



Real-World Guidance from a Leading Practitioner:

How to Define a **Simple Campus Network Security Model**


Topic 3

Introduction

Security has dependencies upon a wide range of factors and lacks standardized high-level risk metrics. Rather than make one single team or person wholly responsible, it helps to embrace security as a distributed quality and mindset. There may still be one single entity who has the authority but an entire organization shares responsibility for protecting its assets, services, employees and customers.

Campuses represent some unique challenges, not least of which are IAM (Identity and Access Management), wired and wireless access edges, and the provision of Internet access to disparate groups of users.

Often large populations of unmanaged devices are expected to connect to the access edge and utilize both campus and multicloud compute resources. Whether it's an infected printer, laptop or mobile device, the spread of malware is an ongoing risk.



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

Risk, Assumptions and Principles


In considering security, one must consider risk. To reason about risk, one must set boundaries and conditions such that the scenarios being considered are finite and tangible.

Difficulty can arise with attributing value to digital objects and services (especially in a networked world). One approach is to begin with a high-level abstraction and classification model that allows multiple teams to reason about data security, assets, the network topology and their associated risks.


Create designs with defined boundaries to maximize reliability and prevent failure propagation. Facilitate secure data transport across many of these boundaries for there to be utility. While exploring concepts of 'closed' versus 'open', each architectural approach is found to have its own merits dependent upon context.

A closed network approach can complicate things like service discovery and increase administrative overheads but is generally considered more secure. 'Closed' implies greater segmentation and stronger boundaries. An open approach can facilitate more rapid innovation and interconnection but can be deemed weaker and more vulnerable to bad actors (and propagation of certain failure types). Should malware or a miscreant gain a foothold in an open network, lateral movement is all but guaranteed.


The balance between 'open' and 'closed' networks depends upon the principles at play and how they manifest within an organization's policies and culture. Generally accepted good security practice leverages various principles, with each carrying associated costs of implementation across applications, systems and networks. These principles include (but are not limited to):




Least privilege



Segregation of duties




Defense in depth



Minimizing attack surfaces



Secure defaults



Failing securely



Compartmentalization

Zoning and Grouping

Group common logical or physical assets based upon their risk profiles, in order to talk about a node, service, or interface's trust level and requisite zoning. Some commonly used trust designations are that of 'trusted', 'untrusted', 'semi-trusted' and 'restricted'. There can also be subsets thereof based upon additional failure and governance domains.

When classifying nodes, consider whether the entity is managed by your organization or deemed unmanaged, and who or what is in control of it. Also, understanding if the flows originating from the device can be controlled, or are deemed uncontrolled, is beneficial in terms of placement within zones. Flows between entities determine the actual (or potential) transitive trusts across network topologies.

If teams can assist with basic trust based classifications of data, assets and flows, an organization can benefit not just from a team's domain-specific knowledge, but from an overall efficiency and situational awareness.

But how does each technology team consider risk without specialist knowledge? Is knowledge of vulnerabilities, exploits, threats and the capability to see all transitive trusts required?

Not necessarily, it starts with them being able to identify their own failure domains and then contributing to a wider view maintained by dedicated security teams. Additionally, with forms of security automation and zero-touch provisioning, assets can initially self-classify and then be continually updated based upon active profiling.

In this regard an initial, simple and accessible model can help by standardizing the language used around trust, dependencies and flows. This shared model and language helps to facilitate communication and discussions on complicated topics such as risk and overall security.

Although there is a trajectory towards zero trust networks, there is still much to be gained from grouping of assets, actors, agents and flows. By appropriately partitioning failure domains (and applying the relevant policies), the resiliency of an organization can be greatly improved. Combined with capabilities like observability, detection, incident response, and good Lifecycle Management, an organization's security posture is strengthened against attack.



Policy enforcement can be implemented throughout fabrics, not just at network choke points, but at the node and endpoint system too. Even when not using micro-segmentation, security automation across a whole fabric is increasingly required to inoculate, enhance response times and bolster remediation activities.



Supporting Decision Making

A good high-level security model should be easily understood by all technology teams. It should allow for a common and well-understood set of terms to be used throughout the organization.

An optimal outcome is that teams become more confident in not just suggesting but arriving at the same classifications and zoning as a dedicated security team.




Checklist:

Defining a Simple Campus Network Security Model

Risk, Assumptions and Principles



Piece together a high-level security model that's easily understood by all technology teams, so they become confident in arriving at the same classifications and zoning as a dedicated security team.



Within the model, standardize the language used around trust, dependencies and flows to facilitate communication and discussions on risk and security topics.



Create designs that leverage good security principles with defined boundaries to maximize reliability and prevent failure propagation.




Group common logical or physical assets based upon their risk profiles, in order to talk about a node, service, or interface's trust level and requisite zoning.



Capture technology teams' feedback on potential failure domains to contribute to a wider view maintained by dedicated security teams.


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

[Read Next](#)

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 2019 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 7400092-002-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 (August 2019).