Build or Buy? The Telecom Executive's Guide to AI-based Network Monitoring

As CSPs realize the benefits they can achieve from AI-based network monitoring they are faced with the immediate question: build our own system — or buy one? We outline the benefits and drawbacks of each approach, providing both the calculations and conceptual considerations you need to weigh in order to achieve the right decision for your organization.

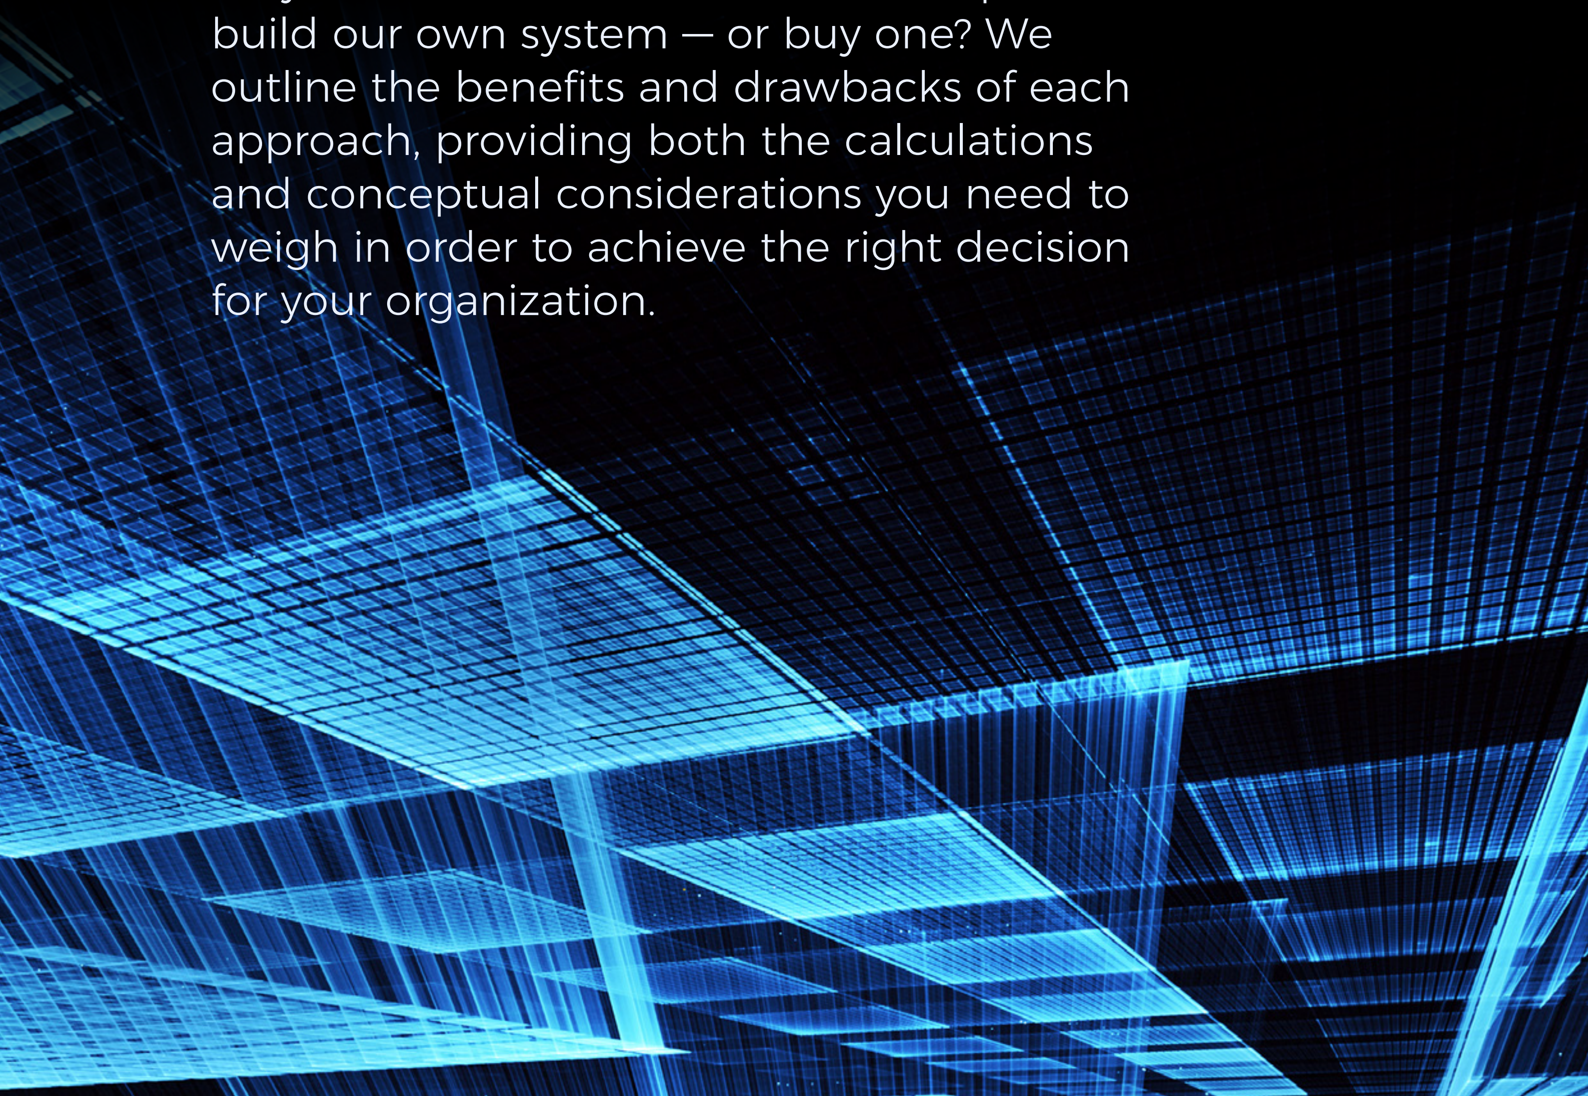

Introduction

In a highly competitive market, CSPs are vying to drive operational efficiency, deliver a better customer experience, and prevent critical performance and quality of service issues across the network, dramatically reducing time to detection and resolution that impact the bottom line. For every operation, the goal is to move from reactive problem solving to proactive monitoring, enabling stakeholders to know more about what is happening across their network elements and domains and fix incidents before minor issues escalate into bigger problems.

CSPs need to stay on top of hundreds of metrics, but with the ongoing growth in operational complexities, effectively managing radio networks, current and legacy core networks, services, transport and operations is becoming a radical challenge.

Static network monitoring gives rise to billions of alarms with a very high rate of false positives, since it's based on manual thresholding for a system that is too complex and volatile to adhere to predetermined states. What is worse - static monitoring leads to late detection of service degraragation and incidents. Even after detection, which often occurs when the incident has already impacted customers and appears in downdetector, there is no context to go on for expedited resolution.

> *The goal is to move from reactive problem solving to proactive monitoring, enabling stakeholders to identify and fix incidents before minor issues escalate into bigger problems*

CSPs monitor their network using a variety of monitoring tools. There are separate fault management and performance management platforms for different areas of the network (core, RAN, etc.), and infrastructure is monitored separately. Although these solutions monitor network functions and logic – something that would seem to make sense — in practice this strategy fails to produce accurate and effective monitoring or reduce time to detection of service experience issues.

The main reason for this dramatic shortcoming is that these tools can't detect service experience degradations. They monitor the network in silos — every network layer as a stand alone and every network type differently — and utilize rule-based or static thresholds. Due to thresholding limits alert storms are common, or alerts aren't generated. The siloed approach prevents effective correlation between related issues. As a result, the NOC team's only way to understand the actual service impact and experience is by collecting customer complaints and looking at "downdetector", which typically takes anywhere from a few hours to a few days, resulting in significant revenue loss and damage to brand reputation.

Modern monitoring is predominantly autonomous, relying on Machine Learning to monitor and correlate huge data streams in order to surface anomalies in real- or near-real time, without human intervention. Typically, autonomous network monitoring systems both alert users of issues and glitches, and provide data visualization for enhanced observability.

In an attempt to bridge the gap in existing OSS tools, CSPs are implementing advanced AI monitoring systems on top of their existing tools.

CSPs that use advanced monitoring systems across their stack typically experience significant costs savings from the early detection of incidents and high ROI resulting from:

- Reduction in time to incident detection
- Reduction in time to incident resolution
- Reduction in total number of alerts
- Reduction in the number of non-actionable alerts
- Reduction in workload on support operations
- Improvement in customer satisfaction scores

As CSPs realize the benefits they can achieve from autonomous monitoring compared to manual / static methods, they're faced with the immediate question: build our own system — or buy one?

In the next pages we will outline the benefits and drawbacks of each approach, providing both the calculations and conceptual considerations you need to weigh in order to achieve the decision that is right for you.

What does it take to build your own monitoring platform?

The build option poses multiple conceptual, technical and resource challenges, and is therefore usually only viable for extremely large, innovative companies with a dedicated team of AI researchers and developers. Depending on the robustness of the solution the CSP chooses to pursue, some build scenarios could take more than four years to develop, particularly for large, complex, and changing monitoring needs.

To achieve the capacity and prowess of a fully matured autonomous monitoring platform, there are three solution maturity levels to be pursued one on top of the other: Basic, Intermediate, and Advanced.

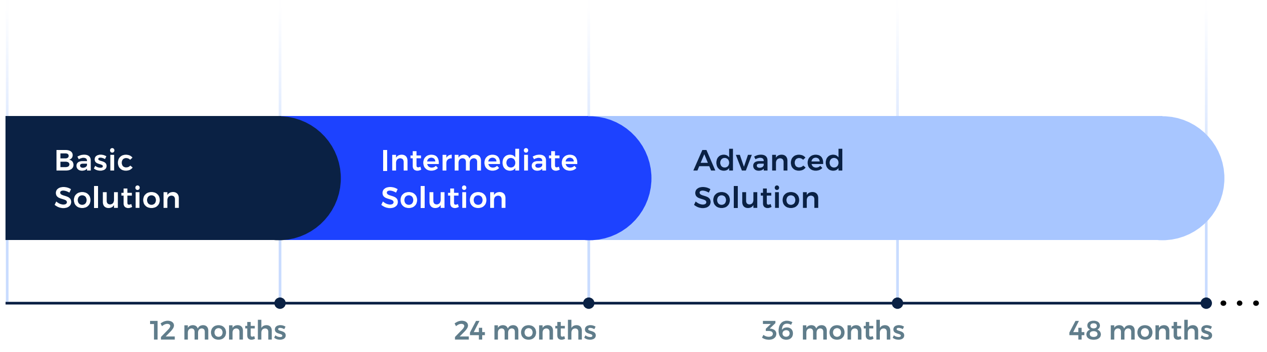

Solution Maturity	Duration (Months)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)	TCO (USD)
Basic	12	2,565	2,009	7,551	$4,052,836
Intermediate (+Basic)	24	5,045	4,034	16,541	$7,377,904
Advanced (+Intermediate)	48	13,250	6,169	36,656	$14,817,699

In the next pages we outline the feature development plan and TCO for each level, including detailed feature descriptions, limitations, alternatives, required resources, and incremental development and maintenance costs.

Building a basic monitoring platform

Solution Maturity	Duration (Months)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)	TCO (USD)
Basic	12	2,565	2,009	7,551	$4,052,836

Data Integration

1. Metric definition: normalize KPIs and measures
2. Provide push API from data source to platform

LIMITATIONS

Requires implementation of API from every new source of data (coding effort)

ALTERNATIVES

Use known API metric normalizations (Graphite)

Alerts

1. Alert types
   - Static threshold alerts
   - Single metric anomaly alerts
2. Conditions
   - Incident duration
   - Incident magnitude (values above/below static values)
   - Anomaly yes/no condition

LIMITATIONS

- No ability to filter anomaly based on significance
- Creates alert storms when real incidents occur (no alert correlation)
- No ability to filter anomaly alerts based on business context
- No ability to consider the effect of external events

Administration

Internal user management

Anomaly Detection

1. Normal behaviour modeling
   - Manual seasonality setting
   - A single baseline using simple statistics, usually “week over week”
   - User driven on/off
   - Adaptation capability: Manually selected during normal learning, usually very noisy
2. Anomaly types
   - Transient anomaly detection

LIMITATIONS

Applicable for small amounts of time series because:

- High false positive rates when manual settings incorrect
- Requires significant user input to exclude metrics that cannot be modeled with simple statistics.

Time Series Analytics

Ad-hoc functions on top of raw metrics

ALTERNATIVES

Perform functions in data source

Visualization

1. Time series and baseline charts
2. Dashboard creation and management capability

ALTERNATIVES

Base your charting on existing time series/dashboarding solutions

Outgoing Integration

Email Alerts

ALTERNATIVES

Perform functions in data source

Investigation

1. Repository of all incidents and alerts
2. Slice & dice visualization
3. Ability to save filtered views

ALTERNATIVES

Existing visualization platforms

Architecture

Data retention policies & implementation

ALTERNATIVES

Perform functions in data source

Feature	Resource	Resource Qty	Effort (Man Days)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)
Data Integration	SW Eng.	1	60	60	7	88
Timeseries Analytics	SW Eng.	1	180	180	7	208
Anomaly Detection	Data Scientist	1	365	365	30	485
Alerts	SW Eng.	3	180	540	30	900
Outgoing Integration	SW Eng.	2	90	180	30	420
Investigation	FE Eng. + DS	2	365	730	30	970
Administration	DevOps Eng.	1	90	90	10	130
Visualization	FE Eng.	2	180	360	30	600
Architecture	SW Eng.	1	60	60	10	100
Product	PM	1	-	-	365	1460
QA	Automation Eng.	0.5	-	-	365	730
UX	UX / Design	0.5	-	-	365	730
R&D Manager	R&D Manager	0.5	-	-	365	730
Data Science Manager	DS Manager	0	-	-	365	0

ASSUMPTIONS			
Platform Size (Metrics)	1,000,000	Hosting Price/Metric/Month	$0.003
Duration (Years)	4	Avg. Annual Employee Cost	$135,000

Building an intermediate monitoring platform

Solution Maturity	Duration (Months)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)	TCO (USD)
Intermediate (+Basic)	24	5,045	4,034	16,541	$7,377,904

Data Integration

1. CLI based connector per data source in organization
2. Complex queries to data source

LIMITATIONS

Requires installation of connector at the source environment (IT effort)

ALTERNATIVES

Assuming use of existing integration platforms

Outgoing Integration

1. Webhook alerts
2. Handle time zone differences, DST changes

ALTERNATIVES

Use existing integration platforms

Administration

Single Sign On based on your IdP

Anomaly Detection

1. Normal behaviour modeling
   - Automated seasonality detection using fourier transform (FFT)
   - 1-2 statistical baseline algorithms (e.g. Holt-Winters, Seasonal Hybrid ESD)
   - Manual or simple rule baseline selection
   - Adaptation capability: simple rule driven normal adaptationy
2. Anomaly types
   - Pattern change detection
3. Statistical confidence test based anomaly scoring
4. Manual rule based anomaly correlation

LIMITATIONS

- Not applicable for large amount of metrics, especially business/ digital experience type metrics
- Does not cover over 60% of metric types - especially irregularly sampled metrics (e.g, usage metrics) which tend to be measured irregularly and are highly non stationary
- FFT based approach does not accurately capture multi-season scenarios - requires significant manual work to fix
- Known algorithms do not adapt well when there is anomalous data - creates false positives and false negatives around real anomalies

Investigation

1. Highlight leading dimensions in the incidents
2. Incident management - acknowledge received alerts

Alerts

1. Alert types
   - Missing data alerts
2. Conditions
   - Send updates on alert

LIMITATIONS

- Creates alert storms when real incidents occur (no alert correlation)
- No ability to filter anomaly alerts based on business context
- No ability to consider the effect of external events

Architecture

1. Production & DR sites
2. Data protection policy – In transit and at rest

Feature	Resource	Resource Qty	Effort (Man Days)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)
Data Integration	SW Eng.	1	30	150	10	190
Anomaly Detection	Data Scientist	3	365	1095	60	1815
Alerts	SW Eng.	2	180	360	30	600
Outgoing Integration	SW Eng.	1	240	240	30	360
Investigation	FE Eng. + DS	1	365	365	20	445
Administration	DevOps Eng.	1	90	90	20	170
Architecture	SW Eng.	1	180	180	30	300
Product	PM	1	-	-	365	1460
QA	Automation Eng.	1	-	-	365	1460
UX	UX / Design	0.5	-	-	365	730
R&D Manager	R&D Manager	0.5	-	-	365	730
Data Science Manager	DS Manager	0.5	-	-	365	730

ASSUMPTIONS			
Platform Size (Metrics)	1,000,000	Hosting Price/Metric/Month	$0.003
Duration (Years)	4	Avg. Annual Employee Cost	$135,000

Building an advanced monitoring platform

Solution Maturity	Duration (Months)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)	TCO (USD)
Advanced (+Intermediate)	48	13,250	6,169	36,656	$14,817,699

Data Integration

1. UI based connectors
2. Self service to data analysts and business users
3. Integration additional capabilities:

- Time Zones
- Daylight Saving Time handling
- Gaps in data
- Delays in data arrival
- Out Of Order data arrival
- Data repair
- Data Readiness – watermarking

Time Series Analytics

1. Composite functions on top of raw metrics
2. Manage computations timing

ALTERNATIVES

Perform functions in data source

Investigation

1. Tools to collaborate & bookmark over the incidents
2. Snooze alerts as a whole, or partially to minimize noise

Anomaly Detection

1. Normal behaviour modeling
   - Robust and efficient seasonality detection (Anodot patent pending) ACF based
   - 6 and more baseline algorithms
   - Advanced classifier based baseline selection
   - Adaptation capability: ML-based normal adaptation, ML-based adaptation during anomaly
   - ML based consideration of event regressors
2. Anomaly types
   - Trend change detection
   - Slow trend detection
3. ML-based anomaly scoring
4. ML-based anomaly correlation

Administration

1. Manage groups of users
2. Provision users based on your organizational user management platform

Alerts

1. Alert types
   - Anomaly alerts
2. Combinations of conditions and automated conditions to minimize number of alerts
   - Correlated metric values
   - Correlated anomalous metrics
   - Number of anomalous metrics in incident above/below value
   - Auto discard low volume alert
3. Correlations
   - Event correlation
   - Alert correlation – minimize number of alerts per incident

Outgoing Integration

1. Additional destinations
2. API calls to consume alerts by 3rd party apps

ALTERNATIVES

Use existing integration platforms

Architecture

Scalable architecture (unlimited)

Feature	Resource	Resource Qty	Effort (Man Days)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)
Data Integration	SW Eng. + FE Eng.	3	410	815	40	1015
Timeseries Analytics	SW Eng.	1	180	180	30	300
Anomaly Detection	Data Scientist	6	730	4380	90	6540
Alerts	SW Eng.	4	365	1460	30	1940
Outgoing Integration	SW Eng.	1	60	180	20	260
Investigation	FE Eng. + DS	1	365	365	30	485
Administration	DevOps Eng.	1	180	180	20	260
Architecture	SW Eng.	3	365	1095	60	1815
Product	PM	1	-	-	365	1460
QA	Automation Eng.	1	-	-	365	1460
UX	UX / Design	1	-	-	365	1460
R&D Manager	R&D Manager	1	-	-	365	1460
Data Science Manager	DS Manager	1.5	-	-	365	2190

ASSUMPTIONS			
Platform Size (Metrics)	1,000,000	Hosting Price/Metric/Month	$0.003
Duration (Years)	4	Avg. Annual Employee Cost	$135,000

Buying Options for AI-based Network Monitoring

Monitoring solutions differ in the area of the business they are designed to monitor. The three main monitoring categories are:

Enterprise Data Monitoring Platforms

A data platform is a complete solution for ingesting, processing, analyzing and presenting the data generated by the systems, processes and infrastructures of modern digital organizations. These solutions offer network, infrastructure, IT, APM and Security and Information Event Management (SIEM) monitoring.

Network Automation Solutions

AI development frameworks for mobile networks with network visibility, anomaly detection, predictive network intelligence, and process automation utilities for network operations and customer care. These solutions are typically siloed to specific network types/layers (few data sources out of the box) and have limited usability (use cases). Integration of additional data sources is very expensive. In addition, solutions of this kind tend to be so complex that they require on-premise service delivery and platform installation.

Autonomous Network Monitoring

Autonomous network monitoring is the brain on top of existing OSS tools, giving CSPs a holistic view across domains (multiple network types, layers and services) for real time detection of service-impacting incidents. These solutions aggregate inputs from network functions and logic such as fault management KPIs, xDRs, OSS/BSS tools, performance management KPIs, probe feeds, counters and alerts for all network types and layers into one centralized analytics platform to analyze 100% of data streams and metrics, regardless of the business's original data architecture and silos.

Here is a list of some of the critical elements to consider when reviewing the right monitoring solution for you:

- **Data coverage.** A monitoring solution is only as robust as the data it can cover. When streams are siloed or cannot be ingested by the solution, holistic visibility is sacrificed as well as the systems' ability to correlate across relevant metrics and dimensions.

- **Level of automation.** While monitoring is autonomously executed by ML algorithms, there is a varying degree of human intervention required to manage and oversee the solution's initial implementation and on-going performance. While some platforms still require manual baselining and correlation definition, other platforms get close to 100% hands-off monitoring.

- **Context.** Monitoring with ML enables not only to surface anomalies, but to also correlate between anomalies in different areas in order to expose the context of what is happening, and, in some cases, the cause. While Time to Detection (TTD) is exclusively determinant on the technology, Time to Resolution (TTR) can be decreased dramatically with good contextual information. While this is a critical feature, current solutions vary widely in the ability to correlate across metrics and dimensions.

- **Noise reduction.** Surfacing critical alerts while preventing alert storms, false positives and false negatives separates the monitoring boys from the men. Monitoring solutions offer different logics and methodologies for noise reduction mechanisms, opting for the sweet spot where no critical alert is silenced — but noise, and the troubleshooting associated with it, is reduced to a minimum.

- **Implementation & time to value.** As with most other data platforms, implementation and positive ROI time can vary greatly from a few weeks to a year. When time is of the essence, this is an important factor to consider.

- **Cost of ownership.** Solutions differ in pricing logic and levels, hosting prices, and scaling costs. Most monitoring solutions can have high costs as you scale due to data volume or host-based pricing models. When considering TCO it's also important to examine the solution's integration with existing monitoring solutions, which can reduce secondary monitoring costs.

Comparison of buying options

	Autonomous Network Monitoring	Enterprise Data Analytics/ Monitoring Platforms	Network Automation Solutions
Vendors	Anodot	Splunk, Microfocus, CrunchMetrics, Guavus	TUPL, Uhana by VMWare, Elisa Automate
Data coverage	Anodot aggregates inputs from network functions and logic such as fault management KPIs, xDRs, OSS/BSS tools, performance management KPIs, probe feeds, counters and alerts for all network types and layers. • No limits on data	Network, Infrastructure, Application Performance Security • Limits on data	Limited to mobile networks and uses streaming telemetry from 4G/LTE, 5G/NR radio access networks (RAN), and the mobility management entity (MME) in the LTE and 5G NSA networks.
Level of monitoring automation	• Autonomous real-time detection and alerting on all data • Auto-learning of seasonality • Autonomous learning of metric behavior • Automatic selection of optimal model from over 20 algorithms • Sequential adaptive learning and feedback • Fast detection time, including detection of small and slow leaks	• Real-time detection and alerting only on manually created alerts • Anomaly detection and specific criteria needs to be manually enabled for each alert • Limited selection of anomaly detection algorithms require manual selection on alert creation • Limits on memory, hardware and number of entities monitored • Pre-defined seasonality selection	When creating an anomaly detection alert, the engineer has to start with a known issue, provide related metrics and begin the process of training the model and providing feedback and additional data, once the model has reached a threshold confidence level, a rule is created and it can be deployed as an alert in the system.
Context	Fully automated, comprehensive event and metric correlation, and root cause analysis via a patented correlation engine	Manually predefined event correlations using time and geographic location, transactions, sub-searches, field lookups, and joins	Correlations between metrics are based on known existing patterns in historical data that an engineer must designate. Engineers need to build out a library of known root causes and solutions that can be used as recommendations.
Noise reduction	Advanced alert scoring, alert reduction, and false positive reduction mechanisms	Manual data enrichment and alert deduplication for noise reduction	Alert reduction via user-defined rule generation and maintenance
Implementation and time to value	Anodot can be implement within 2-4 weeks and can deliver value within the first 30 days	Enterprise data monitoring platforms are very complex to implement, typically taking a year or more before they can deliver value	Network Automation tools typically take 6-12 months to implement and another 12 months before they can deliver value to the organization
Total cost of ownership	Low: Metric-based pricing regardless of data granularity, no limits on data and hardware. Anodot works seamlessly with existing monitoring solutions to improve the quality of alerts generated and reduce secondary monitoring costs	High: Each area of the monitoring stack is typically sold as a stand alone product, which is priced and implemented separately. Sprawling costs with increase in data types, volume and hosts. Can ingest data from other monitoring tools but this incurs additional costs.	High: Implemented for specific use cases. These complex solutions are fully managed and operated by the vendor and offered as a Managed Software as a Service (SaaS) with expensive integrations.

Build vs. Buy Comparison

While viewing the build vs. buy options for Autonomous Monitoring side by side, some key points come to light:

The complexity of autonomous monitoring makes it especially hard to build. That's why generally, build scenarios can only be applicable for very large, innovative CSPs with dedicated R&D and dev teams.

The complexity of autonomous monitoring makes it especially expensive to build and maintain. Estimates show that developing and maintaining a data-driven enterprise software application can cost upwards of $4 million USD per year. Given that real-time monitoring is at the cutting edge of computer science, your project might greatly exceed that figure.

Solution Maturity	Duration (Months)	Production (Man Days)	Maintenance (Man Days)	TCO (Man Days)	TCO (USD)
Basic	12	2,565	2,009	7,551	$4,052,836
Intermediate (+Basic)	24	5,045	4,034	16,541	$7,377,904
Advanced (+Intermediate)	48	13,250	6,169	36,656	$14,817,699

Building your own solution? Expect an exceedingly long time to value. To recap, the duration of building an anomaly detection and monitoring solution is as follows:

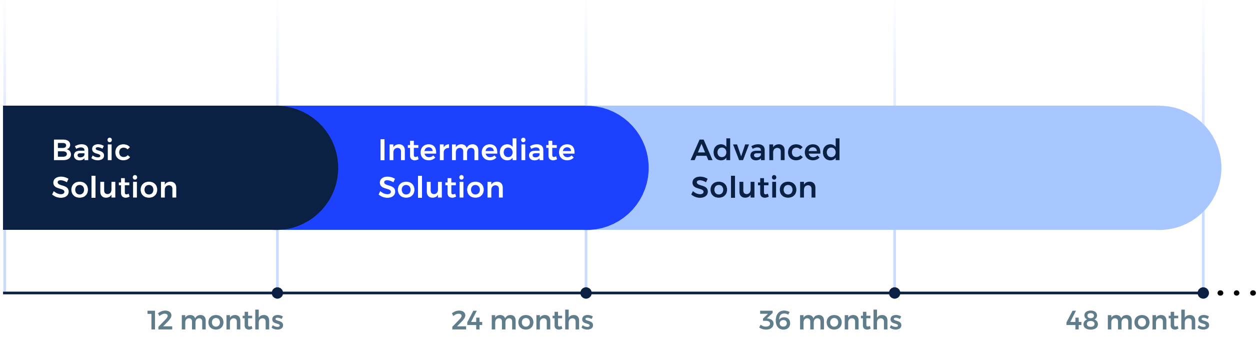

You will struggle to achieve the scale and performance of best of breed dedicated solutions. Dedicated solutions, like Anodot, run in scale that can qualify as continuous monitoring of the 5G RAN and Core cloud native infrastructure, regardless of its extensive number of devices and layers. Even after investing the above resources, the final solution's performance will usually fall behind that of dedicated solutions:

1. Basic home grown solutions usually struggle to scale with the business. As business complexity grows, the number of metrics to monitor may multiply very quickly, and you essentially "scale out" of the feasibility of implementing your own outlier detection approach.
2. More mature home grown solutions (Intermediate and Advanced) will struggle to achieve the results of dedicated solutions built on the cutting edge of monitoring science. Under par results will inevitably translate into:
   - Less accurate detection
   - Longer time to detection and resolution
   - More noise

Most home grown AI solutions fail. According to Gartner, 85% of AI projects ultimately fail to deliver on their intended promises to business. High failure rates of bringing AI to production and keeping it on the rails hinge on multiple factors. Most prominent are the inherent complexity of AI solutions, multi-faceted data challenges, and production challenges related to both maintaining model confidence and scaling the solution.

> *According to Gartner, 85% of in-house AI projects ultimately fail to deliver on their intended promises to business*

With autonomous solutions, customization comes with the territory. While customization is a key driver towards the "build" route, it is actually better achieved by the advanced machine learning algorithms built into mature solutions. This is doubly true in case of Unsupervised ML systems, that instantly adapt to any data architecture, business logic and signal type out there.

Fast and seamless integration and implementation level the playing field for time to value. For many companies, the typical exceedingly long implementation time of monitoring solutions is a trigger for building their own. In the case of some autonomous monitoring solutions this rejection is irrelevant, since they can be up and running within weeks.

Conclusion

AI-based monitoring and anomaly detection is the key to ensuring that businesses can keep pace with the high level of service required for mission-critical applications. Early, contextual detection is a basic requirement for speedy resolution. AI-based monitoring creates more visibility and provides the agility needed to mitigate the outages, blackouts, glitches and issues that do and will happen.

That's why smart communications service providers (CSPs) are investing in AI. By cutting time to detection, reducing false alarms and alert storms, and providing the context for the shortest time to resolution, AI solutions enable CSPs to ensure availability and reliability, deliver more business value, and stay ahead of the competition.

Compared to manual, dashboard-based monitoring systems, ML enables unprecedented scale, accuracy and speed. It enables today's telecom engineers to handle, manage, optimize, monitor and troubleshoot multi-technology and multi-vendor networks. Machine learning enables CSPs to move from reactive problem solving to proactive monitoring and learn more about what is happening across their networks before any minor issues escalate into bigger problems.

> *AI solutions enable CSPs to ensure availability and reliability, deliver more business value, and stay ahead of the competition*

AI enables the transformation of traditional network and service operations towards automation and intelligent operations through three crucial steps that can only be achieved by applying cutting edge machine learning: anomaly detection, correlations and root cause analysis, and, finally - remediation.

There is a wide range of monitoring solutions on the market, and adoption is often correlated with the organization's maturity level. IT monitoring is implemented very early on, and APM usually follows closely. Mature organizations require the monitoring abilities that only Autonomous Business Monitoring can deliver by monitoring and analyzing 100% of the business's data, including complex signals influenced by volatile parameters such as seasonality and human behavior.

CSPs opting to build their own solutions need to understand the costs, staffing challenges, and potential pitfalls to ensure that any home-grown solution not only serves its intended purpose, but also provides a comparable return on investment. While the promise of open-source AI-based solutions is great, so are the challenges associated with implementing them at scale, and, especially, of moving beyond the proof of concept to production - an endeavor which only a fraction of companies building their own platforms successfully achieve.

To start with AI-based monitoring fast it's critical to accelerate time to value by reducing prolonged development and implementation times. In the case of monitoring solutions, reducing time to value works in two channels: less resources are spent on building a solution, while implementing a monitoring solution without delay dramatically cuts costs on faster detection and resolution of incidents that are already happening right now.

Anodot Autonomous Network Monitoring

Anodot's autonomous network monitoring platform is the brain on top of the OSS, giving CSPs a holistic view across domains for real time detection of service-impacting incidents. By monitoring network performance and service experience in real-time, Anodot provides lightning-fast detection of the incidents that impact your customers and bottom line so that you can ensure customer satisfaction, minimize revenue loss and reduce time to repair.

Anodot collects and analyzes data across the entire telco stack, including all data types from all network types, layers and domains at scale. All metrics are actively monitored, enabling CSPs to achieve full visibility of service degradation incidents. Anodot's patented correlation engine correlates anomalies across the network for holistic root cause analysis and the fastest time to resolution, leading to improved network availability and customer experience.

Anodot is completely autonomous. There's no need to define what data to look for or when, no manual thresholds to set up or update. New use cases can be added on the fly, and no monitoring maintenance is needed even as the network configuration changes. That's why Anodot delivers the shortest time to value for AI, and maintains that value over time.

Anodot is built for business users — the platform is ready to use with no data science required. It is easily integrated with any type of data sources, and just as easily applied to new services (IoT, VoLTE, IPTV). CSPs use Anodot to build resilience into their networks. Anodot enables CSPs to reduce the number of alerts by 90% and shorten their Time to Resolve by 30%. This helps operations and NOC teams become proactive in their ability to identify service degradations and outages, improving network availability, customer experience and operational efficiencies.

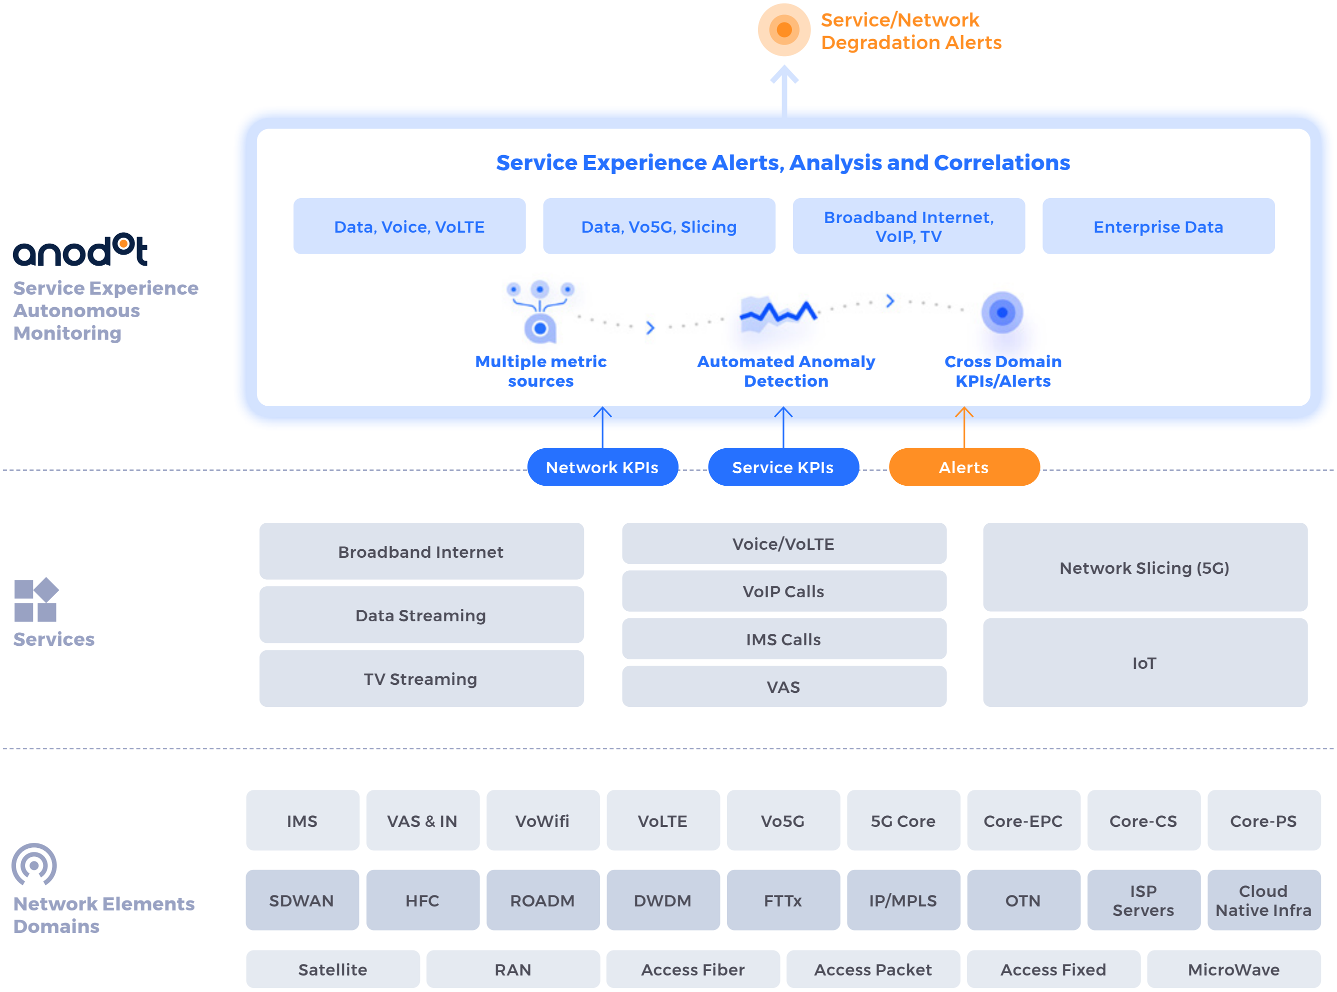

Service/Network Degradation Alerts

anodot Service Experience Autonomous Monitoring

Service Experience Alerts, Analysis and Correlations

Data, Voice, VoLTE

Data, Vo5G, Slicing

Broadband Internet, VoIP, TV

Enterprise Data

Multiple metric sources

Automated Anomaly Detection

Cross Domain KPIs/Alerts

Network KPIs

Service KPIs

Alerts

Services

Broadband Internet

Data Streaming

TV Streaming

Voice/VoLTE

VoIP Calls

IMS Calls

VAS

Network Slicing (5G)

IoT

Network Elements Domains

IMS, VAS & IN, VoWifi, VoLTE, Vo5G, 5G Core, Core-EPC, Core-CS, Core-PS

SDWAN, HFC, ROADM, DWDM, FTTx, IP/MPLS, OTN, ISP Servers, Cloud Native Infra

Satellite, RAN, Access Fiber, Access Packet, Access Fixed, MicroWave