


# AI eGuide: Find Unhappy Users Before they Complain

Real-World Guidance from a Leading Practitioner




# Introduction

## Velocity, Insight, and Impact.

Whether your organization is a fast mover or fast follower, the ability to increase your velocity gives you the advantage in shaping business outcomes. When the correct data and models are used, patterns can be discerned, and insight can be unlocked. With this insight comes the confidence to make decisions, communicate them clearly, to move faster. To gain leverage, one must find these data-backed fulcrums on which to pivot, but with ever-rising IT operational burdens, it often feels like systems and data are slowing you down rather than speeding you up.

Unfortunately, with increasing complexity and the sheer volume of data generated, teams are swamped rather than empowered which results in overall organizational slowdown. By embracing the right forms and interfaces to AI and ML (especially in operational workflows), a revolution is occurring where the promised benefits from richer data is finally being realised, enabling not just better communication, but rapid action, and faster progress.

Whether it's the desire to innovate, accelerate, or gain insight on the bottom line, the network fabrics your organization depends upon, support almost all workflows. With AI for IT, better operations deliver superior and faster outcomes, across the board.



### About the Author:

Donal wears many hats and believes we are all network engineers in one form or another. He consults at [Defensible](#), builds engineering testing tools at [PanSift](#), and grows community at [iNOG](#). Donal hails from a mix of engineering and security roles in telco/mobile, enterprise, vendors, and start-ups. He's previously held multiple industry certifications (including a very early CISSP) and comes from a computer science background. These days he gets most satisfaction when growing communities of practice.

– Donal O Duibhir


# Teams, Talent, and Toil

**Your people are your greatest asset. They are empowered by the devices and digital networks that enable them to work, collaborate, and provide services. Seamless and secure connectivity is the fuel that powers these interactions, but when connectivity and systems fail, the skills and costs required to troubleshoot them are unevenly distributed.**

Communication and confidence are also key for operational staff, as they not only keep the lights on but set the tone for employee and customer experience alike. By empowering operational teams to interface with the network using simple natural language (and by proactively surfacing root causes), there are faster resolutions, fewer escalations, and overall better user and customer satisfaction. Continuous and rapid improvement is empowering for all teams, and it builds confidence throughout an organization.

AI and ML (Machine Learning) can solve and streamline traditional challenges like troubleshooting networks, all while removing toil. This frees up human talent for higher-order creative problems. By unlocking this extra cognitive capacity in individuals and teams, groups can be more forward-looking and proactive.

IT systems and networks also traditionally required operational teams to grow linearly as their managed footprints grew, yet this is not the case with an AI-driven Enterprise where constant learning and optimizations occur.




Proactively identify and fix the root causes of problems using Juniper-Mist's Proactive Analytics and Correlation Engine (PACE).


# User Experience and Assurance

**Positive user experience is crucial for productivity and the sustained use of any business system or service. As edge complexity grows, your organization's success is dependent upon timely and reliable access to these systems and services.**


The user edge, especially with wireless, is one of the most dynamic and complex service surfaces for IT teams to manage.



**Multiple devices**



**Threats**



**IoT**

A user's application experience is contingent upon multiple variables and dependencies, some of which are beyond control, while many are predictable and manageable. As radio energy shifts in response to movements, roaming devices, or interference, multiple device types (be they approved or otherwise) contend for shared airtime.

Disparate client devices, dynamic radio spaces, and an upstream wired edge all combine to form the onramp to services inhabiting or transiting your network fabrics. Correlating point-in-time states for all of these variables and dependencies is a highly complex and challenging task, one that is even more difficult retrospectively. Luckily with the right data being proactively and continuously sampled, ML can surface, not just anomalies and errors, but can also prescribe or take remedial actions that enable service level expectations (SLE) to be achieved.

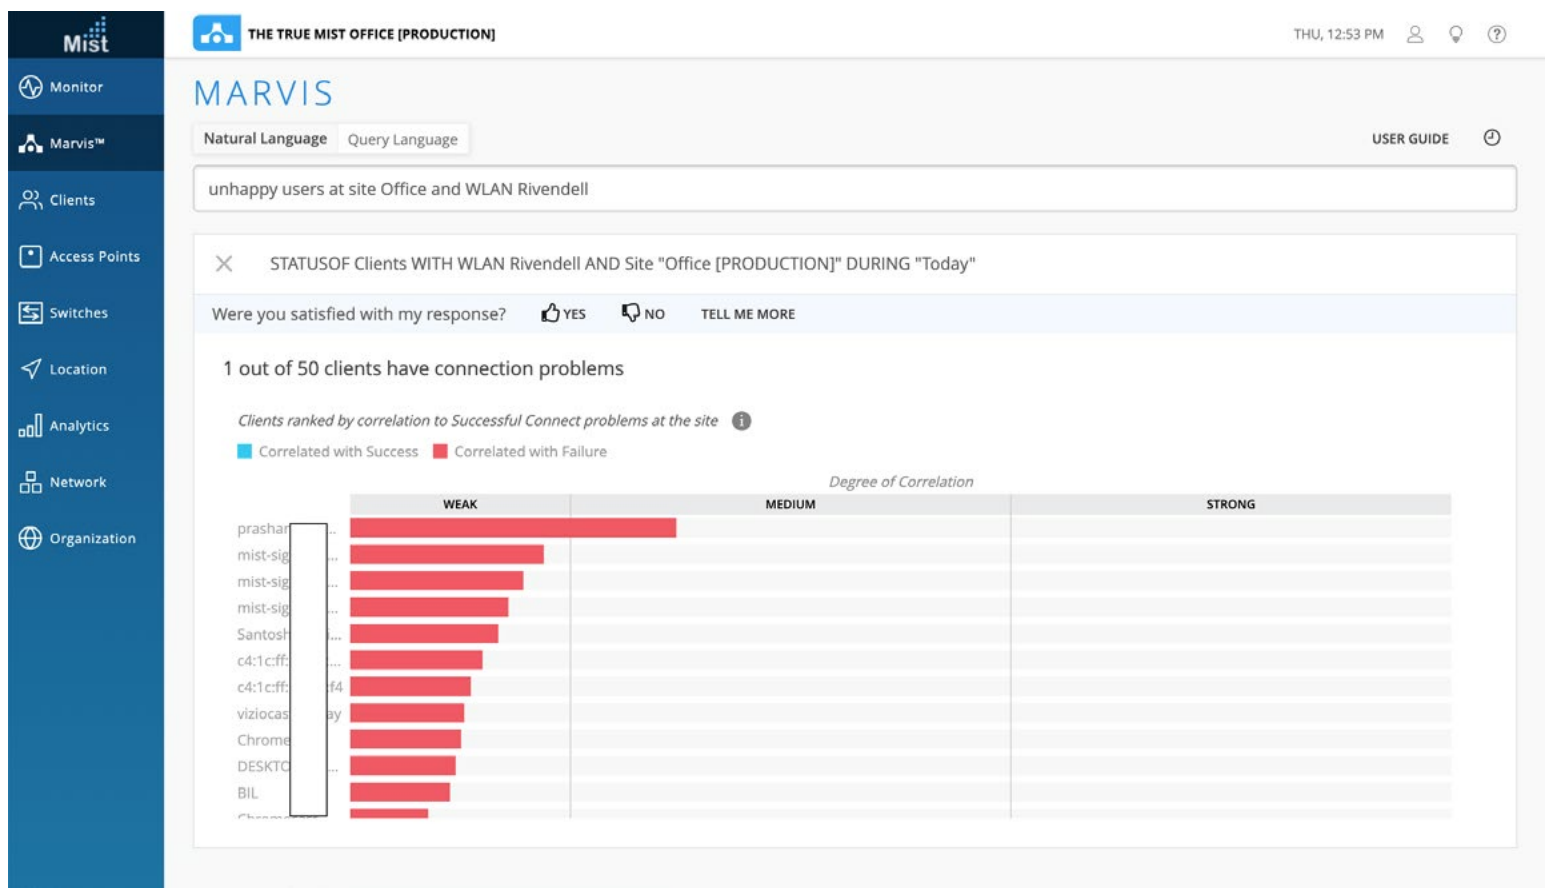
When an AI/ML engine can deeply learn and continuously consider the vast array of metrics that contribute to user experience, a more holistic and intelligent real-time view can be taken for beneficial network health and service assurance. Users understand that sometimes there are problems, but they also expect continuous improvement and timely confident communication. When ML can diagnose and empower teams with the confidence to rapidly troubleshoot (or take proactive steps), operational teams can then assure management and users that promises made will be kept.

# Interfacing with AI

User interfaces are abstractions for observing and controlling systems and data. They are intended to empower us by converting data into usable information, which then facilitates decision making and subsequent actions. As complexity rises, it often outstrips our human ability to spot patterns and correlations (even with the help of visualizations), but this is where ML becomes an operational force multiplier.

The explosion in the volume of this data from nodes and endpoints means managing these assets and their telemetry becomes increasingly burdensome. Greater systems complexity also increases the risk to service levels while hampering a team's ability to reason about end-to-end systems. ML effectively empowers greater observability while democratizing the ability to query complex interactions and events.

When an AI/ML platform can also suggest remedial actions based upon continual learning, there's a real opportunity to embark upon truly self-driving networks. Organizations are increasingly comprised of digital components which themselves are becoming more complex. When the right AI platform is allowed to learn, correlate, and present recommendations or take actions based upon events, a form of AI-Ops emerges as a new interface to complexity.




# Proactive and Intelligent Actions for Automation

The holy grail of automation is a closed-loop system with retained human oversight and veto. Trust is crucial in automation, as is demonstrating positive outcomes for it to be embraced. Where ML shines is with large datasets and forms of training where known positive outcomes can be taught and learned. Once these patterns are recognized, and appropriate models are generated, there is high confidence in answers and prescribed actions (especially so when using structured, well-labeled, and well-understood data).

Whether ML is used for inputs to automation, forms part of the decision logic, or generates triggers about which actions to subsequently take, automation loops are accelerated, leading to more consistent and predictable outcomes. By optimizing for continuous improvement, especially within domains like operations, efficiency gains allow teams to do a lot more with less. In modern markets, whether it's your product, service, or operational capability, AI/ML can accelerate workflows while also enabling the confidence to embrace further automation leaps.

# What to Look for with AI/ML and Where to Apply it

AI is not an arbitrary general level of intelligence; it is mainly applicable to specific and narrow problem domains. AI can perform tasks such as classification, clustering, image recognition, and prediction (to name a few), but it is highly dependent upon the quality of data it is fed, the algorithms used, and how training and validation are performed. Domains with well-known problems and large amounts of readily available and labeled data are often preferable, but the usability of any technology depends upon a more extensive list of factors and variables.




# How to Spot AI Washing?

Any AI/ML platform team should have clear answers and explanations for the following questions (without invoking high-level magic or snake-oil) Like any product or service, it's important to validate claims, so caveat emptor!




Where and how is the data sourced?



What data labeling, cleaning, or additional preparation is required by either party?



How does one avoid biased data and, thus, biased results?



How are models customized for my organization, site, or specific teams?




What questions can I ask of the trained models, and how?




How often are models updated or retrained?



What credentials or experience does your data science team have?



How easy is the solution to use?




How can I integrate this technology with my own systems, platforms, or APIs?


# Elasticity, Scale, and New Architectures

AI/ML is all about learning, and learning is a dynamic and continual action. Fresh data is constantly being produced by dynamic environments, so models need to be updated, retrained, and validated for any utility to be realized. This means ongoing ML uses big data sets, large amounts of compute, and is heavy on storage and network resources. These requirements place heavy demands on any infrastructure, especially when training or updating models at scale.


This leads to some interesting challenges with scale-out designs and modern approaches to system and network architecture that an AI/ML platform team should be able to clarify.



**How does the platform ingest, segregate, and secure data?**



**What is the approach for dealing with elastic demands?**



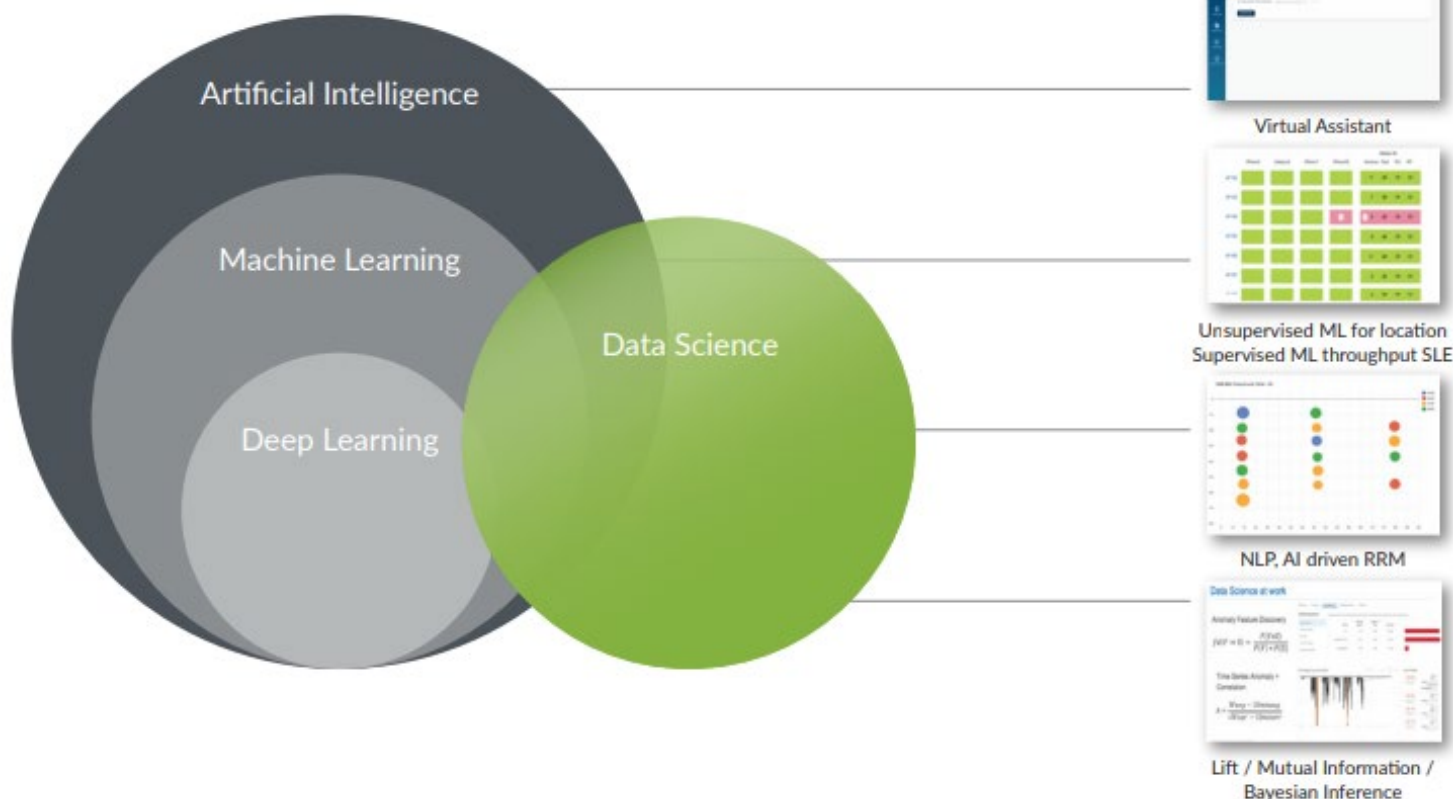
**How are rapid innovations and product features developed, validated, and deployed to production?**



# Naming Things with Confidence

Once AI/ML has been invoked as part of a solution offering, it is essential to identify the class of algorithms and frameworks in use. There are many well-known industry methods and approaches, and even if you don't recognize their names immediately, you can research them to gain confidence in a proposed solution. Once the high-level types of ML being used (such as supervised, unsupervised, or reinforcement learning) are established, you can ask about the algorithms used, methods of validation, and the confidence levels attributed to outputted predictions or results.

If there is a reluctance to name or at least broadly explain the categories of algorithms in play, it can be a signal to be more, rather than less, skeptical and to delve further. The majority of underlying AI/ML approaches are widely known amongst academic and open source communities and have clear use cases for different types of problems or challenges. Similar to the cryptographic community, transparency is considered an asset rather than a liability in AI.




Mist and Juniper Deliver a Full-Access Solution Built on Artificial Intelligence

# Deep Expertise in the Problem Space

Great AI starts with great data, but for the very best results, a team is needed who possess a deep appreciation and understanding of the problem space. Selecting the right models and frameworks is crucial, but so too is knowing how to clean the data, correctly train the models, validate the results, and then fashion them into easily usable tools or actionable outcomes.



**Does the vendor or provider have deep knowledge and expertise in the problem space?**



**Can the provider benchmark data and results against similar industry entities?**



**Does the provider use their own AI engine to improve aspects of their product or service?**

Rather than just attempting to apply arbitrary AI/ML on large data sets, there are many steps required along the way to refine the real value proposition of AI.


# How do Juniper and Mist Deliver?

## Broad and Deep

Juniper's rich and deep history in networking and security, tied to its engineer and automation friendly operating systems, mean it is uniquely positioned to capitalize on AI across IT, all the way from the campus edge through enterprise, to data center and service provider.

On the journey from manual operations, through software-defined networks, to an AI-Driven Enterprise, Juniper has already heavily invested in the key building blocks with the goal of simplifying and improving both engineering and business outcomes. With a clear vision for how AI and automation will not just enhance but enrich our working lives and those of our customers.

Beginning with AI-Ops, the benefits from richer data and AI insights can now be realized without falling foul of complexity, giving organizations and teams the space to move faster, and deliver with confidence, exceeding expectations. By embracing AI network assistants that remove toil and give operational teams the confidence to set accurate SLOs, wired and wireless assurance becomes the new table-stakes rather than aspirational goals.



# Assurance with AI for IT

**Mist is revolutionizing IT with accessible, usable, and actionable AI. Teams can now harness the power to become proactive rather than reactive. With AI/ML that understands your network, a new level of observability and controllability is realized, one that optimizes for user experience and service assurance.**

Mist has reset the industry bar for how infrastructure providers deliver on troubleshooting, lifecycle management, platform innovation, and promises relating to AI's impact on IT. AI becomes a force multiplier when human creativity is unlocked to focus on strategy and proactive pursuits, rather than tactical treadmills and reactive firefighting. Whether it's automation powered by AI (or pure AI-Ops), OPEX can be significantly lowered and/or redeployed elsewhere.

# The Mist Cloud

**The Mist Cloud stack is built on a microservices architecture. It is built from a set of distributed and independently deployable modular services that provide both horizontal and vertical scaling.**


This elastic architecture allows for rapid versioning, non-disruptive upgrades, and new feature development that ensures the velocity of innovation is maximized at all layers. Whether it's the ingestion of streaming telemetry or ongoing machine learning, the Mist Cloud platform provides the resiliency, scalability, and security required for web-scale operations across your fleet.

“ Only you can prevent network fire drills before they happen. Use AI to unlock your creative powers to reduce OpEx. ”


NTT saved over 40% in OpEx costs with Mist's AI-Driven platform.

# The AI Toolbox Applied

The Mist AI engine 'Marvis' uses a range of different ML types from its data science toolbox such as mutual information, Bayesian inference, ARIMA models, and reinforcement learning to solve for different challenges along the path to AI for IT. Mist delivers on a range of innovations including but not limited to:




**Surfacing root cause  
for user experience or  
infrastructure issues**



**Proactive assurance  
checks on the user edge  
(inc. wired health)**



**Truly intelligent RRM  
(Radio Resource Management)**



**Automatic indoor  
location calibration**

Additionally, Marvis uses NLP (Natural Language Processing) to provide simple natural language interfaces for all levels of operational teams to query network health.

By simplifying the interrogation and remediation of complex networks, it becomes easy for teams or agents to diagnose and fix issues.

# The AI Advantage

Born of deep human expertise in wired and wireless networking, Marvis learns, spots anomalies, and proactively recommends remedial actions.

It not only truly self-heals the network but offloads much of the daily operational grind and troubleshooting toil. Everything from anomaly detection to proactive RMAs (Return Material Authorization) means a true virtual network assistant is not just on-call but working 24/7.



## Corporate and Sales Headquarters

Juniper Networks, Inc.

1133 Innovation Way  
Sunnyvale, CA 94089 USA

Phone: 888-JUNIPER  
(888-586-4737) or +1.408.745.2000

Fax: +1.408.745.2100

## APAC and EMEA Headquarters

Juniper Networks International B.V.

Boeing Avenue 240  
119 PZ Schipol-Rijk  
Amsterdam, The Netherlands

Phone: +31.0.207.125.700

Fax: +31.0.207.125.701

Copyright 2020 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. In the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

**PN 7400120- 001-EN**

### Please Note:

This guide contains general information about legal matters.  
The legal information is not advice, and should not be treated as such.

Any legal information in this guide is provided “as is” without any representations or warranties, express or implied. Juniper Networks makes no representations or warranties in relation to the information in this guide.

You must not rely on the information in this guide as an alternative to legal advice from your attorney or other professional legal services provider. You should never delay seeking legal advice, disregard legal advice, or commence or discontinue any legal action because of information in this guide.

Information correct at time of publication  
(February 2020).