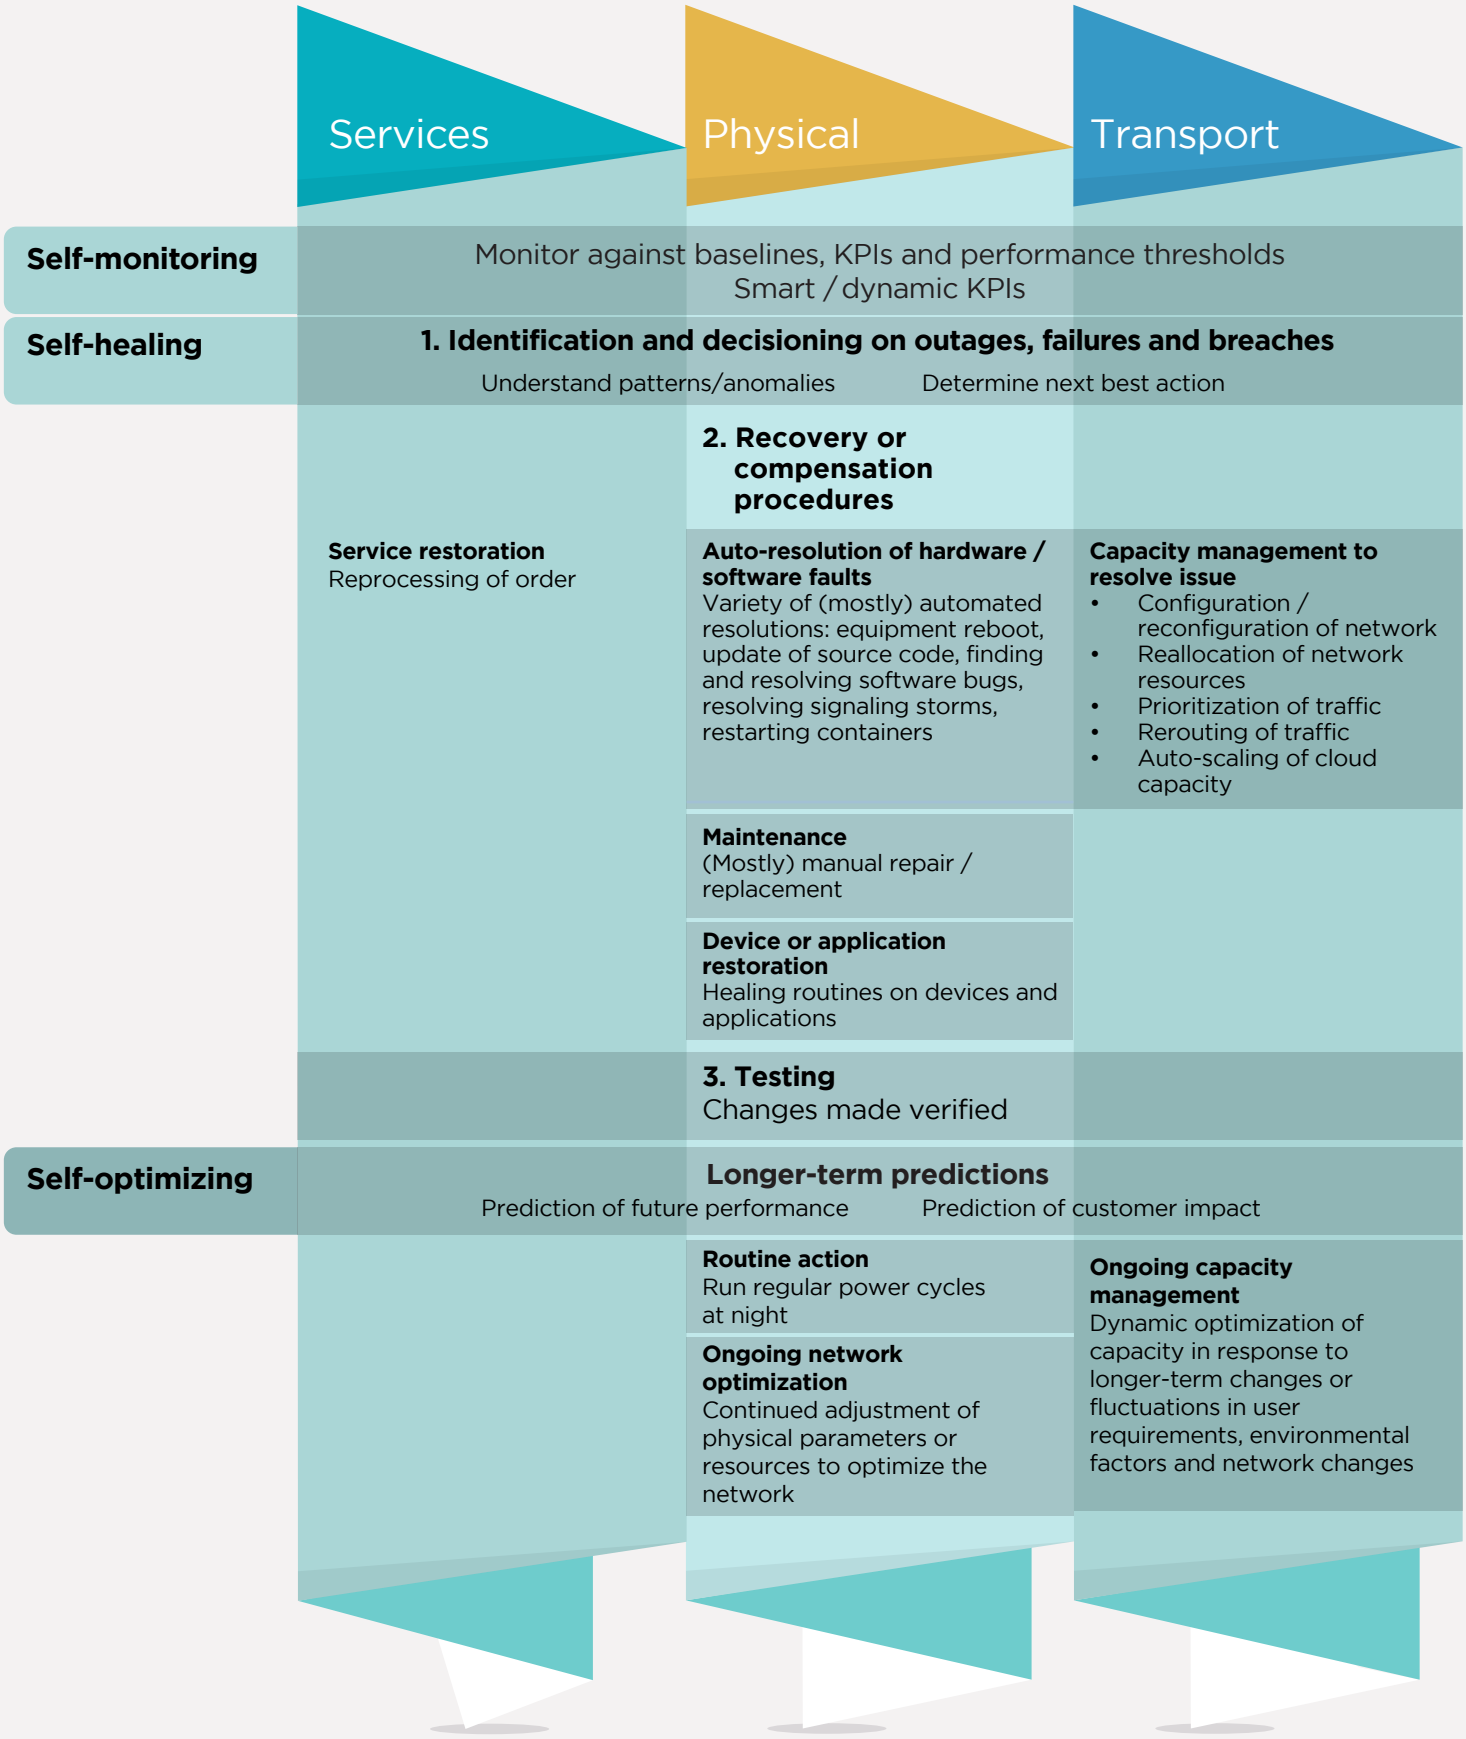

Figure 1 Illustrates three stages of maturity for the self-healing network and provides the self-optimizing vision for the future.

Source: Charlotte Patrick Research

Self-monitoring provides for continuous monitoring of application performance and traffic flows against various network requirements, baselines, and performance thresholds. This stage requires ingesting more data sources from across the network with increased granularity and frequency.

Self-healing provides for more automation - automated detection and diagnosis of full/partial outages, failures and breaches on the network and automated and closed loop resolution or manual recovery when required. The self-healing process monitors for faults and triggers a fault detection alarm when a potential issue is seen and provides decisioning on applying self-healing or manual recovery.

Self-optimizing is the automation of processes that improve efficiency and performance on the network. This stage includes reacting to alerts from self-monitored systems or adjusting systems to continuously align to intent, KPIs and other SLAs. At this stage, the focus is not on reacting or responding. Instead, this stage is focused on why a fault occurred – what are the underlying and recurring issues that can be addressed to obviate failures in the future.

Drivers and barriers to the self-healing network

All businesses are experiencing increasing expectations for network resiliency. In general, network and service operations will enable them to improve the customer experience, reduce operational costs, increase network uptime and support increasing network complexity.

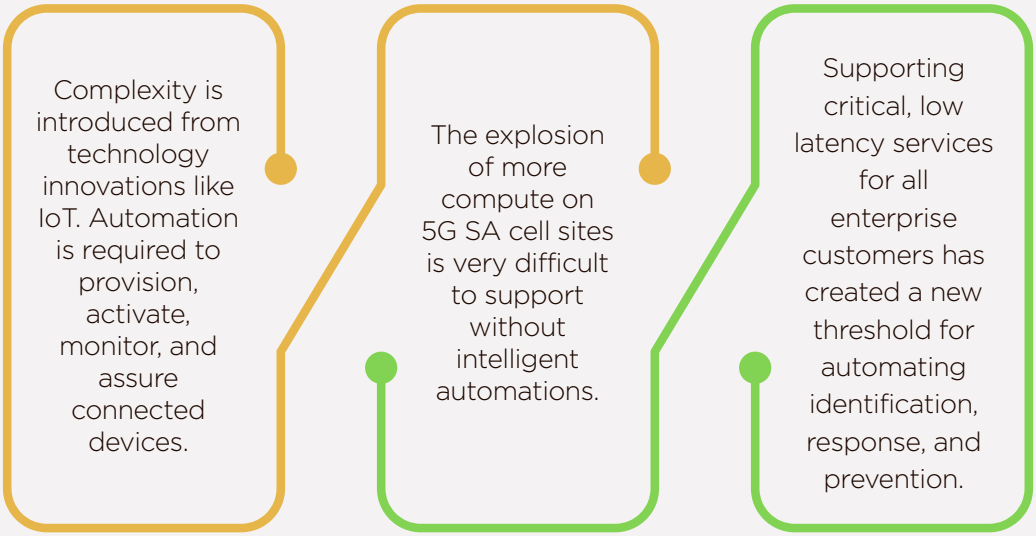

Simplify Network Operations: By improving network self-monitoring organizations may eliminate a significant amount of the mundane and repetitive manual interventions undertaken by network administrators. This would enable the resized NOC/SOC teams to avoid drowning in alerts to focus on high-value tasks like implementing innovations.

Minimize network downtime: Self-healing is shown to speed up remediation of network failures issues impacting quality often without human intervention. This improves availability of business-critical applications.

Improve network efficiency: Self-optimizing networks aim to enhance network efficiency, reliability, and quality of service while minimizing efforts to manually configure networks, in turn reducing operational costs.

Barriers to Self-Healing

During interviews conducted by Charlotte Patrick, independent analyst, participants discussed a wide variety of technology that will enable self-healing. These include good quality data, advancements in model development and process improvements. As shown in Figure 2, operational intelligence and quality data contribute to the evolution of self-healing.

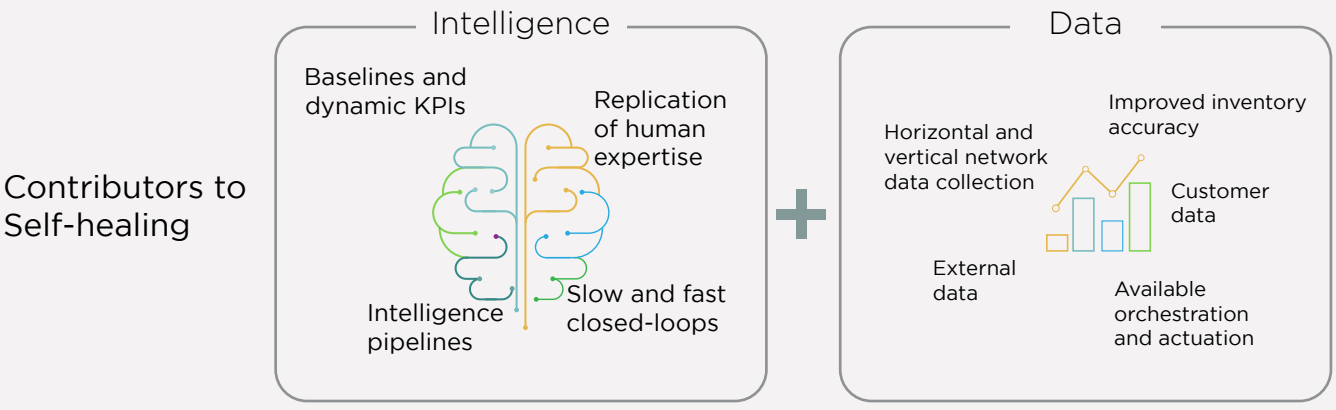

Figure 2 Source: Charlotte Patrick Research

Intelligence

Baselines and dynamic KPIs: By using baselines and dynamic KPIs operations can implement fault detection and optimization. Dynamic KPIs move in line with expected network behavior and real-time network conditions, technologies, and user demands.

Replication of human expertise: A primary challenge to more advanced self-healing is that it requires the replication of human expertise. These are some examples of models which will be deployed to replicate human-like thinking:

- Machine reasoning 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 learns from feedback in the form of rewards or penalties and adjusts actions to achieve better outcomes over time.

- Intelligence pipelines use machine learning agents working in concert to combine "captured intelligence" with "local observation" to solve the more complex problems in a self-healing network. The models in a pipeline solve one specific problem then pass the answer to another model; or one model acts together with other models to solve problems, in parallel. One model might 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*** describes entities like databases that hold all gathered wisdom on a topic. A graph database may hold a list or inventory on the network but also model the relationships between the items in the inventory. Time series databases model evolving networks by tracking changes over time. In contrast 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. They collect the data to fulfil a specific request or configure network elements on request.

- Slow and fast closed loops: Self-healing requires the use of **fast closed loops** to provide rapid response to immediate issues – for example providing dynamic demand allocation. **Slow closed loops** provide data trends from looking at data over a longer period of time, providing responses that are more gradual and less reactive. Slow closed loops are often deployed for long term optimization use cases – like capacity optimization.

Data

Self-healing requires the collection of diverse data that may not be used by other processes in the network. Data types include:

Horizontal and vertical network data collection: This represents a full range of passive and active assurance data that stretches horizontally and vertically, as required. Fault management may require frequent collection, while diagnostic routines will need the right data set at a granular level. A survey done by the TM Forum found that multi-vendor environments presented the most challenge.

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 current. In a self- healing network, where there are short-lived cloud network functions (CNFs) vendors create shadow inventories to enable accurate root-cause analysis and orchestration of necessary changes.

Customer data: More organizations are augmenting network insight with customer data, both demographic and usage data, to enable decisioning.

External data: Ingestion of third-party data like weather conditions, information from supplier websites or web tools provided by other utilities brings additional dimensions to diagnostics.

Available orchestration and actuation: Completing self-healing prescriptions require availability of top-level, domain orchestrators and other actuators in the network. This is typically still a 3-4-year journey and will underpin the deployment of self-healing decisions.

The Future and the Role of AIOps

Self-healing is the latest automation in network management. While incremental improvements are made each year, analysts like Charlotte Patric believe it could be at least eight years before the journey is complete. However, Self-healing networks require three core functions:

Data collection: Continuously gather data on various aspects like connectivity, traffic flow and device health.

Analysis with AI/ML algorithms: These analyse collected data to identify patterns and predict potential problems before they impact service.

Automated remediation: If an issue is detected the in-place system takes corrective action like rerouting traffic or reconfiguring devices.

AIOps offers a range of functionality from diagnostic to pre-emptive identification of issues followed by the initiation of prescriptive actions. To be successful you need a process oriented AIOps solution which harnesses the power of AI to automate tasks and digitally transform services. There is a lot of value to be gained by implementing an AIOps solution that combines Automation, AI, and Analytics like VIA AIOps by Vitria.

Andrew Colby, VP at Vitria, explains the vital role VIA AIOps can play in achieving the level of automation required to deliver on the promise of the self-healing network. He provides this advice based on customer experiences:

A Start small in new areas by applying basic automation to repetitive back-office processes.

B Let teams learn from multiple automations to experience what good looks like and apply learnings to more complex implementations.

C Measure effectiveness against a range of goals including near term savings and long-term industry goals and mandates.

D Create a robust feedback loop to track impact and compare projects.

Self-healing networks are real and continuously maturing. They are improving operational effectiveness. Implementing VIA AIOps by Vitria provides the automation and analytics foundation for organizations serious about developing a self-healing network.

Self-healing networks offer a range of functionality from diagnostic to pre-emptive identification of issues followed by the initiation of prescriptive actions. To be successful you need a process oriented AIOps solution which harnesses the power of AI to automate tasks, digitally transforming services provided by the NOC. Today the capabilities have evolved to include enabling complicated closed-loop activities in real time – to respond to faults or enable new demands for network services.

Self-healing processes monitor for faults, triggers detection and initiates **decisions for manual or automated recovery** (like a reboot) or enabling a **compensating procedure** like load balancing. Vitria's VIA AIOps provides a foundational platform for realizing the benefits of the self-healing network.

VIA AIOps delivers the process automation capabilities required to transform operations and dramatically lowers cost. VIA delivers intelligent automation from a powerful platform that combines AI, analytics, and machine learning in real time. VIA provides Telcos with a modern operating model that reduces costs and enables a superior customer experience to support a leaner, more efficient, and effective operations staff.