

Defending Against Malware Propagation with TrustSec Software Defined Segmentation

Cisco TrustSec® simplifies the provisioning and management of secure access to network services and applications. Compared to access control mechanisms that are based on network topology, Cisco TrustSec defines policies using logical policy groupings, so secure access is consistently maintained even as resources are moved in mobile and virtualized networks. De-coupling access entitlements from IP addresses and VLANs simplifies security policy maintenance tasks, lowers operational costs, and allows common access policies to be applied to wired, wireless, and VPN access consistently.

Business Issue

Every business or entity today needs to be concerned about the spread of Malware within their organization. Many attacks come through sophisticated hacking using various probing techniques resulting in the eventual exploit of vulnerable systems or security policies as implemented in firewalls and other security devices. These attacks look to take advantage of unpatched systems or security holes within an enterprise's DMZ to gain access to the organizations data.

Some of the most devastating attacks recently however have occurred as a result of an employee's PC having been compromised resulting in the spread of malware amongst user machines in order to collect information used in later attacks. Some of these attacks are the result of phishing attacks or may be more sophisticated spear-phishing attacks where the hacker is posing as a co-worker requesting that the employee look at a document that may be pertinent to a project they are working on. The following provides an example of just such a scenario.

Your employee is working at home on their corporate issued laptop. In checking their email, outside of the safety of the corporate email and web security solution, they open an email from a co-worker who references a particular URL to a document on the web in relation to a project they are working on. The employee clicks the URL to read this document. Unfortunately, unbeknownst to them, the URL to which they were redirected was not a reputable site but a malicious website from which sophisticated malware has now been installed. This wasn't just any spam-based phishing attack but one that was carefully researched and executed by an organization attempting to breach your corporate network and obtain highly confidential, intellectual property. You have been targeted!

Sound like something out of a spy novel? Hardly, some of the most destructive intrusions have occurred in similar fashion.

When the employee arrives at the office the next day, the malware launches and begins propagating throughout the network as seen in Figure 1 below.

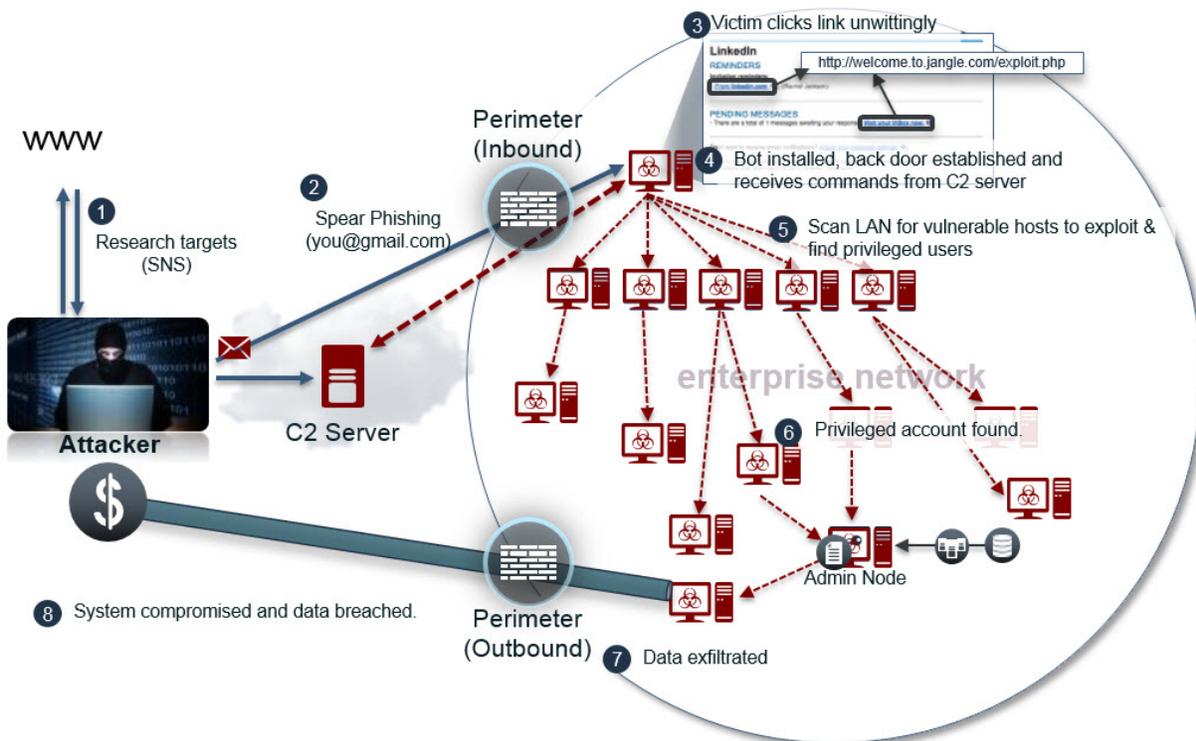


Figure 1 Typical Phishing attack

With TrustSec Software Defined Segmentation implemented in your network however, when your employee arrives at the office and logs into the TrustSec enabled network, something occurs that the group launching the attack hadn't anticipated. Upon successful authentication at the Cisco Identity Services Engine, your employee is identified and associated with a specific role, or Security Group, with a policy associated with it. This policy can block communications between users.

Benefits of TrustSec Software Defined Segmentation in preventing Malware propagation

- Restricts communications between your users or various devices within a VLAN.
- Replaces your policies built on MAC and IP Addresses by specifying "Roles".
- Eliminates the need to create complex, IP-based access lists to control traffic.
- Common policy regardless of where, when, or how a user or device accesses the network.
- Simplifies operational complexity and lowers operational expenses through the use of centrally managed role-based policies vs distributed access lists constructed to conform to local IP addresses.
- Reduces time to create and deploy new policies to address new and ever changing threats

Use Case Details

TrustSec software defined segmentation is an evolutionary security technology integrated in Cisco products that changes the paradigm by which networks have been secured both in the past and today through the use of access lists. With Cisco TrustSec technology, it is no longer necessary to configure access lists by geographically sensitive and oft times obscure IP Addresses of users and devices. With TrustSec, implementing an enterprise security policy changes to a “role-based” model that can be applied ubiquitously across your enterprise whether at you corporate offices or a remote campus or branch located on the other side of the globe.

When a user attaches to the network whether via wired, wireless, or VPN they will be authenticated and authorized by the Cisco Identity Services Engine. As part of the authorization process configured within ISE, the user is associated with a role or Security Group. This Security Group information, is then communicated via RADIUS to the network device to which they are attached in the form of a TrustSec Security Group Tag, or SGT. This SGT is a numerical value representing the Security Group and is associated with their IP Address. All IP traffic sourced from the user will thus be associated with this SGT. Should the user or device move to a different location in the network, although their IP Address may change, their Security Group membership will remain the same based on their authentication to the network.

The SGT associated with the user or device’s IP address can be propagated throughout the network in one of two ways. The first method is by adding an eight byte header to the Ethernet frame which contains the SGT value. This method is known as inline tagging which requires hardware support on the network device’s Ethernet ports connecting to other network devices. The second method which can be used in lieu of inline tagging is through the use of a TCP-based protocol known as the Scalable-Group Tag eXchange Protocol or SXP which can advertise the IP Address to SGT mapping of the user to peer devices.

The role and policies associated with the employee are instantiated in the network through the TrustSec policy centrally configured at the Cisco Identity Services Engine and communicated to the network device automatically upon successful user or device authentication. These TrustSec Policies are a role-based policy permitting and denying communications between security groups through the use of a simple rule such as “permit ip” or “deny ip” or may include more granular rules or access control entries permitting and denying specific protocols or ports. When deployed to network devices such as switches and routers, these role-based policies are known as Security Group Access Control Lists, or SGACLs. These SGACLs differ from the standard access control lists in that they are not built using IP Addresses but are based on the TrustSec Security Group Tag.

The TrustSec Policy Matrix within ISE, is used to define a TrustSec policy and any associated SGACLs. This matrix can be seen in Figure 2 below. In the matrix below simple policies such as “Permit IP” and “Deny IP” can be seen as well as policies in blue with an SGACL attached such as “Anti_Malware”

Source	Auditors (9/0009)	Developers (8/0008)	Development_Ser... (12/000C)	Employees (4/0004)	Finance (20/0014)	Production_Serv... (11/000B)
Auditors (9/0009)	Anti_Malware	Deny IP	Deny IP	Deny IP	Deny IP	Permit_HTTP_HTTPS
Developers (8/0008)	Deny IP	Anti_Malware	Permit IP	Deny IP	Deny IP	Permit IP
Development_Ser... (12/000C)	Deny IP	Permit IP	Permit IP	Deny IP	Deny IP	Permit_HTTP_HTTPS
Employees (4/0004)	Deny IP	Deny IP	Deny IP	Anti_Malware	Deny IP	Permit_HTTP_HTTPS
Finance (20/0014)	Deny IP	Deny IP	Deny IP	Deny IP	Anti_Malware	Permit_HTTP_HTTPS
Production_Serv... (11/000B)			Deny IP			Permit IP

Figure 2 TrustSec Policy Matrix

As new IP to SGT mappings are learned at a network device, the policies associated for that specific security group will be retrieved dynamically from ISE. With the policy distributed throughout the network, devices such as switches and routers can both use a common policy for traffic forwarding decisions between different Security Groups or even between members of the same Security Group. It is through this ability to restrict communications between users within the same Security Group and even within the same VLAN, that lateral movement of malware between users can be prevented.

Figure 3 below depicts the TrustSec framework. In this figure you can see that three different roles or Security Groups have been defined, those being Auditor, Finance, and Developer. Additionally, servers have also been classified and associated with specific Security Groups and policies can then be defined restricting access to those assets. Along with the policies controlling access to data center resources though, TrustSec policies can be created restricting access between users in different Security Groups as well as within the same group as seen in the matrix in Figure 3.

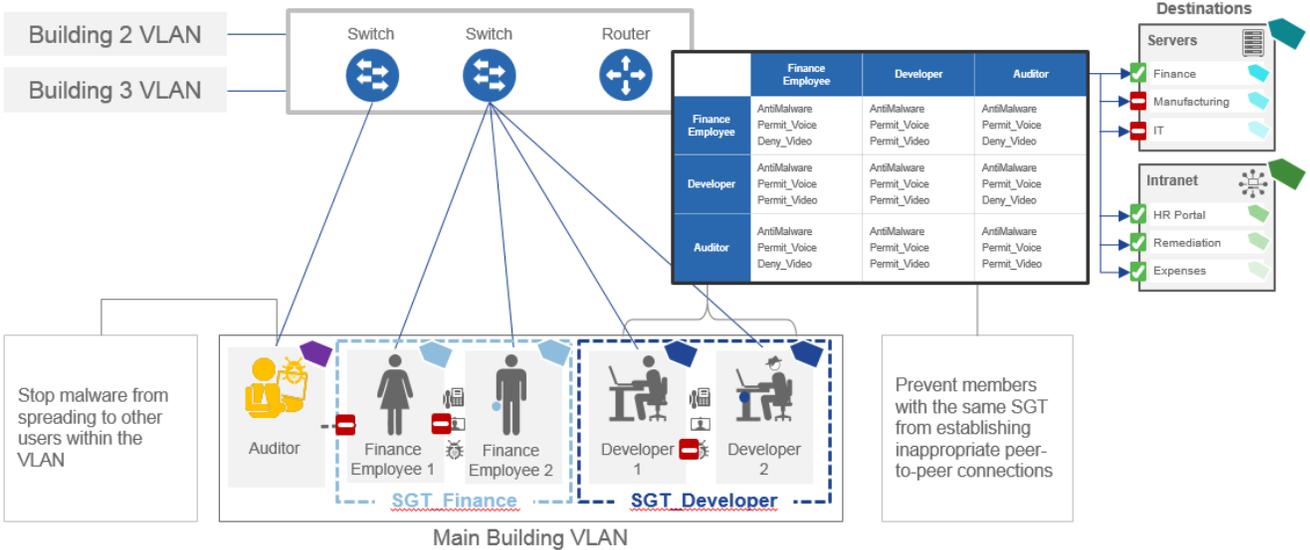


Figure 3 TrustSec Microsegmentation

In the matrix, the SGACL's for all users contain access control entries restricting communications over specific protocols between users most commonly associated with malware exploits and shown as "AntiMalware" in the example. Voice is permitted between all user Security Groups but Video is only permitted between Developers and between Developers and Finance. An example of an actual Anti-Malware policy that can be deployed can be seen in the sample below:

```
deny icmp
deny udp src dst eq domain
deny tcp src dst eq 3389
deny tcp src dst eq 1433
deny tcp src dst eq 1521
deny tcp src dst eq 445
deny tcp src dst eq 137
deny tcp src dst eq 138
deny tcp src dst eq 139
deny udp src dst eq snmp
deny tcp src dst eq telnet
deny tcp src dst eq www
deny tcp src dst eq 443
deny tcp src dst eq 22
deny tcp src dst eq pop3
deny tcp src dst eq 123
deny tcp match-all -ack +fin -psh -rst -syn -urg
deny tcp match-all +fin +psh +urg
permit tcp match-any +ack +syn
permit ip
```

In the anti-malware policy above, services that would typically never be found running on a workstation such as DNS, NTP, SQL Server, and Web are defined and denied, as well as communications between workstations such as Telnet, SSH, RDP and FTP. Additionally, traffic displaying various port-scanning techniques such as “Christmas Tree” scans where all or a number of TCP flags are set to inspect the host’s response in an effort to determine the OS running on a device can be thwarted by the two “deny tcp match-all” access control entries in the sample above.

Creating and implementing these SGACLs and then incorporating them in a TrustSec policy is an easy two-step process. The first step is to define the actual SGACL that will be used. Figure 4 below shows the creation of the SGACL within ISE.

1. At Cisco ISE navigate to “Work Centers” > “TrustSec” > “Components” > “Security Group ACLs”.
2. Provide a name, description, and IP version for the SGACL.
3. Define the access control entries that will comprise the SGACL.

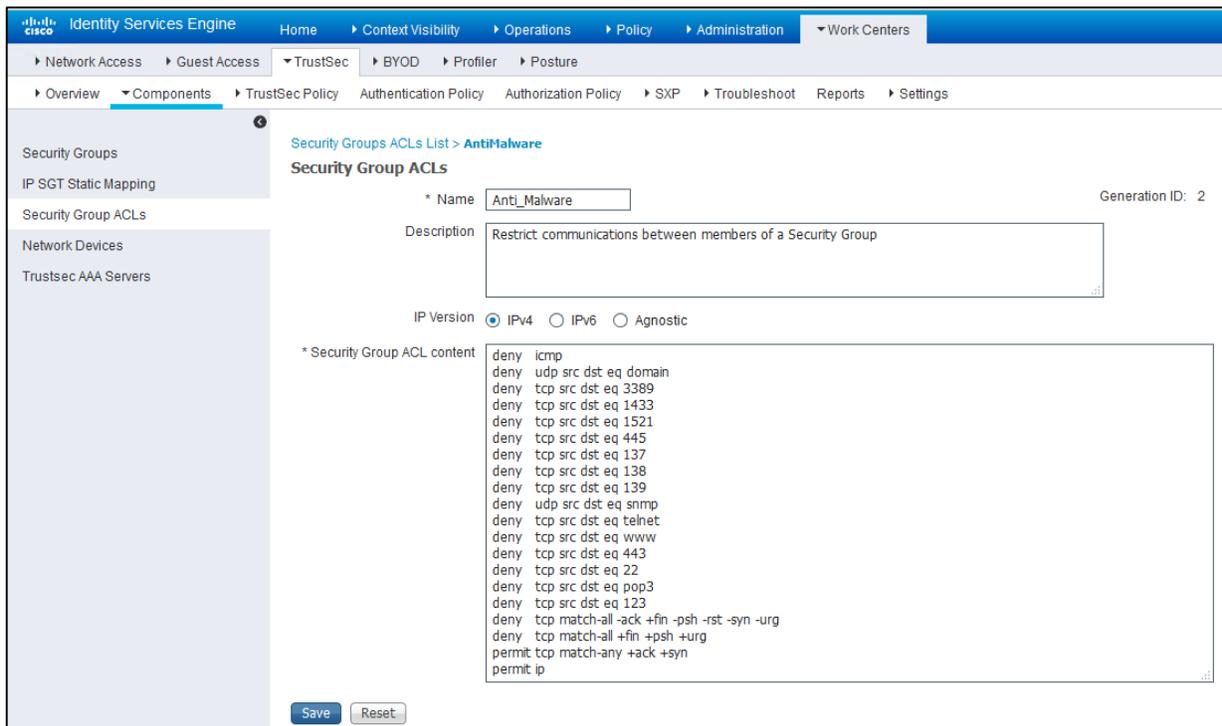


Figure 4 SGACL Creation

Note: The SGACL depicted above can be used on any Catalyst switch. Nexus switches based on NX-OS however may not be able to use the operands above and should be tested prior to deployment. The TrustSec Policy containing this SGACL will only be downloaded to those devices that contain IP-SGT mappings for the Security Group to which this SGACL is deployed to such as “Employees” for example. In this example then, only Catalyst switches would download these policies as they would be the only devices with mappings for Employees.

Once the SGACL has been created it will then be used in defining a TrustSec policy. Creation of the policy is performed using the TrustSec Policy Matrix as seen earlier in Figure 2.

1. Navigate to "Work Centers" > "TrustSec" > "TrustSec Policy" and the matrix view as seen in Figure 2 will open.
2. Using the mouse, hover over the cell where "Employees" intersect as source and destination and a pencil will appear as seen in Figure 2.
3. Click on the pencil and the dialog box as seen in Figure 5 will open.
4. The source and destination Security Groups should be listed. Click on the "Assigned Security Group ACLs" drop-down box and select the SGACL created earlier.
5. Click "Save" and once back at the Policy Matrix select "Deploy" from the menu bar.

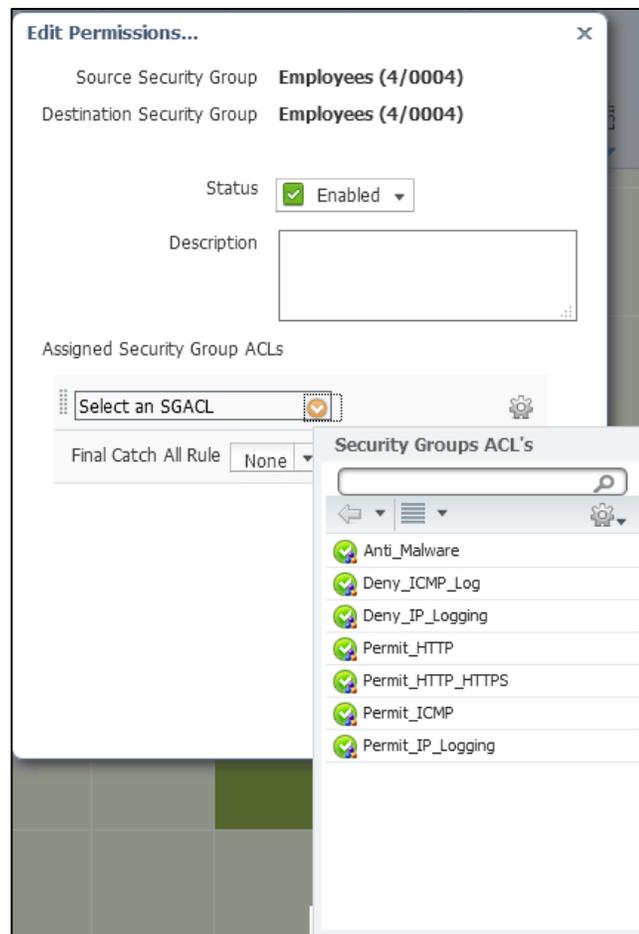


Figure 5 TrustSec policy creation

Once deployed, the new TrustSec policy will block the lateral propagation of Malware as seen in Figure 6 below. Having implemented TrustSec Software Defined Segmentation in the network, centralized definition of granular policies can now be easily implemented and or changed as new threats occur without the need to re-write numerous IP-based ACLs. These changes can then be pushed to the network through the press of a single button drastically reducing operational costs and considerations while securing the network in a most expedient fashion.

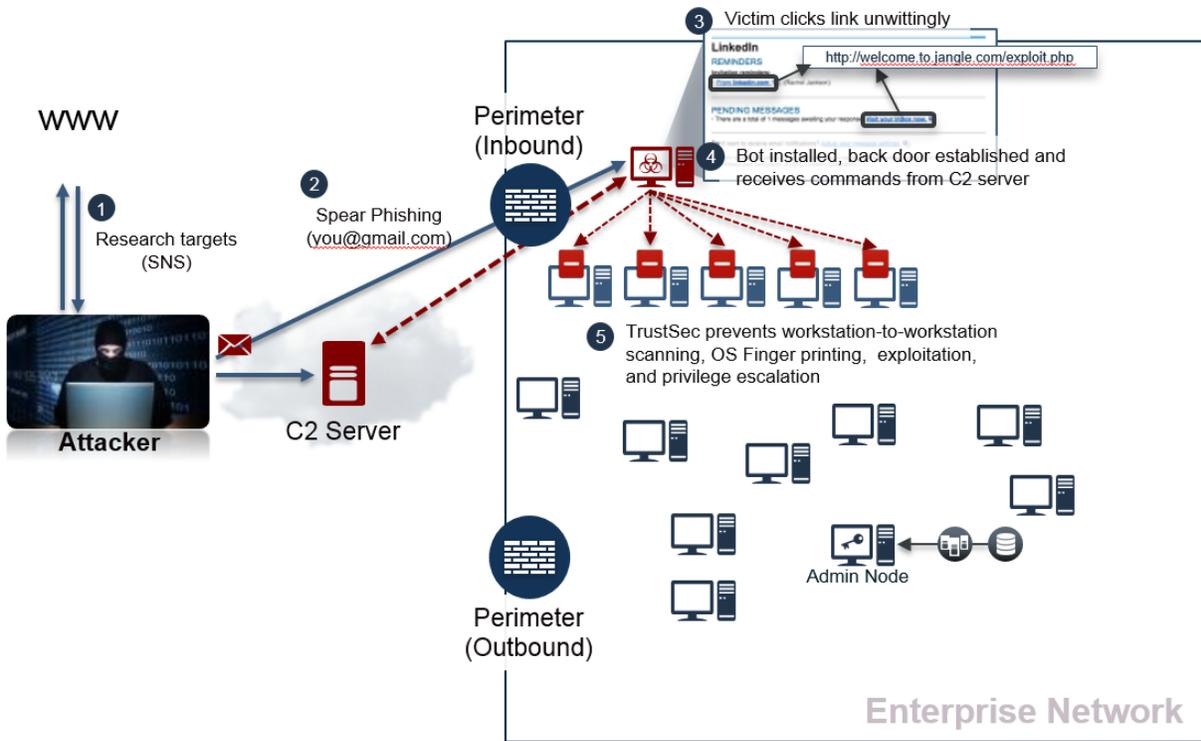


Figure 6 Malware lateral propagation denied

For More Information

Please visit our page on Cisco.com at <http://www.cisco.com/go/trustsec> or the TrustSec Security Community at <https://communities.cisco.com/community/technology/security/pa/trustsec>.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Recycling symbol: 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: 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)