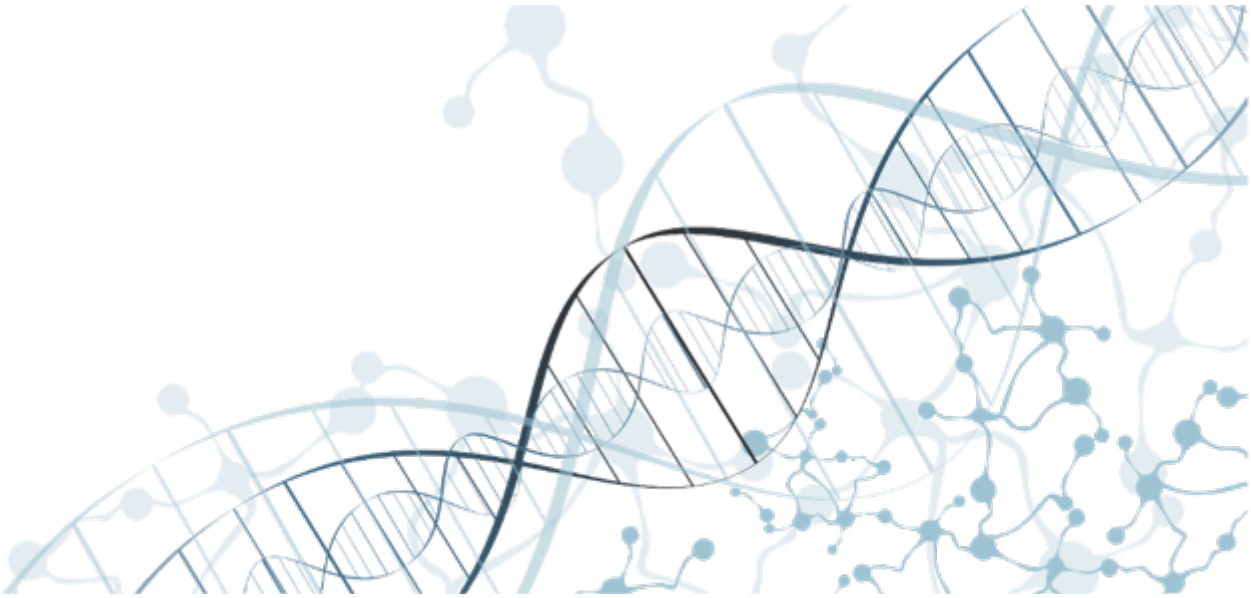

Whitelist Policy Considerations for SD-Access

Fay-Ann Lee, Technical Marketing Engineer

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESSED OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.
THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.
NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED 'AS IS' WITH ALL FAULTS. CISCO AND THIRD-PARTY SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.
IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.
Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: http://www.cisco.com/go/trademarks.
Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)
Copyright 2012 Cisco Systems, Inc. All rights reserved.

Table of Contents

Overview

SD-Access Segmentation

Segmentation within SD-Access is enabled through the combined use of both Virtual Networks (VN), which are analogous to VRFs, and Cisco Scalable Group Tags (SGTs). VNs, like VRFs, provide complete isolation between traffic and devices in one VN and those in other VNs. While segmentation can be accomplished through the use of virtual networks alone, SGTs provide logical segmentation based upon group membership. Thus, SGTs provide an additional layer of granularity, allowing you to use multiple SGTs within a single VN to provide micro-segmentation within the VN.

Figure 1. Sample use of VNs and SGTs

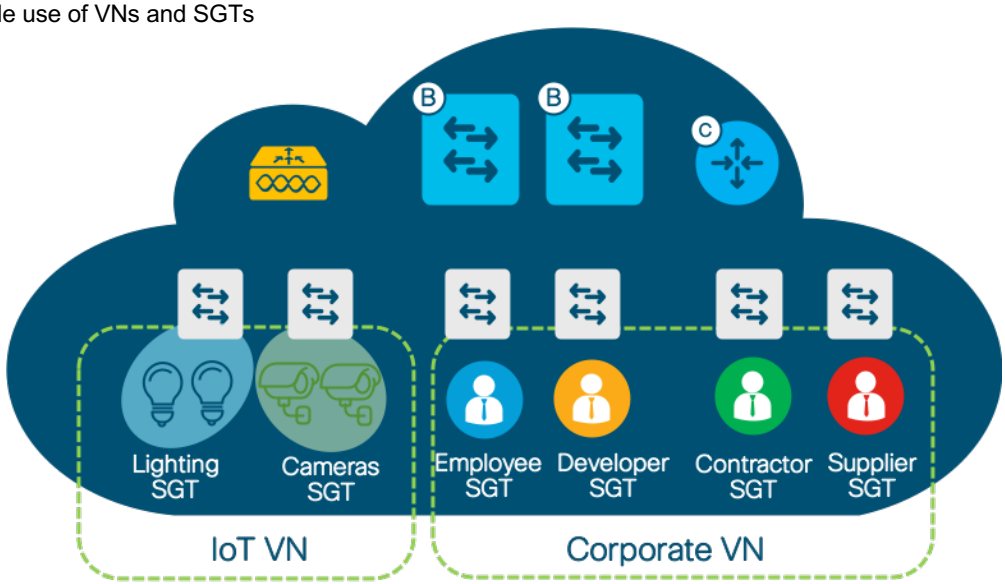

Per the above diagram, traffic from groups within the IOT are completely isolated from the groups with the Corporate VNs. However, by default, traffic between groups (SGTs) within each VN can freely communicate. For example, employees can freely communicate with the developers, the contractors, and the suppliers. In a SD-Access fabric, micro-enforcement with SGTs, also known as Group-Based policy enforcement, is used to filter communications between SGTs.

About Group-Based Policy Enforcement

A group-based policy consists of a source SGT, destination SGT, and a contract. A contract may be as simple as permit/deny ip or it may be based on Layer 4 access control entries explicitly permitting/denying specific TCP/UDP ports.

Cisco DNA Center uses a matrix view to define these policies.

Figure 2: Group-Based Policy Matrix (illustrative purpose only)

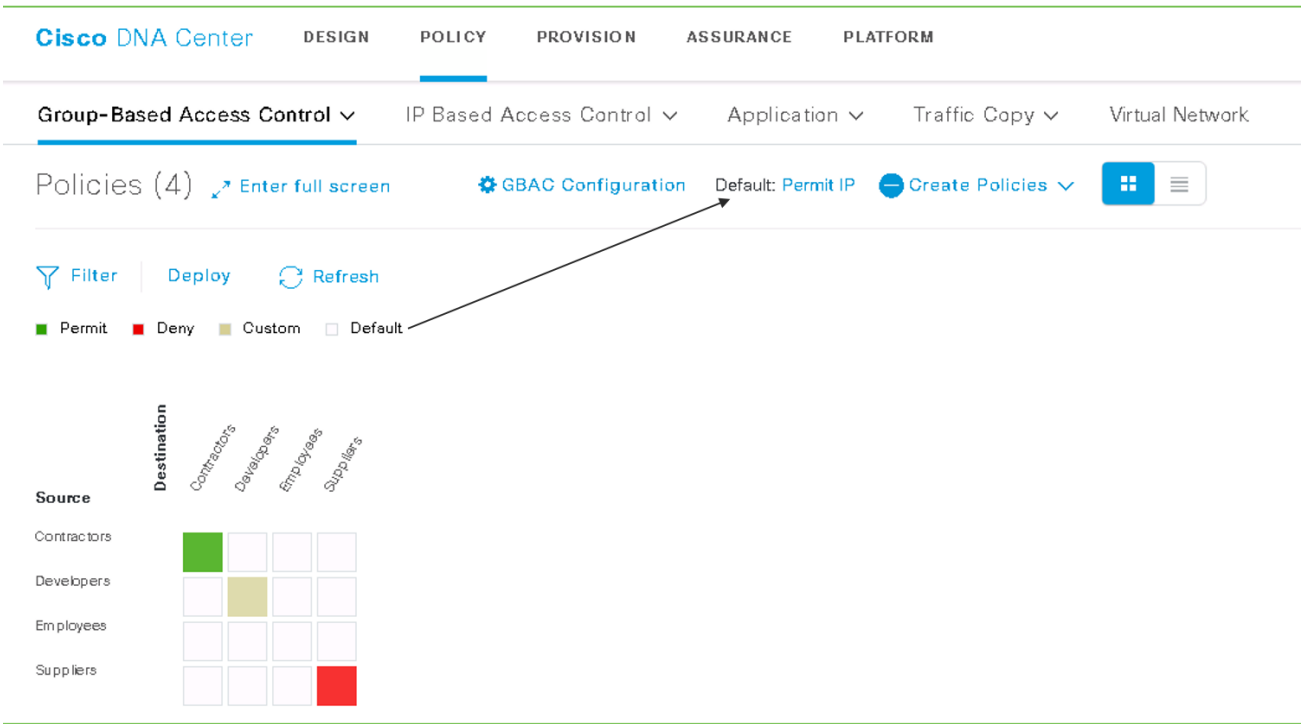

The policies are deployed to ISE and then ISE updates the edge nodes with only those policies for SGTs associated with the attached devices. Enforcement occurs upon egress where the destination is attached (more detail on egress enforcement will be covered later).

Note: Policy cannot be defined for broadcast or multicast traffic.

Policy Enforcement Models within a VN

In the world of computing and network security enforcement, there are generally two types of policy enforcement models. These are often referred to as blacklist (default allow) and whitelist (default deny). When a whitelist policy model is used all traffic is denied access, except that which is explicitly included in the whitelist policy.

Whether to choose a default permit or deny policy, it largely depends on the requirement. If the requirement is to permit most traffic on the network, then there will be a smaller number of total policies required by using a default permit with explicit deny rules. Conversely, if the requirement is to deny most traffic on the network then fewer policies will be needed by using a default deny with explicit permit rules.

Blacklist vs Whitelist Policies

	Blacklist	Whitelist

Default Action	Everything allowed	Everything is forbidden
Benefits	Specifically blacklist those applications and traffic that you are concerned about.	Write access controls just for those types of traffic and applications that want to permit
Potential problems	Reactive. Someone needs to put the problematic item in the list. For example, if it is a virus, the IT specialist will add it to the list after detection which could be too late.	Preventative. Prevents communications unless entries whitelisted items. It can stop work because a needed item is not on the list.

Blacklist

Let's say that we have SGTs for *employees*, *developers*, *quarantined_users* and *network_services*. Logically, *quarantined_users* should be restricted from almost every group with bi-directionally defined policies.

Figure 3: Group-Based Policy Matrix- Blacklist

In this example, we can see that a blacklist policy is being used because the default policy is to Permit IP. Because all traffic between SGTs is permitted by default, we need to add deny policies for the *quarantined_users* SGT and all other SGTs that have been defined. This example uses a small number of SGTs, but in a network with hundreds (or whatever), the number of policy definitions would be large.

Alternatively, with a whitelist approach there would be fewer policies to define, as shown in Figure 4.

Figure 4: Group-Based Policy Matrix-Whitelist

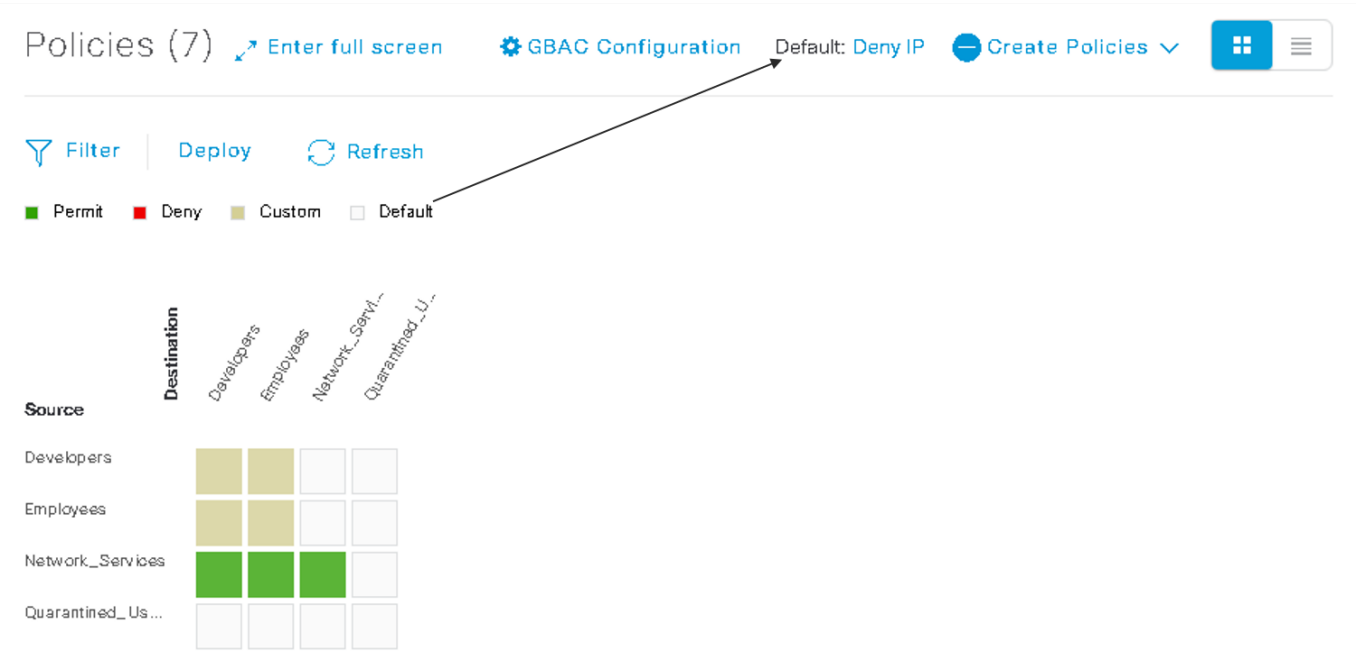

In this example, we can see that a whitelist policy is being used because the default policy is to Deny IP. Because all traffic between SGTs is deny by default, so by default quarantined_users can't communicate with anything. This example uses a small number of SGTs, but in a network with hundreds (or whatever), the number of policy definitions would be large.

SD-Access Lab topology for test environment:

Figure 5: Test Topology

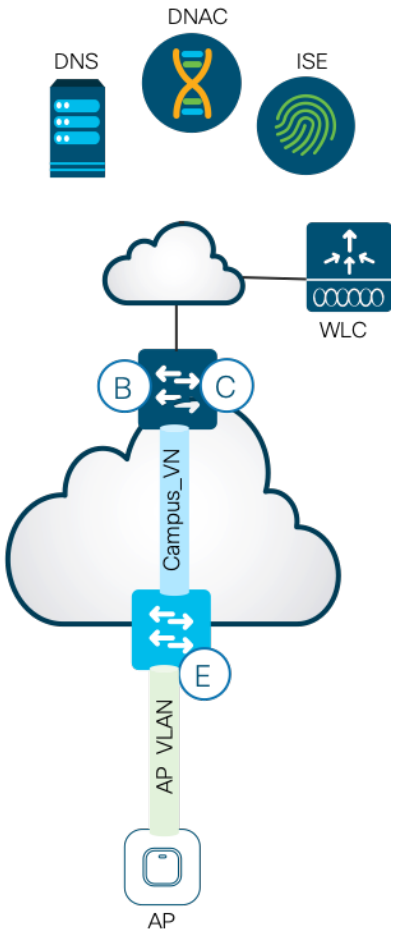

Code versions:
DNA Center 1.3.3
Catalyst 9K running 17.1.1+

Fabric Edge Device

As a SD-Access best practice, ISIS is the recommended protocol for underlay traffic in a SD-Access network. However by design, SGT based enforcement blocks broadcast traffic. Therefore, it is a MUST to disable enforcement for switch-to-switch communications (underlay) by configuring "`no cts role-based enforcement`" on all switch-to-switch links.

Figure 6: Fabric Edge Configuration

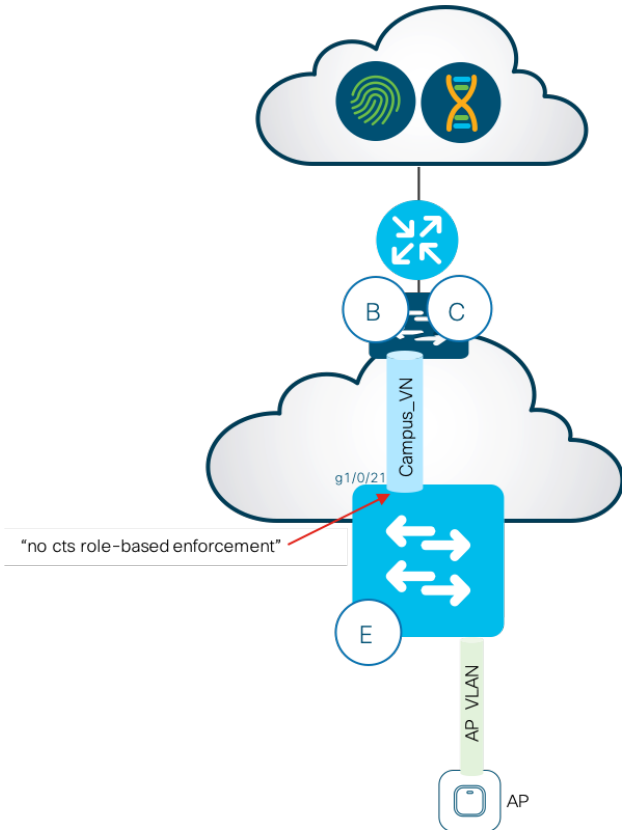

Note: See Appendix for logs generated when this configuration is skipped

Wireless Access Point Considerations

Currently Cisco DNA Center enables enforcement on the AP VLAN (2045). As a result, SGT assignment and pre-configured policies are required to allow the AP to connect when the default policy denies communication. You will need to make configurations in two places. First on the fabric edge that the AP is connected to. Secondly, configuration of a group-based policy in Cisco DNAC to allow communications between the unknown SGT(0) and the SGT assigned to the APs as shown below.

Note: Enforcement on VLAN 2045 can be disabled. However, a re-provisioning attempt from DNA Center would simply re-enable enforcement

Note: SGT=0 (Unknown SGT) is used here since the WLC is unclassified

Configuration steps:

1. Configure a VLAN (2045) to SGT mapping via Cisco DNA Center by navigating to the **Host Onboarding** tab within the site provisioning flow

Figure 7: Cisco DNA Center Access Point SGT Configuration

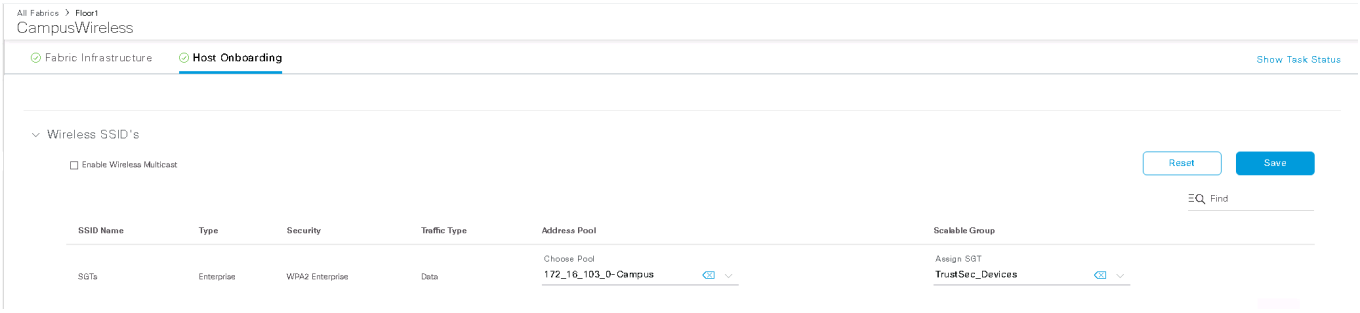

Or via SSH:

```
Edge(config)#cts role-based sgt-map vlan-list 2045 sgt 2
```

Note: This command assigns the SGT=2:TrustSec_Devices as the SGT for the AP. This SGT could be anything but for the reasons cited in the "Recommendations" section, I've used SGT=2.

2. Configure "permit ip" policies to allow Unknown-SGT to SGT=2 and vice-versa to communicate by navigating to **Policy→Group-Based Access Control→Policies**

Figure 8: Group-Based Policy for Access Point

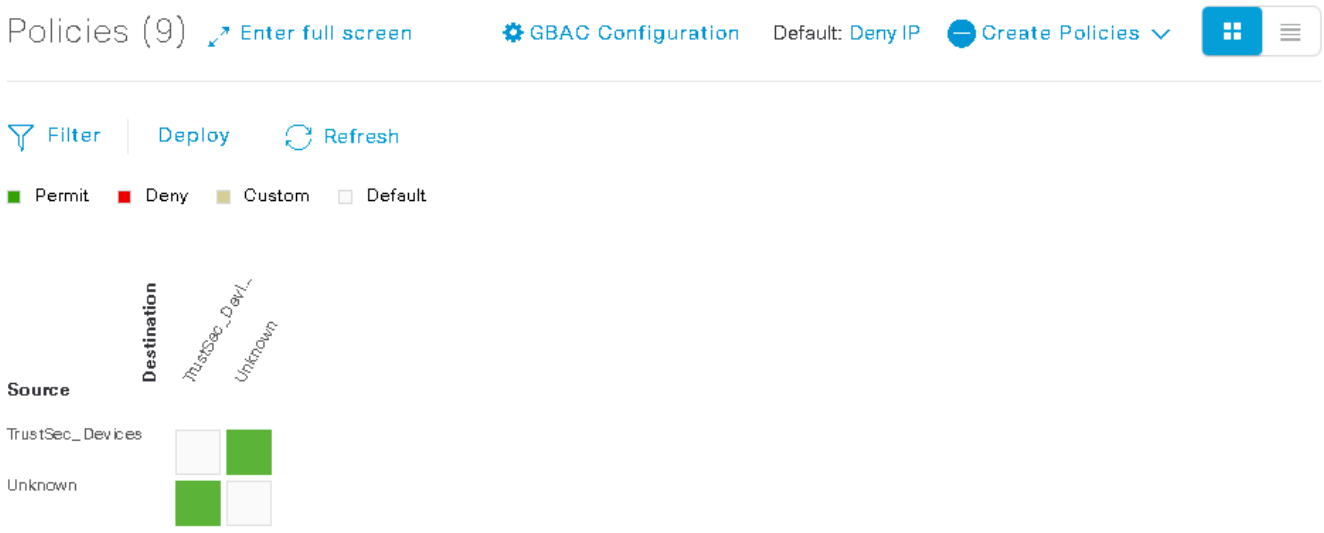


```
Note: A contract that's more specific can be used in place of "permit ip". Please reference the
"Securing Communications Further" section below for more details
```

Policy Extended Node Considerations

Currently Cisco DNA Center enables enforcement on VLAN chosen for policy extended node (PEN) management. However, there is not a reserved vlan id for a PEN and Cisco DNA Center does not provision a SGT to a PEN which make it difficult to create a policy to allow the PEN to initalize. Therefore, the current workaround is to disable enforcement on the uplink that connects the PEN to the edge.

In the diagram below, the PEN is connected to interface g1/0/1. Configure “no cts role-based enforcement” on the interface PRIOR TO going through the PEN configuration on Cisco DNA Center.

Figure 9: Policy Extended Node Configuration

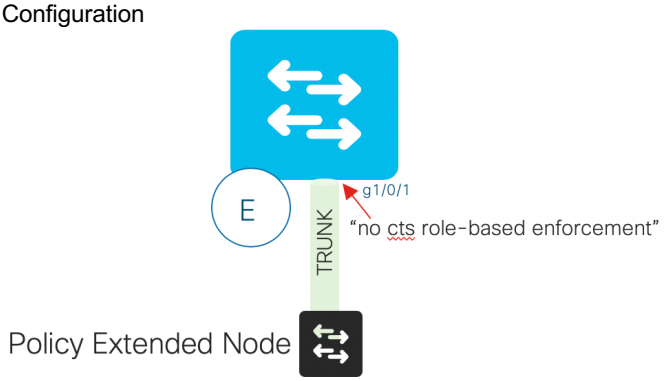

Allowing Communication Outside of the Fabric Considerations

By default, the SGT values of things outside a fabric are unknown (SGT=0). This includes site-to-site and north-to-south communications. Therefore, you must configure contracts to allow outside traffic (SGT=0) to any SGTs assigned within the fabric.

For example, for an employee (SGT=4) to communicate with a DNS server that is hosted outside of the fabric, a contract is necessary to permit SGT 0 to SGT 4.

Note: Using SGT=0 in a contract allows any type of communication from any unclassified device.

General Guidelines for Securing Communications Further

ISE

Like other network endpoints, SDA fabric devices must be authenticated and authorized by ISE to download SGTs and group-based policies. As part of this process, the fabric devices be assigned to an SGT. By default, ISE configures this SGT to SGT=0 (Unknown). This SGT is known as the “device SGT”.

Currently, the device SGT has no relevance in the fabric.

However, when considering a whitelist policy model, it is recommended to assign SGT with the value =2: TrustSec_Devices or some SGT value other than 0 to avoid the need to have a catch all policy allowing 0:Unknown to 0:Unknown communications.

Note: In a future ISE release, the device SGT will be 2:TrustSec_Devices by default

Navigate to **Work Center-->TrustSec-->TrustSec Policy--> Network Device Authorization** and set the default rule to assign the security group to "TrustSec_Devices"

Figure 10: Cisco ISE Device SGT Configuration

Appendix

Fabric Edge Details

Commands that enable enforcement:

```
Edge#
cts role-based enforcement←Enables SGT-based enforcement globally
cts role-based enforcement vlan-list 1027,1030← Enabled SGT-based enforcement on specific VLANs
```

When a default deny is deployed:

```
Edge#
*Feb  6 11:27:05.572: %BFDFSM-6-BFD_SESS_DOWN: BFD-SYSLOG: BFD session ld:1 handle:1,is going Down
Reason: RX DOWN
*Feb  6 11:27:05.576: %CLNS-5-ADJCHANGE: ISIS: Adjacency to Border (GigabitEthernet1/0/21) Down, bfd
neighbor down
```

```
Edge(config)# int g1/0/21 ←Uplink to next hop switch
Edge(config-if)#no cts role-based enforcement
Edge(config-if)#
*Feb  6 11:30:50.997: %BFDFSM-6-BFD_SESS_UP: BFD-SYSLOG: BFD session ld:1 handle:1 is going UP
*Feb  6 11:30:51.006: %CLNS-5-ADJCHANGE: ISIS: Adjacency to Border (GigabitEthernet1/0/21) Up, new
adjacency
```