JUNIPER NETWORKS | Engineering Simplicity

eGuide

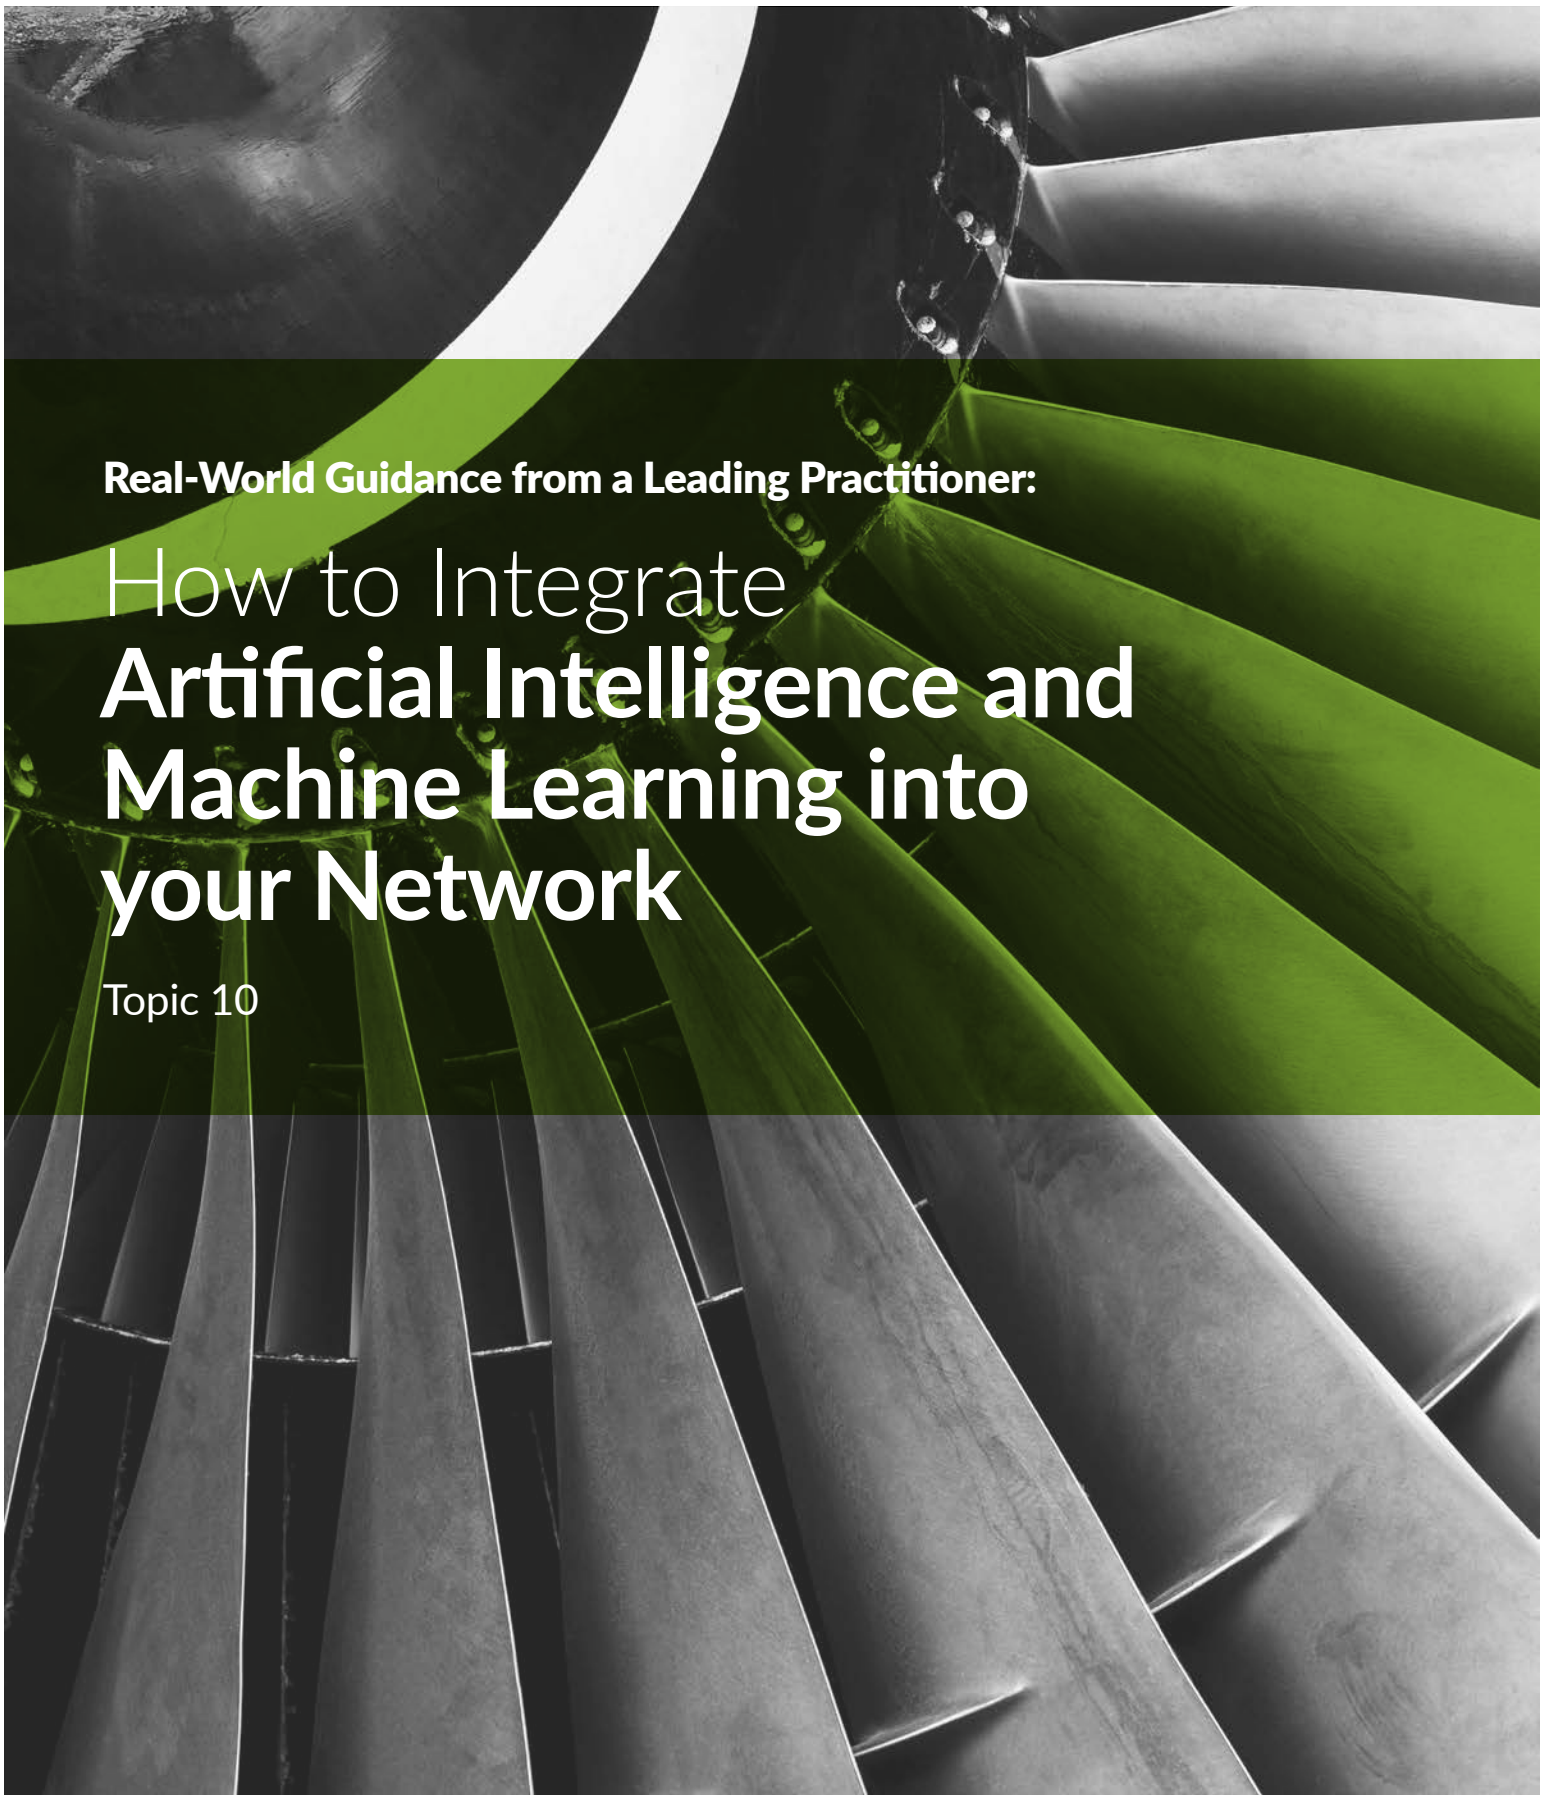

Real-World Guidance from a Leading Practitioner:

How to Integrate **Artificial Intelligence and Machine Learning into your Network**

Topic 10

August 2019

Introduction

Artificial Intelligence (AI) is a broad term that encompasses a range of technologies which mimic forms of human intelligence.

AI can enable computers to solve problems, interpret complex inputs, or even iteratively improve at tasks by learning from experience. Much of what is construed as AI today is actually a subset called Machine Learning (ML) which is underpinned by statistical models and mathematical methods.

The penetration and adoption of AI across multiple domains is rapidly accelerating. This is due to a convergence of factors which lower the boundaries to AI. From low-cost utility compute and storage, to free training and open-source AI frameworks, AI capabilities are becoming readily available to a growing set of individuals and organizations. As such, everyone from scientists and engineers to farmers and doctors are using AI to deliver improved services (or themselves depend upon systems and tools that utilize forms of AI).

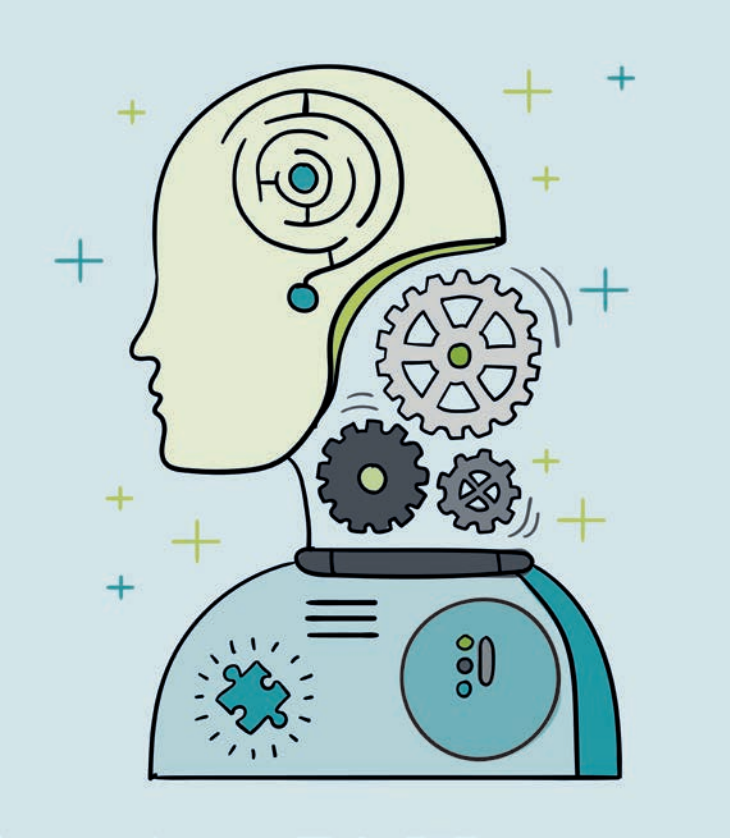

About the Author:

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

How AI works

One of the most common misunderstandings around AI is that machines are becoming self-aware, or that they can rival the creativity and understanding of the human mind. This type of advance would be coined Artificial General Intelligence (AGI or Strong AI) and is still deemed to be decades away, if at all possible.

What is currently available is a set of more narrow AI technologies that focus on specific problem domains and perform tasks such as image recognition, classification, clustering, prediction, language translation, and voice recognition. More recently, new breeds of AI are being used to train each other to actually generate feature-rich data rather than just identifying or labeling it. Each application of AI tends to embody the early goals taught to it, so avoiding early bias in training is extremely important for results not to be skewed.

But what parameters do AI algorithms optimize for and what types of specific real-world applications do they have? The applications of AI abound and are making inroads into network engineering and campus edge services with greatly improved:

Interfaces
(Including Natural Language Processing)

Operations
(Including Spatial-Temporal Correlation for Root Cause Analysis)

Anomaly Detection

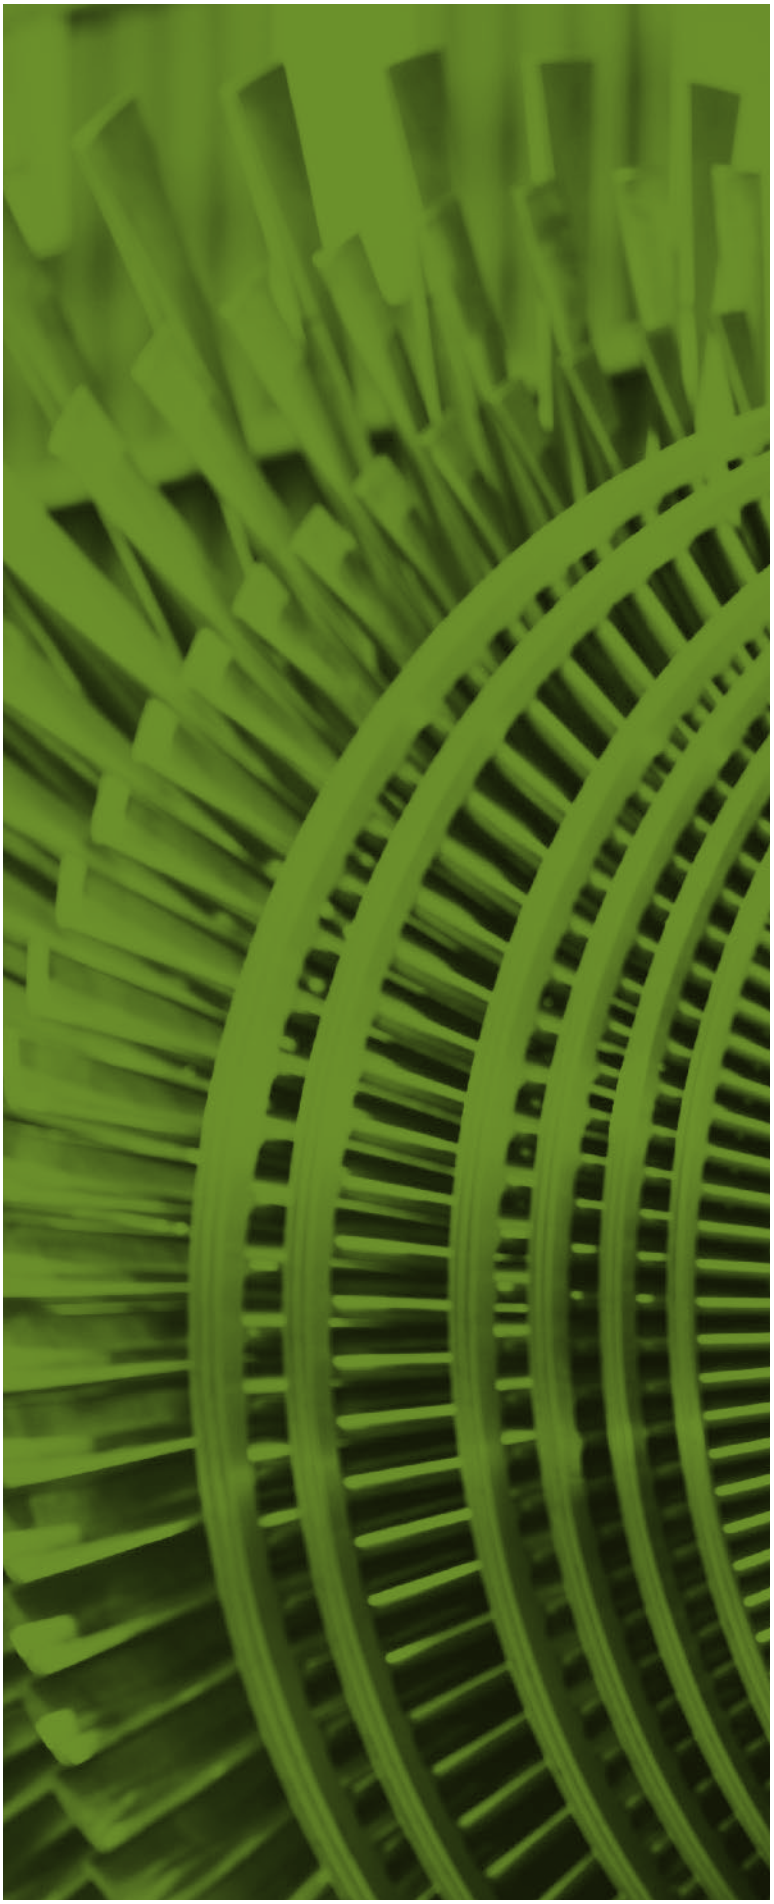

How machines learn

Many of these exciting advances depend upon a form of AI called Machine Learning (ML), and more specifically one of its subsets called Deep Learning.

ML is not just one of the fastest-growing AI domains, but it often underpins advances in other AI technologies. ML can be applied across a range of either structured or unstructured data, but to understand what it is, this quick definition is useful:

> **"A computer program is said to learn from experience E with respect to some task T and some performance measure P, if its performance on T, as measured by P, improves with experience E."**
>
> Well-posed Learning Problem, Tom Mitchell, 1998

Some high-level categories of ML include:

Supervised learning
(using labeled data)

Unsupervised learning
(using unlabeled data)

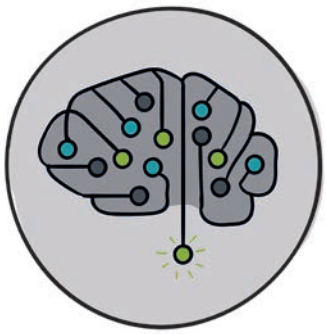

Reinforcement learning
(using rewards)

Each requires some form of either training, validation, or signaling to operate. The raw material for ML is indeed data, but it's the application of specific algorithms that model and minimize for conceptual cost functions that do the real heavy lifting. Common use cases of ML perform the regression, prediction, classification, and clustering of data, and can even control systems in real-time.

Machine learning for security

In a network or security context, the data used by ML is often some representation of network state or telemetry.

This data frequently contains labels such as protocol types or key-value pairs which can help to expedite ML efforts. The data, whether consumed at rest or via streaming telemetry, can be used to make operational decisions, initiate optimizations, trigger workflows, or detect real-time anomalies. By closing OODA loops **(view guide 7)**, ML can help drive faster response times and enact more efficient use of resources.

In this age of 'big data', it has become increasingly difficult for security analysts and IT operations teams to find the proverbial needle in the haystack, let alone correlate complicated issues across multiple giant haystacks. The scale, complexity and criticality of modern networks mean that engineers require better context and data mining capabilities to ensure rapid incident detection and integral service restoration. This is especially true for dynamic and fluid fabrics such as WLAN where Radio Frequency (RF) and user experience is a moving target.

Once trained, ML algorithms and their associated parameters can assist with tasks such as anomaly detection and complex root cause identification. ML can also become an even greater force multiplier when trained across large datasets from discrete but similar production networks.

Thinking machines

Artificial neural networks (ANNs) are a form of machine learning that tries to mimic aspects of the human brain.

They have enjoyed a resurgence in popularity since the 1980s and 1990s due to the diminishing cost of specialized compute and increased success in healthcare and environmental monitoring. They are still somewhat computationally expensive to employ and train, but are becoming more accessible due to open-source frameworks.

These ANNs model a virtual network with a full mesh between layers of nodes (neurons). Each conceptual edge connection has weights that help determine subsequent output values of each node per layer. Results can then be observed as values in the final output layer. For deep learning to occur, there are multiple hidden layers between the original input layer and the final output layer which can use either a forward or backward propagation of values. The neural network learns by amending weights and biases across the network to minimize a cost function that best matches training data.

In the network and security world, we can leverage telemetry data as it often already contains useful labels or named values. This added context is indeed useful but does not solve all training challenges, nor does it automatically elicit answers from a neural network. More steps are required to pose the right questions and determine useful answers. This is where the true value of domain-specific platforms shine when employing deep learning combined with simple interfaces and actionable outcomes. One such example is in the 802.11 WLAN domain, where models can be trained on Information Element (IE) fields and observed user device performance to determine how to provide the best possible user experience.

Your data foundation

Once networks are observable, data can be acquired about internal or external states. But how much data is enough, and what data types are actually required for AI/ML initiatives? Does your current instrumentation lack enough fidelity for any meaningful answers to be reached?

If so, there are some options or steps to think through:

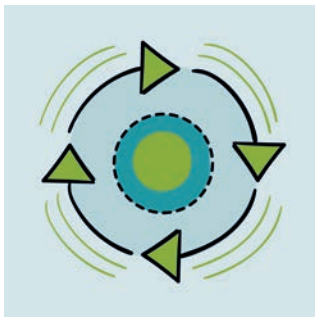

What outcomes or processes are you trying to improve?

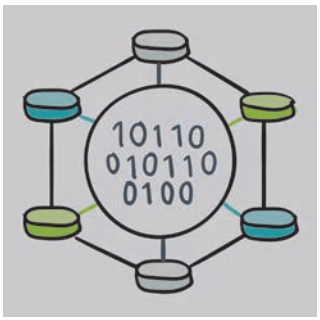

Have you currently got access to the right data, where, and how?

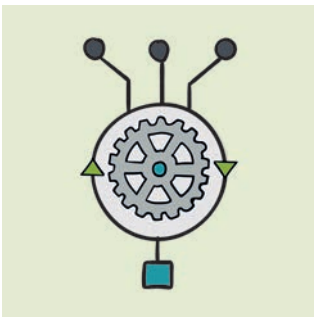

What algorithms are used (either by you or third-party services) to answer what questions?

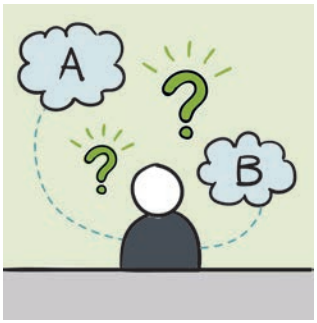

Is the data readily consumable or must it be processed first?

How much work is involved in data acquisition, cleaning, labeling and training?

Can you leverage any third-party platforms, products, or services?

It may be that you already have a data science (or business analytics) team with ML capabilities that can help guide you. But whether you build or buy, it's still imperative that you understand what can be instrumented and how the data is going to be used.

The AI difference

As engineers, we struggle to keep pace with ongoing changes and endeavor to make better evidence-based decisions (all the while attempting to eliminate toil). We seek deeper insights from simpler and more natural system interfaces, yet the volume and complexity of data keeps growing.

Tied to this, the expectations of a new generation of engineers have been shaped by consumer innovations such as voice-activated mobile devices, virtual assistants and smart home automation. Fortunately, new breeds of AI-powered analytics and assistants are taking shape to help us improve more than just our precision or speed when operating campus networks.

Aggregated data sets and deep learning are powering a renaissance in real-time system optimizations, as well as interfaces that provide better and faster operational outcomes which also lead to overall greater efficacy. Many do not realize that the ability to query a production network using natural language already exists today. Operators can diagnose issues or uncover potentially service impacting anomalies by asking simple questions while still being able to delve into high fidelity data when required.

In the battle to make sense of growing complexity, we are turning to machine learning for everything from predictive analytics to autonomous workflows and virtual assistants. That way, we can ensure we maintain our edge and are able to scale to support the soaring numbers of devices and applications that require reliable connectivity and predictable performance.

Checklist: Integrating AI and ML into your Network

Define your data across access, location and whether it's labeled or unlabeled.

Gain an understanding on how data will be used and what can be instrumented.

Identify where ML can deliver operations, insights and better outcomes.

Assess whether you can utilize in-house ML skills, or whether upskilling or external resource is required.

Consider the extent to which ML can assist your security team in incident detection and service restoration?

Continue reading the series

Choose from the other topics available in this series to find out more on how to architect your campus network:

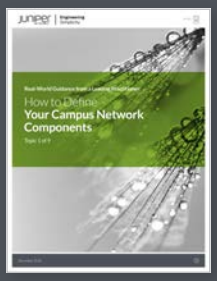

Topic 1: How to Define Your Campus Network Components

Read Next

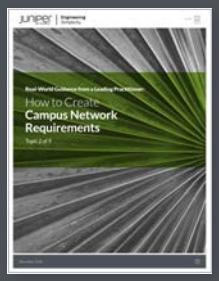

Topic 2: How to Create Campus Network Requirements

Read Next

Topic 3: How to Define a Simple Campus Network Security Model

Read Next

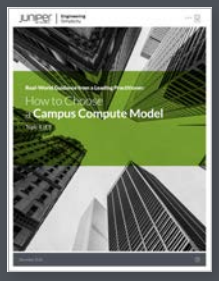

Topic 4: How to Choose a Campus Compute Model

Read Next

Topic 5: How to Accommodate Mobility on a Campus Network

Read Next

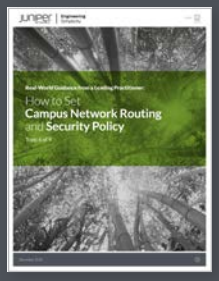

Topic 6: How to Set Campus Network Routing and Security Policy

Read Next

Topic 7: How to Build Observability into a Campus Network

Read Next

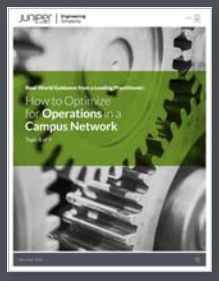

Topic 8: How to Optimize for Operations in a Campus Network

Read Next

Topic 9: How to Develop a Smart Building with IoT

Read Next

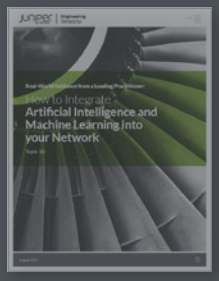

Topic 10: How to Integrate Artificial Intelligence and Machine Learning into your Network

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

PN 7400112-002-EN

JUNIPER NETWORKS | Engineering Simplicity

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 (August 2019).