Cisco TrustSec How-To Guide: Phased Deployment Overview

For Comments, please email: howtoguides@external.cisco.com

Current Document Version: 3.0

August 27, 2012

Table of Contents

Introduction

What Is the TrustSec System?

Cisco TrustSec®, a core component of the Cisco SecureX Architecture™, is an intelligent access control solution. TrustSec mitigates security risks by providing comprehensive visibility into who and what is connecting across the entire network infrastructure, and exceptional control over what and where they can go.

TrustSec builds on your existing identity-aware access layer infrastructure (switches, wireless controllers, and so on). The solution and all the components within the solution are thoroughly vetted and rigorously tested as an integrated system.

In addition to combining standards-based identity and enforcement models, such as IEEE 802.1X and VLAN control, the TrustSec system it also includes advanced identity and enforcement capabilities such as flexible authentication, Downloadable Access Control Lists (dACLs), Security Group Tagging (SGT), device profiling, posture assessments, and more.

Figure 1: TrustSec Architecture Overview

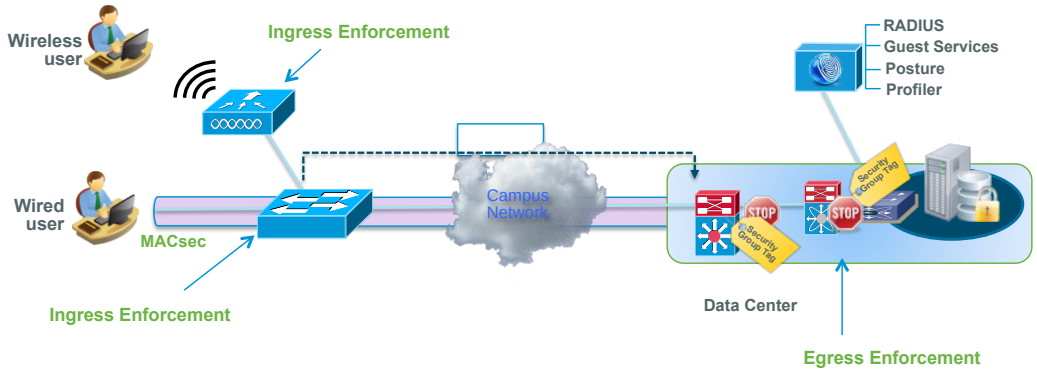

About the TrustSec How-To Guides

The TrustSec team is producing this series of How-To documents to describe best practices for TrustSec deployments. The documents in the series build on one another and guide the reader through a successful implementation of the TrustSec system. You can use these documents to follow the prescribed path to deploy the entire system, or simply pick the single use-case that meets your specific need.

Each guide is this series comes with a subway-style "You Are Here" map to help you identify the stage the document addresses and pinpoint where you are in the TrustSec deployment process (Figure 2).

Figure 2: How-To Guide Navigation Map

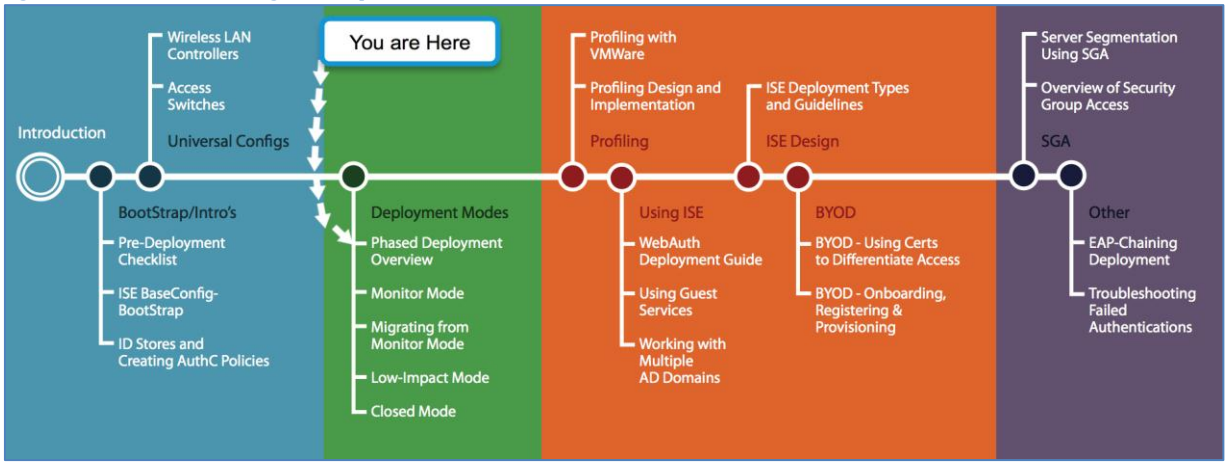

What does it mean to be 'TrustSec Certified'?

Each TrustSec version number (for example, TrustSec Version 2.0, Version 2.1, and so on) is a certified design or architecture. All the technology making up the architecture has undergone thorough architectural design development and lab testing. For a How-To Guide to be marked "TrustSec certified," all the elements discussed in the document must meet the following criteria:

- Products incorporated in the design must be generally available.
- Deployment, operation, and management of components within the system must exhibit repeatable processes.
- All configurations and products used in the design must have been fully tested as an integrated solution.

Many features may exist that could benefit your deployment, but if they were not part of the tested solution, they will not be marked as "TrustSec "certified". The TrustSec team strives to provide regular updates to these documents that will include new features as they become available, and are integrated into the TrustSec test plans, pilot deployments, and system revisions. (i.e., TrustSec 2.2 certification).

Additionally, many features and scenarios have been tested, but are not considered a best practice, and therefore are not included in these documents. As an example, certain IEEE 802.1X timers and local web authentication features are not included.

Note: Within this document, we describe the recommended method of deployment, and a few different options depending on the level of security needed in your environment. These methods are examples and step-by-step instructions for TrustSec deployment as prescribed by Cisco best practices to help ensure a successful project deployment.

Deployment Modes

Overview

Cisco TrustSec is a system of multiple Cisco® products deployed to secure the access layer. The main types of access layer include WLAN, LAN, and VPN. WLANs have Service Set Identifiers (SSIDs), which endpoints or users select and use for access. A typical network has a guest SSID that provides Internet access only and an internal SSID through which access to the internal network is provisioned. The other benefit of SSID is that the IT team can decide to deploy a more secured WLAN by setting up new SSID and directing selected users to the newly created SSID for evaluation purposes.

With the LAN access layer, a single interface has to deal with different endpoints and users; there is no concept of SSIDs in LAN switchports. So when the interface is enabled with TrustSec, it must be able to address different endpoints and users without the benefit of the SSIDs that are used in WLANs. On the wired access ports, switches or network access devices (NADs) are responsible for enforcing permissions based on credentials provided by endpoints. When 802.1X is enabled on the interface, the TrustSec system expects the endpoint to provide credentials to access the network. However, not all endpoints on the network support 802.1X. For instance, certain legacy devices on the network, such as printers, fax machines, and IP cameras, may not support 802.1X and are therefore denied access. Also, when you enable 802.1X on the switchport, the switchport may enforce a single-device-per-interface standard policy, which is likely to interfere with how the network is used by users. If switchports are enabled with 802.1X without consideration of such use cases, users with IP phones or unmanaged hubs will have trouble connecting to the network.

You can resolve all of these challenges by following a phased approach to TrustSec deployment. With a phased deployment, you can provide secure network services with little-to-no impact on end users.

As Figure 3 illustrates, the three main TrustSec deployment phases are Monitor Mode, Low-Impact Mode, and Closed Mode. Deploying Monitor Mode first allows the administrator to step through all the issues, gaining visibility into successful and failed authentications, with minimal impact to the users and endpoints. Once issues have been addressed through Monitor Mode in Phase 1, you can provide secured network access in Phase 2 through Low-Impact Mode or Closed Mode.

Figure 3 Phased Deployment of TrustSec

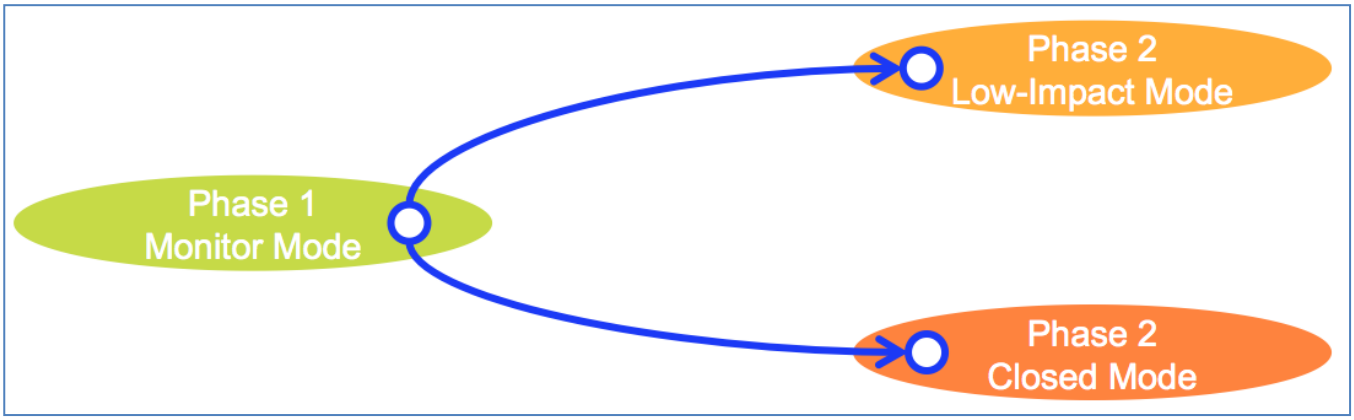

For detailed guide on each respective phase, please refer to the specific corresponding how-to document.

Phase 1: Monitor Mode

Monitor Mode works like an audit mode. Using logging data for validation, administrators use this mode to ensure that all devices are authenticating correctly, either with 802.1X or MAC Authentication Bypass (MAB). At the same time, the open authentication command used on the switch interfaces in Monitor Mode makes it possible to provide network access across your wired and wireless infrastructure, without impacting your wired users or devices. If a device is misconfigured or is missing an 802.1X supplicant, the Open Authentication feature ensures that access will not be denied and simply logged (Figure 4). When they deploy TrustSec in Monitor Mode, most organizations are surprised at what devices they find connected to the network that they were unaware of previously.

Figure 4 Monitor Mode Port Behavior Using the Open Authentication Feature

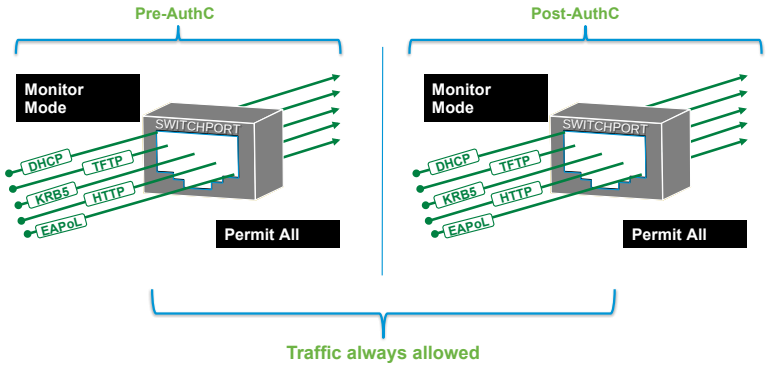

Figure 5 shows a high-level flow of authentication in Monitor Mode.

Figure 5 Monitor Mode Flow

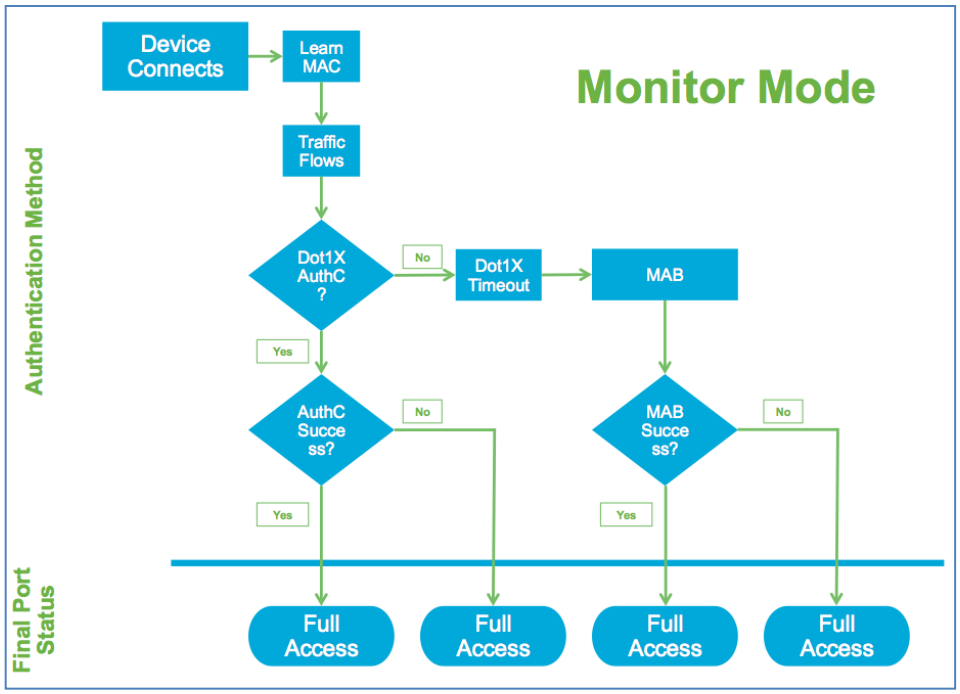

Wireless environments with 802.1X are binary (just like 802.1X was designed to be), so when a user is unable to authenticate, they simply do not get access to the wireless network. Most users can accept this behavior and are willing to find a location with a physical network connection (wired) instead. While end users are mostly willing to accept an inability to join a wireless environment, they are much less understanding when faced with a lack of access to a wired network port.

As Figure 5 indicates, Monitor Mode is a process, not just a command on a switch. The process will use a combination of RADIUS accounting packets and Open Authentication and Multi-Auth features on your Cisco infrastructure, coupled with device profiling, in order to provide visibility to the administrator into who and what is connecting to the network and from where. If a device should be authenticating successfully but fails due to a misconfiguration, the administrator will be informed based on logging data and can correct the issue without denying network access to the user. The goal of Monitor Mode is to address any possible authentication issues prior to moving to next phases.

Figure 6 Monitor Mode Process

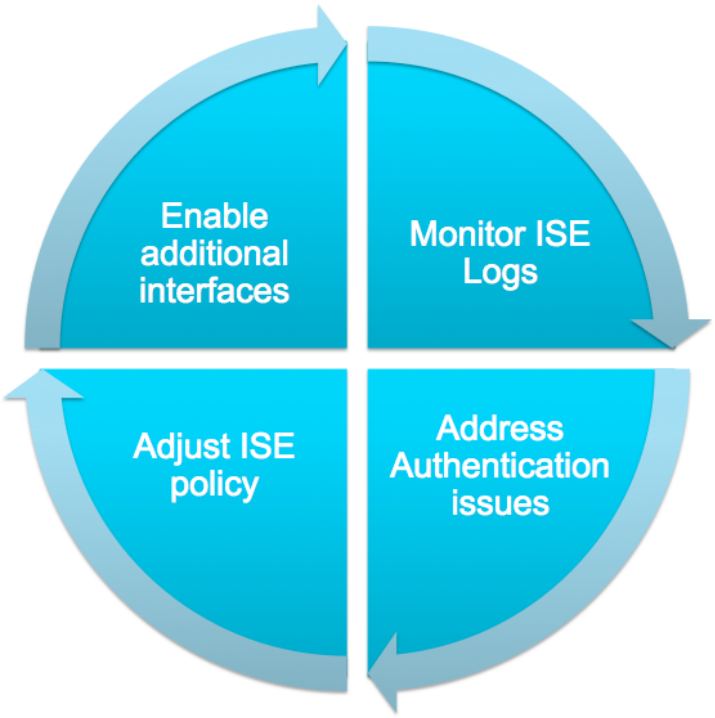

Note: It is not possible to implement Monitor Mode with wireless networks. Therefore, we will introduce wireless in the Low-Impact Mode phase.

For more information on moving from Monitor Mode to Low-Impact mode, please refer the TrustSec How-To Guide: Migrating from Monitor Mode.

Phase 2: Low-Impact Mode

In the Low-Impact Mode, we will add security on top of the framework that we built in Monitor Mode by applying an (ACL) to the switchport to allow very limited network access prior to authentication. After a user or device has successfully authenticated, they will be granted additional network access.

For example, Low-Impact Mode may be used to give any host attaching to the network the ability to use Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and perhaps get to the Internet, all while blocking access to internal resources. When a device connected to that same switchport passes authentication, a downloadable ACL (dACL) is applied that will permit all traffic (see Figure 7).

Figure 7 Low-Impact Mode Port Behavior 1

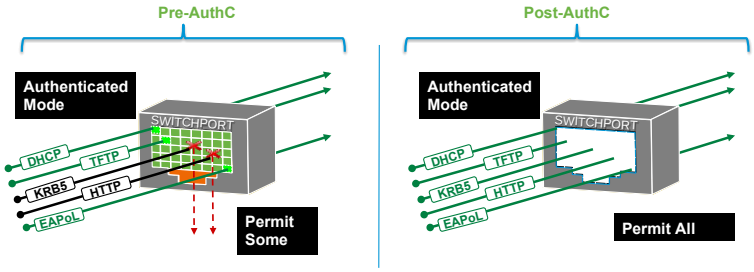

This phase continues to use Open Authentication on the switchports, while providing very strong levels of security for non-authenticated devices. However, since a limited set of traffic will always flow regardless of the authentication state of the device, this mode is a practical solution for today's enterprises because it allows regular IT operational activities to occur, such as the reimaging of workstations with Pre-Execution Environment (PXE) type solutions.

Figure 8 shows a high-level flow of authentication in Low-Impact Mode.

Figure 8 Low-Impact Mode Flow

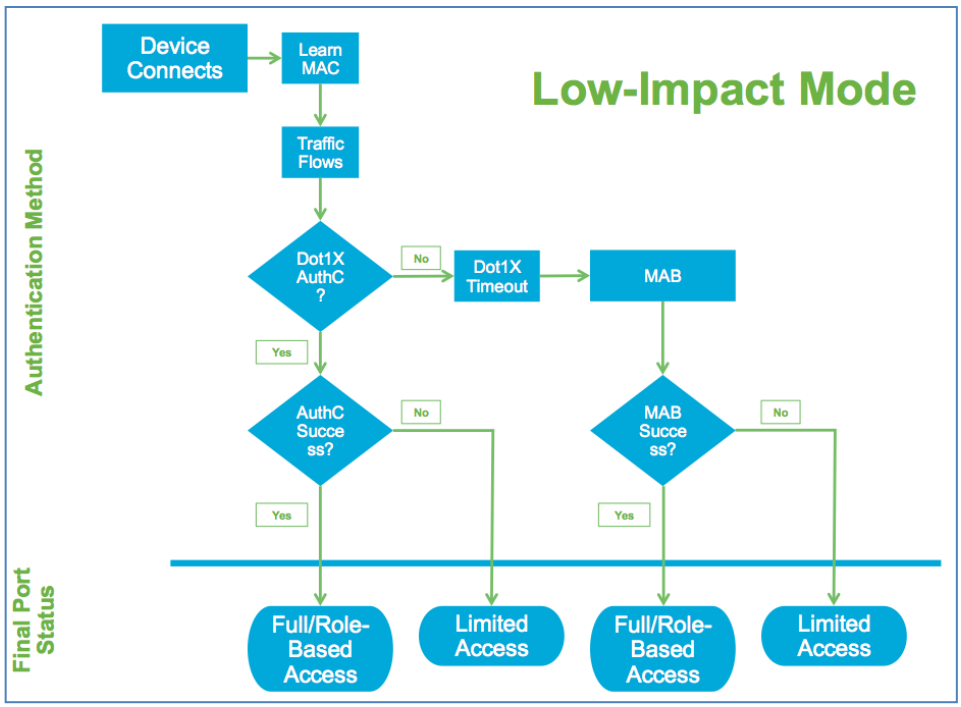

Wireless access in Low-Impact mode follows a very similar flow. A user or device authenticating to wireless with valid credentials will be authorized for full network access. This gets tightened down with additional security and specific access based on the user or device's role.

Depending on the requirements, the administrator can increase the security level by adding more granular security and differentiated access to the Network. Within the Low-Impact Mode, we replace the dACL or wireless ACL (wACL) that permits traffic with a more specific dACL or wACL that is assigned based on the user's group membership or other attributes of the user's context. Following diagram depicts granular access control based on authorization (Figure 9).

Figure 9 Low-Impact Mode Port Behavior 2

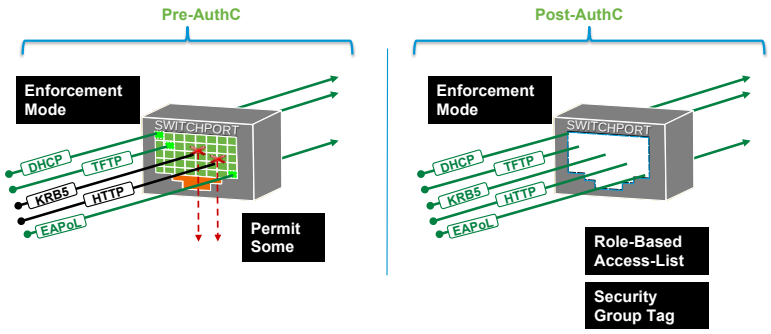

Phase 2: Closed Mode (formerly High-Security Mode)

The default 802.1X mode, which was previously called High-Security Mode, is now referred to Closed Mode. Closed Mode is recommended only for IT environments that are experienced with 802.1X deployments and have considered all the nuances that go along with it. Closed Mode should be deployed with caution (Figure 10).

Figure 10 Closed Mode Port Behavior

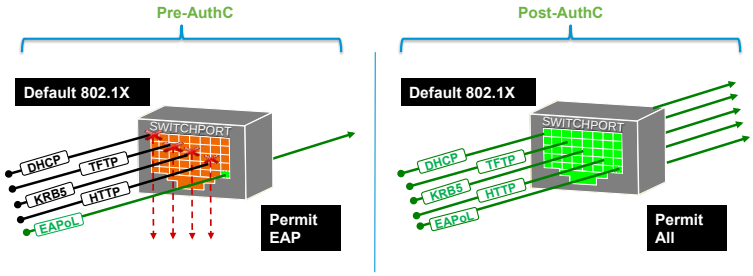

The main difference between Closed Mode and Monitor Mode or Low-Impact Mode is that interface command **authentication open** is not used. That means any traffic prior to authentication will be dropped, including DHCP, DNS, and Address Resolution Protocol (ARP) traffic. Some endpoints without a supplicant will need to wait for the interface to time out before MAB authentication starts on the interface. This could cause some endpoints to give up on DHCP process, even after MAB succeeds. To address this problem, the 802.1X timer needs to be tweaked to accommodate for such endpoints. For detailed information about Closed Mode, see the TrustSec How-To Guide: Closed Mode.

Figure 11 shows a high-level processing flow for Closed Mode.

Figure 11 Closed Mode Flow

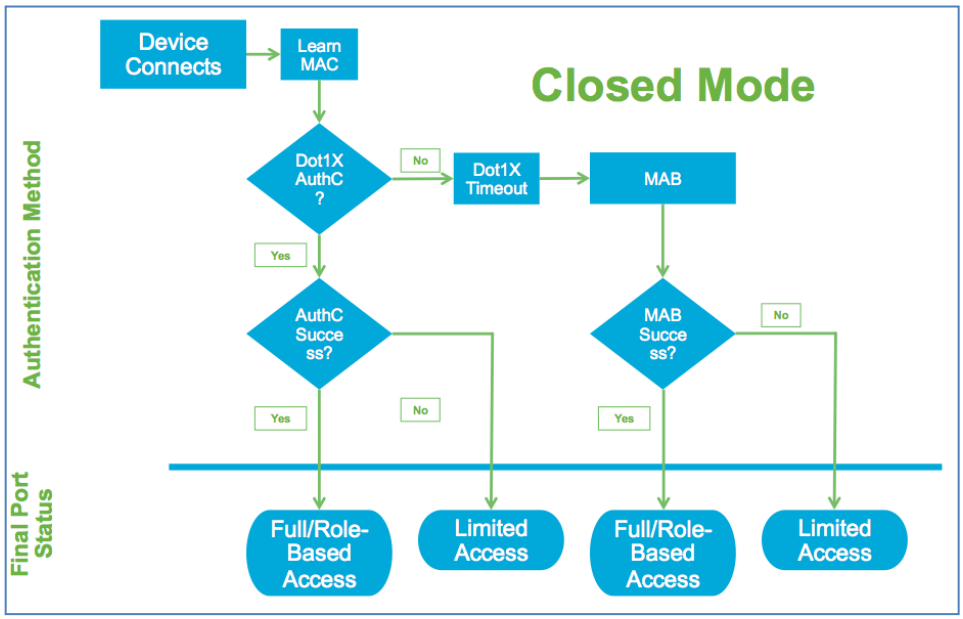

Appendix A: References

Cisco TrustSec System:

- http://www.cisco.com/go/trustsec
- http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/landing_DesignZone_TrustSec.html

Device Configuration Guides:

Cisco Identity Services Engine User Guides:
http://www.cisco.com/en/US/products/ps11640/products_user_guide_list.html

For more information about Cisco IOS Software, Cisco IOS XE Software, and Cisco NX-OS Software releases, please refer to following URLs:

- For Cisco Catalyst 2900 series switches:
  http://www.cisco.com/en/US/products/ps6406/products_installation_and_configuration_guides_list.html
- For Cisco Catalyst 3000 series switches:
  http://www.cisco.com/en/US/products/ps7077/products_installation_and_configuration_guides_list.html
- For Cisco Catalyst 3000-X series switches:
  http://www.cisco.com/en/US/products/ps10745/products_installation_and_configuration_guides_list.html
- For Cisco Catalyst 4500 series switches:
  http://www.cisco.com/en/US/products/hw/switches/ps4324/products_installation_and_configuration_guides_list.html
- For Cisco Catalyst 6500 series switches:
  http://www.cisco.com/en/US/products/hw/switches/ps708/products_installation_and_configuration_guides_list.html
- For Cisco ASR 1000 series routers:
  http://www.cisco.com/en/US/products/ps9343/products_installation_and_configuration_guides_list.html

For Cisco Wireless LAN Controllers: http://www.cisco.com/en/US/docs/wireless/controller/7.2/configuration/guide/cg.html