

From Human-Centric to AI-Assisted Operations

Tim Raphael

Regional Product Manager
IP Automation

Agenda

- Today's challenges
- What does "AI Assisted" mean?
- What exactly are LLMs?
- What do we need to get to AI assisted?
- What about the standards bodies?
- Nokia's view



Current Challenges

The human-centric approach

Today's fault management processes give operators control, judgement and accountability but...

- Scale and complexity causes:
 - High alarm volumes without insight into root causes, distributed among engineers
 - Limited visibility into historical trends
 - Reactive monitoring and break/fix after the fact
 - Manual fix execution with the potential for human error
 - Automation limited to predefined rules and known patterns.

Operators tend to fall into 3 groups

Reactive

- Alerts are handled individually with the tendency to "flood".
- Issues arise from threshold breaches .
- Change freezes used as risk control measure.
- Remediation is manual, usually via CLI.

Minor Automation

- Some amount of automated alert handling / pre-triage.
- Standardised remediations executed with playbooks / scripts.
- Still human-driven, tiered teams for triage, handling and resolution.


ML Enhanced

- Alerts are correlated by time, target and alert category.
- Trends are identified before threshold breach.
- Criticality is generally triaged before humans see the alert.
- Historical trends are reported upon

What AI Assisted means

What “AI-assisted” means

- AI-assisted does not **replace operators** or decision authority
- Improves **access to contextualised information** leading to more informed decisions.
- **Humans remain accountable** for operational outcomes and stay “in the loop”.
- Leads to **better MTTRs** and **clearer lessons** learnt from outages and failures.



But what exactly is an ~~AI~~LLM?

- LLMs are statistical predictors for the next token based on a huge corpus of **text**.
- This makes LLMs **NON-DETERMINISTIC**
- They're good at: summarisation, translation, pattern recognition, fuzzy matching, classification and structured extraction.
- Agents and agent harnesses use tool calls to **add context**, which **steers the prediction model** towards a goal.
- This combination of **language "understanding"** in the right context makes LLMs good for:
 - Root cause analysis
 - Semantic searching
 - Reasoning ($A + B \Rightarrow C$)
- Humans are typically not very good at these en-mass, we have (comparatively) low parallelism and working memory.

Potential of LLMs

- Replace humans having to **understand the flow of events** and **use tribal knowledge** to determine the **correct action**.
- **Agent Pipelines** are the answer
- There are **good/cheap** ways of applying LLMs to operations, and there are **less-cheap options**:
 - Throwing all the raw inbound information into an LLM and asking basic questions like “what caused this?”
 - Loading LLM context with operations manuals, how-to guides, metric definitions etc. etc. (everything)

Ingredients for AI-Assisted Operations

Agent Pipeline for Network Operations

1

Use deterministic tools to pre-process alerts, discard noise for focus on realistic signals.

2


Use classical AI (Machine Learning) to find trends, correlations and patterns in datasets.

3


Use smaller/cheaper models to classify for priority, routing or severity.

4

Use a cost-balanced model for root-cause analysis, action recommendations, tool calling for enrichment and to operate a HITL cycle.



Alarm/Event Correlation




ML-based and dynamic solution helping to reduce the noise

ML based Analytics – Baselining, Clustering, Predicting

Reporting



Thresholding



Baselining



Clustering



Predicting



Passive

Reactive

Active

Pro-active

Predictive

Use of
AI/ML

Collect current and historical data for effective visualization and analysis


Detect situations where customized indicators violate defined thresholds

Learn the expected range of values for key resources

Identify things that appear to be behaving differently than others

Anticipate behavior that might happen in the future

Identify patterns that represent real issue




Smart TCA

Multi-Variate analysis



Outlier detection


Before becoming service affecting




Trend prediction

Unified access to network data

- Context should be carefully controlled – **more does not mean better**.
- Data should have **sufficient metadata** to aid understanding and improve correlation results.
 - Origin /source
 - Validity / age / expiry
- Data aggregation improves speed of access and benefits more than just agentic use-cases.
 - Operators can cross-check agentic conclusions.
 - Dashboards and observability signals can be pre-calculated to improve time-to-diagnosis.



LLM of Choice?



Supported LLM Options

- Operator-hosted LLMs (On-prem / Private Cloud)
 - Full data control and regulatory alignment
 - Requiring investment ie GPUs
- External / Hyperscaler LLMs
 - Rapid access to state-of-the-art models (for PoC or targeted use)
 - Running cost to be considered over time
- Small Language Models (SLMs)
 - Lower cost, domain-focused inference
 - Suitable for high-volume, repeatable tasks

Integration Principles

- LLM-agnostic architecture via gateway / API abstraction
- Model selection per use case, not hard-wired dependency
- Support for switching models over time
- Guardrails for cost, accuracy, and data governance

Hybrid combinations

SLMs for day-to-day operations

LLMs for complex reasoning and summarisation

What do the standards bodies think?

Standards Perspective



AI Operations (AIOps) Programme
Autonomous Networks Reference Architecture
(IG1251)
AIOps Lifecycle & AI Agent concepts
IG 1274 & IG 1274M



I E T F[®]

RFC 7575 — Autonomic Networking: Definitions and Design Goals

RFC 7576 — General Gap Analysis for Autonomic Networking

RFC 8993 — A Reference Model for Autonomic Networking




ETSI GS ZSM 002

ETSI GS ZSM 009-1

Strong alignment across standards on where AI can/will assist:

1. Fault correlation / denoising
2. Anomaly and outlier detection, trend analysis and prediction
3. Impact analysis and prioritisation
4. Context enrichment & event reconstruction
5. Decision support (options, recommendations, confidence levels)
6. Closed-loop readiness (without execution)

Path to Higher Levels of Autonomy




Solve today's challenges and pave the way for self-managing networks.

AUTONOMOUS NETWORKS REFERENCE ARCHITECTURE (IG1251)

Evolving to higher levels of autonomy requires a well-designed management and control architecture

1. Decoupled Operational Layers
2. Intent-driven, Open Interfaces
3. Closed Loop Automation
4. Endogenous Intelligence
5. Single Domain Autonomy
6. Cross Domain Collaboration
7. Supports interactions between Autonomous Domains of different Autonomy Levels



In Summary

- AI-assisted does not **replace operators** or decision authority
- Improves **access to contextualised information** leading to more informed decisions.
- **Humans remain accountable** for operational outcomes and stay “in the loop”.
- Leads to **better MTTRs** and **clearer lessons** learnt from outages and failures.

In Summary

- Agentic operations give operators a productivity boost, improving operational metrics
- Success relies on:
 - Structured processes and workflows
 - Reliable access to high-quality network data
 - Access to advanced LLMs, strapped into agentic pipelines with guardrails
 - Human-in-the-loop style workflows

NOKIA

Copyright and confidentiality

The contents of this document are proprietary and confidential property of . This document is provided subject to confidentiality obligations of the applicable agreement(s).

This document is intended for use by 's customers and collaborators only for the purpose for which this document is submitted by . No part of this document may be reproduced or made available to the public or to any third party in any form or means without the prior written permission of . This document is to be used by properly trained professional personnel. Any use of the contents in this document is limited strictly to the use(s) specifically created in the applicable agreement(s) under which the document is submitted. The user of this document may voluntarily provide suggestions, comments or other feedback to in respect of the contents of this document ("Feedback"). Such Feedback may be used in products and related

specifications or other documentation. Accordingly, if the user of this document gives

Feedback on the contents of this document, may freely use, disclose, reproduce, license, distribute and otherwise commercialize the feedback in any product, technology, service, specification or other documentation.

operates a policy of ongoing development. reserves the right to make changes and improvements to any of the products and/or services described in this document or withdraw this document at any time without prior notice.

The contents of this document are provided "as is". Except as required by applicable law, no warranties of any kind, either express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose, are made in relation to the accuracy, reliability or contents

of this document. SHALL NOT BE RESPONSIBLE IN ANY EVENT FOR ERRORS IN THIS DOCUMENT or for any loss of data or income or any special, incidental, consequential, indirect or direct damages howsoever caused, that might arise from the use of this document or any contents of this document.

This document and the product(s) it describes are protected by copyright according to the applicable laws.

is a registered trademark of Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

Autonomous Networks Reference Architecture

The diagram illustrates the progression of AI agents from manual operations to full autonomy, categorized into Business and Service Operations.

Business Operations:

- Initial state: **Single Business Agents** (Planning, Deploy, Monitor, Optimization).
- Intermediate state: **Business Operations** (Planning, Deploy, Monitor, Optimization) and **Local Intelligence**.
- Final state: **Knowledge and Intelligence** (Application Catalog, Model Training, Knowledge Base, Data Services).

Service Operations:

- Initial state: **Resource Operations** (Planning, Deploy, Monitor, Optimization) and **Local Intelligence**.
- Intermediate state: **Resource Operations** (Planning, Deploy, Monitor, Optimization) and **Local Intelligence**.
- Final state: **Resource Operations** (Planning, Deploy, Monitor, Optimization) and **Local Intelligence**.

AI Agent Evolution:


- L0: Manual Operations & Maintenance** (Human Agent).
- L1: Assisted Operations & Maintenance** (Human Agent).
- L2: Partial Autonomous Networks** (Human Agent).
- L3: Conditional Autonomous Networks** (Human Agent).
- L4: High Autonomous Networks** (Human Agent).
- L5: Full Autonomous Networks** (Human Agent).

AI Agents:

- Human Agent** (Human Agent).
- AI Agents** (AI Agents).

- Fault management is a primary entry point for AI and autonomy
- AI is first applied in an assistive role
- AIOps provides correlation, anomaly detection, and decision support
- Humans remain in control and accountable for operational decisions
- Closed-loop remediation is a governed, later evolutionary step

IG 1274 & IG 1274M



Standards Perspective - ETSI ZSM

Closed-loop Governance and Control

ETSI GS ZSM 002 V1.1.1 (2019-08)



**Zero-touch network and Service Management (ZSM);
Reference Architecture**

ETSI GS ZSM 009-1 V1.1.1 (2021-06)



**Zero-touch network and Service Management (ZSM);
Closed-Loop Automation;
Part 1: Enablers**

- End-to-end, adaptive closed-loop management for networks and services
- AI used to analyse conditions and support decision-making
- Strong emphasis on policy, governance, and operational safety
- Human supervision explicitly supported, especially in early automation stages
- Incremental evolution toward zero-touch operations

Standards Perspective - IETF / IRTF Perspective

Internet Research Task Force (IRTF)
Request for Comments: 7575
Category: Informational
ISSN: 2070-1721

M. Behringer
M. Pritikin
S. Bjarnason
A. Clemm
Cisco Systems
B. Carpenter
Univ. of Auckland
S. Jiang
Huawei Technologies Co., Ltd
L. Ciavaglia
Alcatel Lucent
June 2015

Autonomic Networking: Definitions and Design Goals

Abstract

Autonomic
is self-
self-healing
function
central
across

This document
what a
reference
Network
Management

Internet Research Task Force (IRTF)
Request for Comments: 7575
Category: Informational
ISSN: 2070-1721

S. Jiang
Huawei Technologies Co., Ltd
B. Carpenter
Univ. of Auckland
M. Behringer
Cisco Systems
June 2015

General Gap Analysis for Autonomic Networking

Abstract

Internet Engineering Task Force (IETF)
Request for Comments: 8993
Category: Informational
ISSN: 2070-1721

M. Behringer, Ed.
B. Carpenter
Univ. of Auckland
T. Eckert
Futurewei USA
L. Ciavaglia
Nokia
J. Nobre
UFRRGS
May 2021

analysis
distributed
showing the
extent to
human
features
needed in the
Research

A Reference Model for Autonomic Networking

Abstract

This document describes a reference model for Autonomic Networking for managed networks. It defines the behavior of an autonomic node, how the various elements in an autonomic context work together, and how autonomic services can use the infrastructure.

- Long-standing research on autonomic and self-healing networking
- Introduction of self-* properties: self-configuration, self-healing, self-optimization
- Intent-based interaction to express “what” rather than “how”
- Hybrid operational models where autonomic functions and human operations coexist
- Strong emphasis on safety, gradual adoption, and predictability