IMM Configuration Sharing and Policy Cloning Across Organizations

Understanding new IMM Capabilities that support the creation of efficient and scalable policy models

October 2023

Contents

Introduction

This guide discusses several much-anticipated features recently added to Cisco Intersight® Managed Mode (IMM) that allows logical resources (pools, policies, and templates) in one organization (org) to be shared with resources in another organization, coupled with ways to clone these logical resources across organizations.

While many customers transition to IMM from spending years with Cisco UCS® Manager (UCSM) and Cisco UCS Central (UCS Central) and having a deep understanding of these management solutions, some customers will adopt IMM with their initial introduction and implementation of Cisco UCS – perhaps with a new Cisco UCS X-Series deployment, for example – and do not come from the traditional UCSM/UCS Central knowledge background. For new Cisco UCS customers, please feel free to skip the initial section below: "Context – Where we were and where we are now…" and resume with the following section "Explanation – configuration sharing across organizations," as you see fit.

Context – Where we were and where we are now

Before we embark on explaining some exciting new functionalities in IMM, let's briefly review some important capabilities and differences as compared to UCSM and UCS Central, particularly for existing customers who struggle with the lack of organization policy resolution in IMM.

With UCSM and UCS Central, there is an organizational (org) hierarchy that has existed with these Cisco UCS management solutions from their inception. While an entire paper could be dedicated to this subject alone, I want to focus and discuss two of the main capabilities in UCSM/UCS Central, namely policy resolution and default policies that have proved at times to be a barrier to IMM adoption for some customers.

Policy resolution is an automatic, built-in capability within UCSM and UCS Central, in that each org level can have nested sub-orgs that are logically connected to one another for name-based policy resolution. If the service profile (SP) is located inside a nested sub-org structure, then the resolution path for each policy lookup will extend from the sub-org where the SP lives, all the way up the hierarchy to the /root level, which is at the top of the UCSM/Central organizational hierarchy. The /root org is system-defined and cannot be altered or deleted.

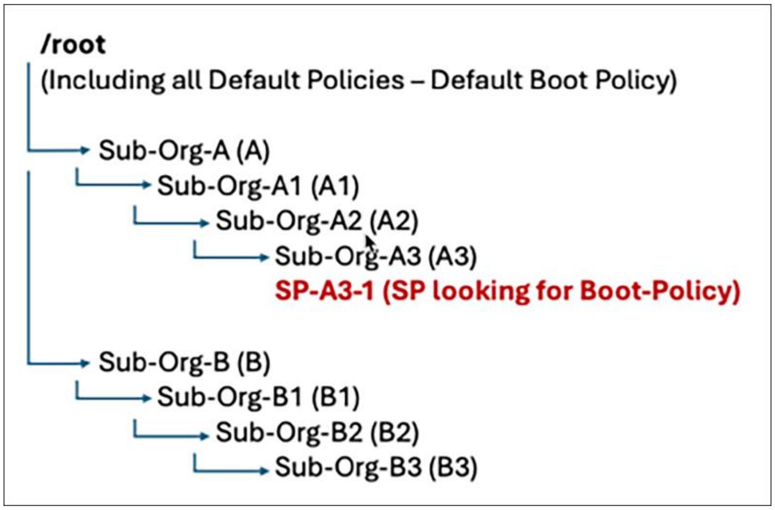

Figure 1.
UCSM/UCS Central policy resolution example

As an example, a simple organizational structure is displayed in Figure 1. An org structure in UCSM/UCS Central starting from root (as described by: /root -> nested A -> nested A1 -> nested A2 -> nested A3), with the service profile (SP) instantiated in nested Sub-Org A3. SP-A3-1 will use name-based resolution to search for a custom policy named “Boot-Policy,” and that policy can be instantiated in Sub-Org A3, or A2, or A1, or A or root – or all five levels! If there is a policy named “Boot-Policy” in level A3, then the SP will resolve to that policy first and foremost (because it is closest to the SP). If there’s another “Boot-Policy” in level A2, and if the policy instance in level A3 is deleted for some reason, then the resolution will automatically switch to “Boot-Policy” in Sub-Org A2, and this behavior will continue upward all the way to /root. If all the instances of “Boot-Policy” are the same, then nothing disruptive would ever happen to the server if a given lower-level policy instance is deleted, and the automatic resolution will select the next higher instance; however, if there is the slightest difference, then a server reboot will be triggered, only intercepted by a User-Ack Maintenance Policy that is properly configured. With UCSM/UCS Central architecture, the endpoint state (server) must match the desired state (the configured SP, and the code always tries to enforce that.

Now, what if all the instances of “Boot-Policy” (one or more) are deleted? The SP still needs a boot policy, and that is where the default policies in the /root org come into play. Seemingly, the purpose of these policies is to give the end user an example or guide to how to configure the different policies, and while they can certainly help in that effort, many policies are more complex with many different configuration parameters, so it is not possible to provide meaningful canned-examples for all use-cases. This brings us to the primary reason default policies exist in UCSM and UCS Central. With an SP in UCSM/UCS Central, there are required policies that must be attached to the SP for the SP to deploy successfully to the server. For example, you must have a BIOS policy, you must have a Boot policy, etc. Given that, if the SP is looking for a custom-created policy named “Boot-Policy,” and if that policy does not exist anywhere in the resolution hierarchy (for example, someone mistakenly deleted the Boot-Policy), or if there was no custom-created boot policy created to begin with, then that default-boot policy at the /root level would automatically slide into place in the SP upon SP deployment to the server. We absolutely need a Boot policy, so in the absence of a custom named Boot policy, the default policy will be used.

But will this be disruptive? Well, in the case where the custom “Boot-Policy” is deleted, and the settings in the default-boot policy differ from the settings previously configured in “Boot-Policy,” then, yes, it could be disruptive (if User-Ack is not configured properly in the maintenance policy); it is also doubtful that many administrators checked the default policy settings and understood this automatic policy substitution behavior. Again, we are talking about complexity and a level of detail above and beyond what a lot of UCS configurations have entailed.

In discussing policy resolution and default policy behavior, as good as these capabilities are in UCSM and UCS Central, they were complex and, with the slightest misconfiguration, could cause disruption to the server. Certainly, many customers did not take advantage of all these capabilities to their fullest extent, and thus opted for a simpler configuration setup. However, those who took advantage of this policy resolution architecture and understood their use, crafted powerful, scalable policy-model designs. So, this raises the question: what use cases did they address?

First, creating different orgs and sub-orgs is very useful in rule-based access control (RBAC) segmenting the UCS domain or multiple UCS domains in the case of UCS Central. With UCSM and UCS Central, minimizing “Blast Radius” through segmentation is important because of the possibility of a single configuration error being disruptive to a lot of running workloads on servers. If you had 100 Service Profiles all using the same VSAN for Storage, and that VSAN is mistakenly deleted, you could disrupt 100 Servers. As such, the org hierarchy could achieve segmentation, isolating the configuration of one group of servers from another, and allowing the use of RBAC for appropriate users/roles to access each org/sub-org in that hierarchy.

Secondly, with multi-level org/sub-org model, you could place certain common pools/policies at a higher level in the hierarchy secured with appropriate RBAC and have lower sub-orgs hosting the SPs that would access those higher-level pools/policies, again with appropriate RBAC assigned to those lower-level SPs. Referring to Figure 1 above, here is another simple example: SPs in the Sub-Org-A hierarchy are separate from SPs in the Sub-Org-B hierarchy, and yet, both could resolve to the /root level for pool/policy access. This capability greatly reduces policy sprawl, and a single configured pool/policy can service dozens, hundreds, even thousands of servers (SPs).

This raises the question, "What about IMM, how does it compare?" Up to this point, organizations have been completely siloed in IMM; thus, all the pools, policies, VLANs, VSANs, templates, and server profiles need to be instantiated in the same org, and this siloed-approach contributes heavily to policy sprawl. If I have ten organizations in IMM, they will all need to have their own self-contained pools, policies, and SP-templates/SPs. Additionally, if they all use the same firmware-policy configuration, I will have to create that firmware policy ten times, one for each organization. And, when I need to make a change and edit that firmware-policy, I will have to edit all ten firmware-policies. This equates to an overhead burden on the administrator, and to policy sprawl. Cisco Intersight's Configuration Sharing Across Organizations, discussed in the next section, directly addresses this concern.

And what about default policies? This is a feature frequently requested by Intersight customers and is something we hope to deliver in the future; albeit as stated earlier, there will not be auto-policy resolution/hierarchy behavior with Intersight. IMM may have a simpler architecture, but it is more powerful and more flexible and has a safer change-control mechanism than UCSM and UCS Central. We don't want any customer workloads to suffer any unintended disruptions.

Having communicated some foundation and context, I strongly believe the remainder of this guide will be beneficial to customers using IMM today and those considering migrating to IMM. New, powerful features are being introduced to IMM to close some of the biggest gaps with legacy UCSM/UCS Central. With IMM, we are not just trying to meet the expectations of existing UCSM/UCS Central customers but to far surpass them and strengthen IMM's features.

Explanation – Configuration Sharing Across Organizations

To address the policy sprawl dilemma with the original implementation of IMM, I am pleased to discuss with you an exciting capability new to IMM and mentioned at the end of the previous section: Configuration Sharing Across Organizations. Many will know and refer to this feature as "shared-orgs," which is at the crux of the feature's capabilities.

Configuration Sharing Across Organizations (CSAO) directly addresses the customer-desired segmentation of organizations with appropriately applied RBAC to facilitate the network designs of large-scale enterprises and managed service providers (MSPs) alike, while minimizing the policy sprawl alluded to above.

Unlike UCSM and UCS Central, there is no need for built-in policy-resolution mechanisms. In IMM, server profiles (SPs) do not have the same mandatory policy requirements to successfully deploy a SP to an endpoint and at no time are server endpoints automatically rebooted because of some policy edit/change for the accompanying SP. IMM server profiles have better change management control and are safer to manage and operate. Consequently, automatic policy resolution has never been a design intent of Cisco Intersight, and with this new CSAO functionality, it is simply not required. Conversely, this feature leverages manually assigned sharing rules to allow or disallow resource sharing across organizations from one "source" organization to other "consumer" organization(s). This offers a simple, well-defined model and offers better control by administrators. This functionality also allows the flexibility to facilitate transitioning from pre-existing siloed policy-model architectures to a more efficient, scalable, shared architecture without server/workload disruption. Just as your business and network topologies are dynamic and changing in real life, so too is Cisco Intersight in being agile and flexible in supporting these required changes. Transitioning will be covered later in this guide.

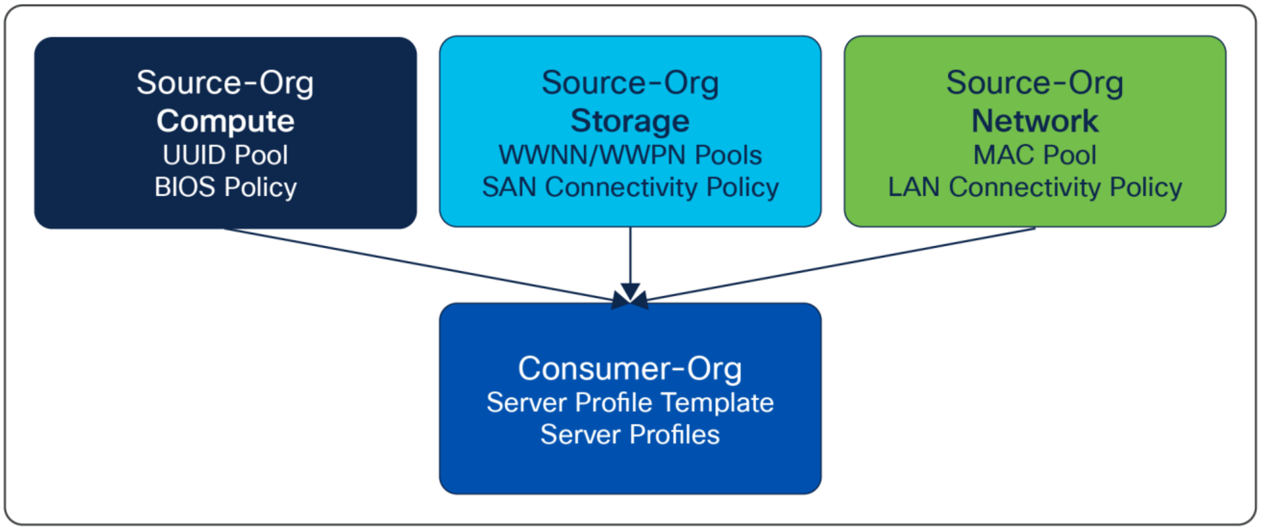

Figure 2.
Configuration Sharing Across Organizations

Example – Configuration Sharing Across Organizations

The best way to explain this feature is to reference an example in IMM. In Figure 2, we have a very simple topology of IMM organizations and their relationships. It is important to understand that, while it appears there is some type of “hierarchy” in the example, with IMM there is in fact no hierarchy between organizations. The graphic is depicted vertically, but this is just a way to show there is a sharing relationship between the organizations. All organizations are equal, and the rules established for sharing (or not), and the related RBACs, are constructs of the IMM org model. Nothing is implicitly shared between organizations; the rules to share between orgs must be established manually.

In our example above, we have 3 x organizations that align with typical disciplines in the UCS ecosystem, namely compute, storage, and network. Each of those organizations will have an appropriately defined RBAC giving network administrators access to the org to create and configure certain policies to be later “consumed” by the consumer organization. This is a use case of higher-level expert administrators creating and editing policies to be used and consumed by adjacent organizations with the SP templates/SPs that are defined and configured in the consumer organizations, with their appropriate users and roles.

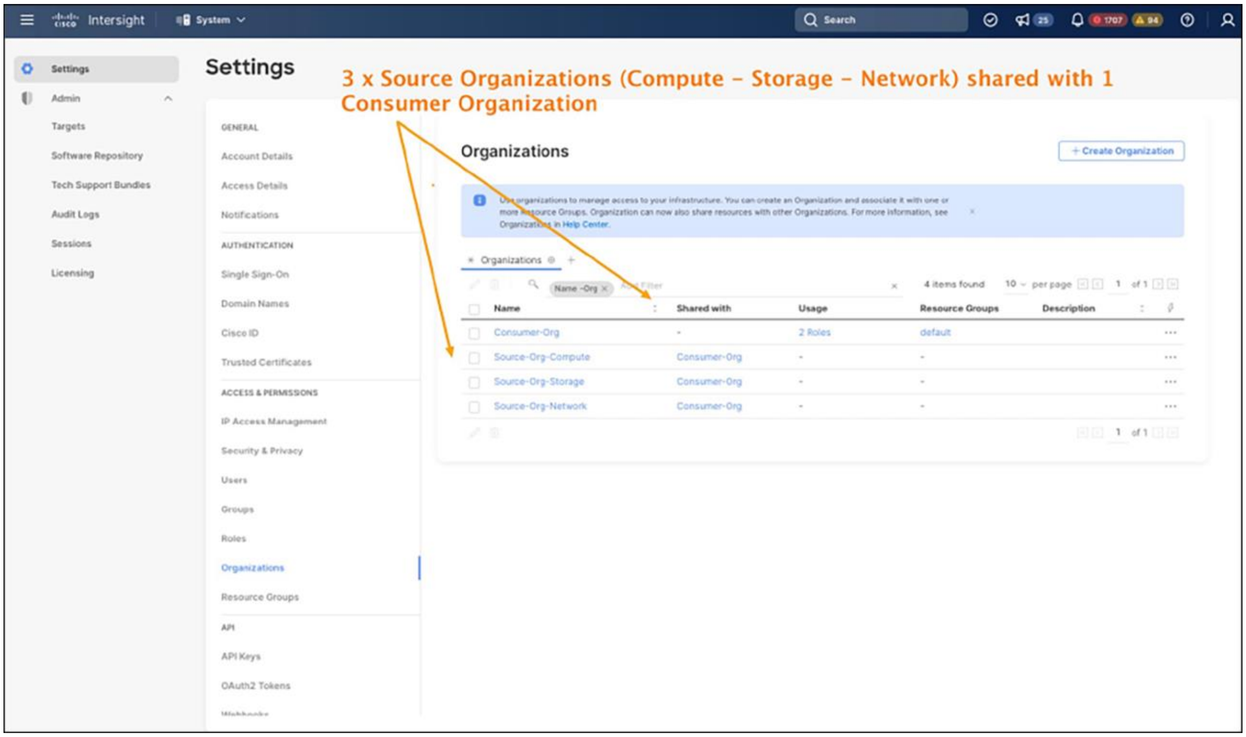

3 x Source Organizations (Compute – Storage – Network) shared with 1 Consumer Organization

Settings

Organizations

Use organizations to manage access to your infrastructure. You can create an Organization and associate it with one or more Resource Groups. Organization can now also share resources with other Organizations. For more information, see Organizations in Help Center.

Name	Shared with	Usage	Resource Groups	Description
Consumer-Org	-	2 Roles	default	
Source-Org-Compute	Consumer-Org	-	-	
Source-Org-Storage	Consumer-Org	-	-	
Source-Org-Network	Consumer-Org	-	-	

Figure 3.
Source Orgs (Compute, Storage, Network) and Consumer Org

It is also important to understand that this guide is not suggesting the simplified design described above is a “best practice,” because many different customers will craft their organizational models according to their own business needs, with some taking a more geographic approach to org creation while others focus on business units, platform, or workload use-cases, or, in the case of managed service providers, individual customers.

To keep this guide manageable and brief, we will focus on a compute organization and consumer organization in our scenario. For the compute organization, the administrator of that organization has created the UUID pool and BIOS policy and has taken the steps to share that organization with the intended source organization.

In the Cisco Intersight System settings, and with organization creation, there is a checkbox to enable you to share configuration resources with adjacent organizations. The next screen allows you to select one or more of the adjacent organizations to share with. This is depicted in figures 4 and 5 below. Currently, you need to have an account administrator role to create and edit these organizations.

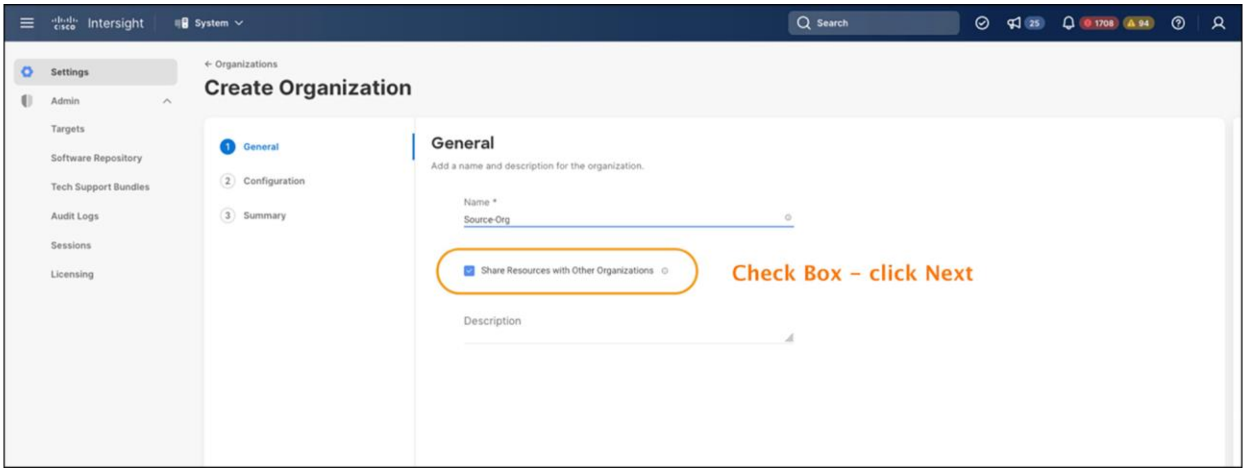

Create Organization

General

Add a name and description for the organization.

Name *
Source-Org

☑ Share Resources with Other Organizations

Check Box – click Next

Description

Figure 4.
Creating Organization – enable organization sharing

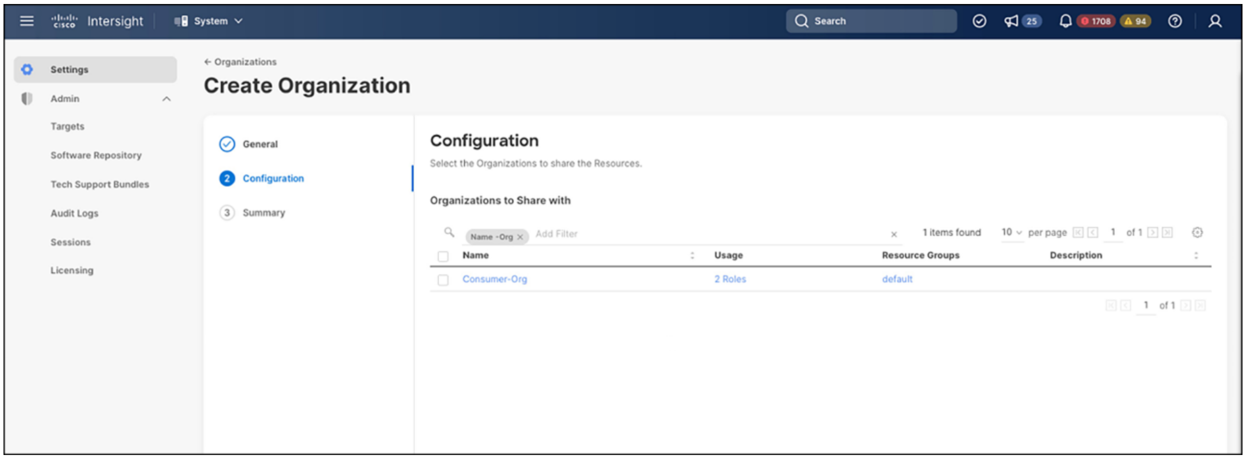

Create Organization

Configuration

Select the Organizations to share the Resources.

Organizations to Share with

	Name	Usage	Resource Groups	Description
□	Consumer-Org	2 Roles	default	

Figure 5.
Organization sharing, second screen – selecting organizations to share

Using this guidance, and if building from scratch, you may want to consider creating the consumer organizations first, and then the source organizations so that you can subsequently enable sharing and then select the intended consumer organizations while creating the source organizations. However, you can always go back and edit an existing organization and enable sharing, then select the consumer organizations. In Figure 3, you can see the 3 x source organizations and the designated consumer organization. While each of the three source organizations can have their appropriate administrators, or a single combined administrator for all three, the user assigned to the consumer organization (perhaps with a server admin role) will need at least read-only access to each of the source organizations for the consumer organization to access the source organization pools/policies, because you cannot consume what you cannot see.

Looking at Figure 6, we can see the Consumer-Org-Server-Admin Role assigned to the user of the Consumer-Org has server administrator privileges to the Consumer-Org. After all, we want this user to be able to create and craft the SP template/SP for that org. But, as we have stated previously, this user also needs read-only privileges to the Shared-Org-Compute. Referencing back to Figure 2, this user will also need read-only privileges to Shared-Org-Storage and Shared-Org-Network in order to access the pools and policies that are created in those other source orgs.

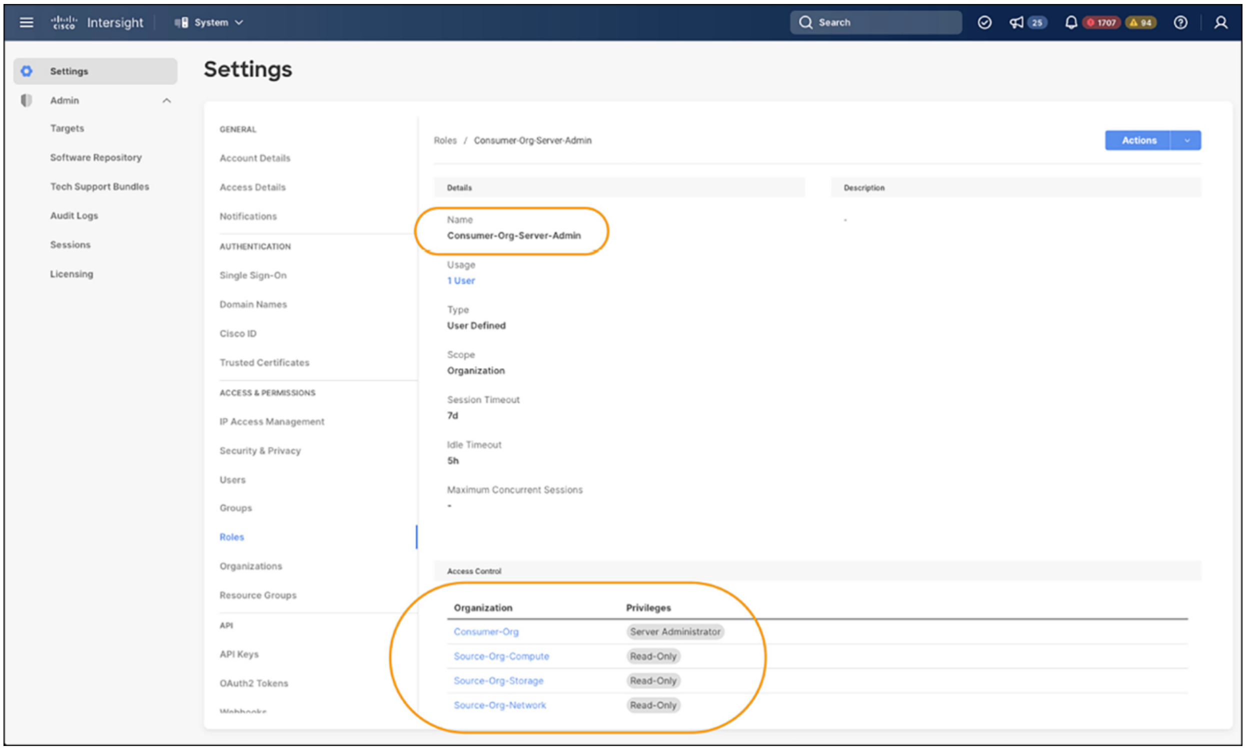

Name
Consumer-Org-Server-Admin

Usage
1 User

Type
User Defined

Scope
Organization

Session Timeout
7d

Idle Timeout
5h

Maximum Concurrent Sessions
-

Access Control

Organization	Privileges
Consumer-Org	Server Administrator
Source-Org-Compute	Read-Only
Source-Org-Storage	Read-Only
Source-Org-Network	Read-Only

Figure 6.
Organization and privilege access

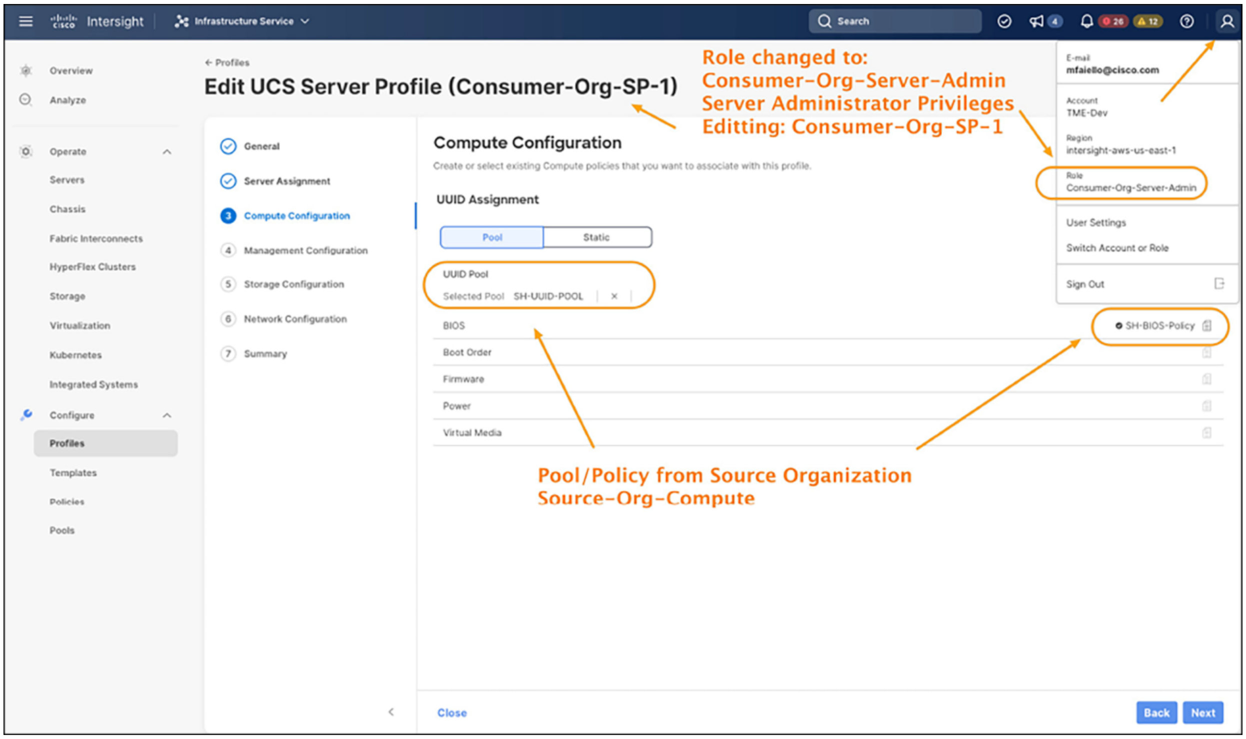

Edit UCS Server Profile (Consumer-Org-SP-1)

Compute Configuration

Create or select existing Compute policies that you want to associate with this profile.

UUID Assignment

UUID Pool

Selected Pool SH-UUID-POOL

SH-BIOS-Policy

E-mail
mfaiello@cisco.com

Account
TME-Dev

Region
intersight-aws-us-east-1

Role
Consumer-Org-Server-Admin

Role changed to:
Consumer-Org-Server-Admin
Server Administrator Privileges
Editting: Consumer-Org-SP-1

Pool/Policy from Source Organization
Source-Org-Compute

Figure 7.
Organization and privilege access

Per Figure 7, when the user for Consumer-Org logs into their account, with server administrator privileges, and creates an SP template, or, as in our case, the SP (Consumer-Org-SP-1), you can see that the user has access to the SH-UUID-Pool and SH-BIOS-Policy created by the server administrator in Source-Org-Compute.

To circle back to the aforementioned “Blast Radius” discussion in the first section, with UCSM/UCS Central, you might be asking, “What about IMM?” With Configuration Sharing Across Organizations, and linking resources between orgs, is there a legitimate concern that a single mistake – a wrong edit, a deleted policy – will take down my workloads in IMM? The simple answer is emphatically “No!” Welcome to the beauty and power of Cisco Intersight! Because of the safer change-controls enjoyed with the IMM architecture, you no longer need to worry about servers accidentally rebooting with a configuration change. All changes made to the SP (desired state) do not pass automatically to the servers (endpoint state). Simply put:

1. You edit the desired policy/SP and save.
2. You receive a change of status on the SP from green (“OK”) to orange (“Inconsistent”) Pending Changes.
3. You then have the means to verify (as a higher-level administrator) what exactly changed in the policy/SP (annotated on the first screen of the SP in yellow).
4. Once the change is verified, you deploy the changed SP with the option of rebooting the server immediately (Activate – if any changes require reboot) or simply stage the disruptive change at the endpoint.
5. Administratively reboot the server when convenient (for example, during a maintenance window, etc.) to resolve any inconsistency between the server endpoint and the SP change.

If you mistakenly delete a policy, or make the wrong configuration edit, then simply do not deploy the SP change, and edit the policy back to its original configuration. This is easy, and the server endpoint will never know the difference.

This architecture is superior to how things were done in UCSM/UCS Central. Customers will have to experience this to truly appreciate it.

In the next section, we can review the existing rules and limits that govern sharing configuration resources across organizations.

Rules - Configuration Sharing Across Organizations

- Organization creation and editing requires a user with an account administrator role.
- Cisco Intersight accounts are currently limited to 50 organizations.
- A maximum of 200 sharing rules per account is currently allowed.
- An organization can be shared with up to a maximum of 50 organizations in an account, and in compliance with the maximum sharing-rule limit of 200, given above.
- A hypothetical example: four source orgs can each share with 46 consumer orgs (4+46=50 total orgs) or (4 x 46 = 184 sharing rules)
- A consumer organization user is required to have at least read access to the shared source organization.
- No implicit, automatic read access will be given to the source organization from the adjacent consumer organization users upon creation.
- Configuration Sharing Across Organizations can be disabled only after removing all the relationships formed using the resources. In the consumer organizations, remove any SP/Template references to shared pools, policies, profiles, and templates.
- An organization associated with a resource group cannot be shared with other organizations.
- The port policy and its sub-policies do not support Configuration Sharing Across Organizations in phase-1 support of this feature. A port policy and all its sub-policies must exist within the same organization as a single entity. A port policy can include embedded policies: for example, Ethernet network group policy, flow control policy, link control policy, Ethernet network control policy, and link aggregation policy.
- Transitive sharing is not supported currently. It is 1:1 sharing, no multilevel sharing. Org-A can share with Org-B, but, in turn, Org-B cannot share with Org-C.
- Partial user access: a user has read-write privileges to a consumer organization but does not have read access to the source organization.
    - The user may have access to resources in the consumer organization but not to the shared resources.
    - The sharing rules are created without creating the user roles correctly.
    - Example: when a user has access to the server profile or template in the consumer organization but does not have visibility or access to a consumed policy shared with that consumer organization.

Migrating considerations - Configuration Sharing Across Organizations

For existing IMM customers, they have their Org structures siloed due to how Intersight operated at the time of org creation. Now, with Configuration Sharing Across Organizations, they can take steps to redesign their org structures to leverage the sharing capability and reduce some of the policy sprawl they are exhibiting in their Cisco UCS environments. Policy sprawl is simply creating the same policy repeatedly, which, in turn, requires the administrator to edit all those policies when it is time to make a change. Ideally, when the same policy is consumed by many workloads, it is better to create that policy once (or as few times as possible), and then make it available to consumer workloads. It is a matter of efficiency and reduced administrator burden and OpEx.

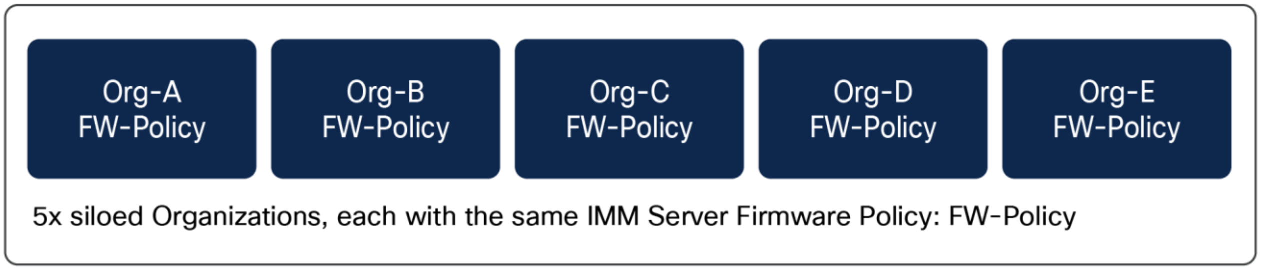
Org-A
FW-Policy
Org-B
FW-Policy
Org-C
FW-Policy
Org-D
FW-Policy
Org-E
FW-Policy
5x siloed Organizations, each with the same IMM Server Firmware Policy: FW-Policy

Figure 8.
Siloed organizations in existing IMM org structure

Unlike UCSM and UCS Central, making changes to an org structure is simpler and more flexible with IMM, and can be done without server or workload disruption.

Considering a simple example in Figure 8, we have five siloed pre-existing IMM organizations (Org-A, Org-B, Org-C, Org-D, and Org-E), each with its own pools, policies, and SP-templates/SPs. A small part of that policy configuration across all the organizations uses the same IMM server firmware policy: FW-Policy. The administrator had to create that policy five times, once for each organization (in this simple example, and likely many more times in real-world examples). Additionally, the administrator would need to edit each of these policies every time there is a change to the server firmware policy that is being supported or deployed.

By creating a shared source organization, as depicted in Figure 9, and establishing sharing rules to each of the previously mentioned siloed organizations, they now become consumer organizations. The administrator can create a Shared-FW-Policy in the Source-Org, which becomes selectable to the SP-template/SPs in the respective consumer organizations. The administrator can simply change the SP-template/SP reference for each, and reference the new Shared-FW-Policy in the Source-Org. When complete, the orphaned FW-Policies in the separate Consumer-Orgs can then be deleted. The next time the administrator needs to make a FW policy change, that administrator only needs to edit the single Shared-FW-Policy in the Source-Org. The Consumer-Orgs will automatically receive the Shared-FW-Policy change.

1. Create New Source Org and share with Consumer Orgs
2. No HW Resources in Source Org - Allows Sharing
3. Create same FW Policy, Shared-FW-Policy in Source-Org

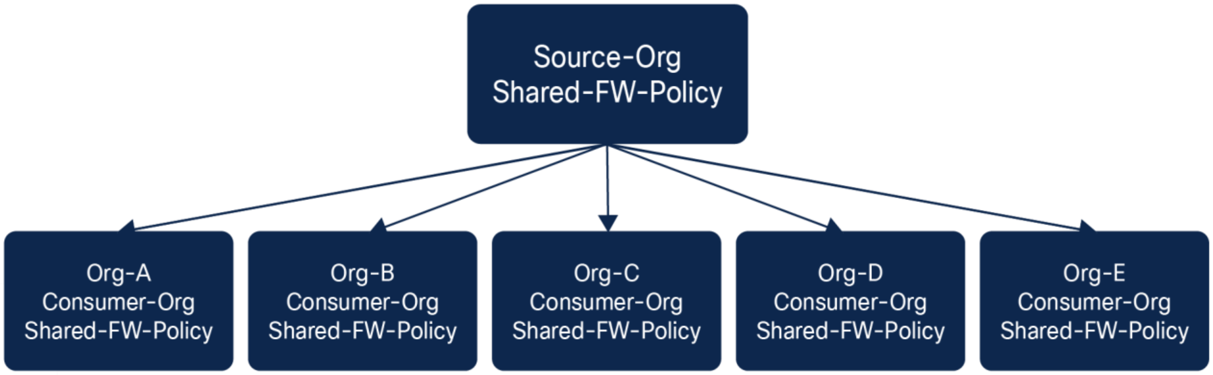

Source-Org
Shared-FW-Policy

Org-A
Consumer-Org
Shared-FW-Policy

Org-B
Consumer-Org
Shared-FW-Policy

Org-C
Consumer-Org
Shared-FW-Policy

Org-D
Consumer-Org
Shared-FW-Policy

Org-E
Consumer-Org
Shared-FW-Policy

4. Consumer Orgs maintain their existing Resource Groups - No Impact
5. Within each SP-Template/SP in each Consumer Org, edit and change reference for the FW-Policy to the Shared-FW-Policy in the Source-Org
6. After all SP-Templates/SPs have switched their references to the Shared-FW-Policy in the Source-Org, remove the individual FW-Policies in the Consumer-Orgs (Cleanup)

Figure 9.
Migrated org structure using source and consumer orgs

Now, the reality of real-world environments would be to maximize this capability to take advantage of efficiency and scale. While starting off with a few policies to get a better feel for how sharing works here, transitioning the bulk of your pools, policies, and SP-templates to a shared organizational model would be helpful in reducing both policy sprawl and administrative OpEx overhead. After all, it is easier to make a change in one place versus making those changes in five places, or fifty! One can create the pool/policy/SP-template constructs in the shared source orgs, and simply allow the separate consumer orgs to derive SPs from the templates in their respective consumer orgs. Or, if migrating, detach each SP from its existing template in the consumer org and reattach it to its like complement in the source org. You can think of this as a hub-spoke model.

Perhaps you will want to transition your ID pools to a shared-org model. In that case, build the shared source org and share with your consumer orgs. Take the individual UUID, MAC, and WWNN/WWPN pools and recreate them in the source org with the same ID parameters. Once this is done, you can take your SP-templates/SPs and change the resolution to the applicable pools in the source orgs. You will find that ID accountability will be maintained with those pools in the source orgs. The IDs in use will be properly audited. After you have “freed” your individual, original consumer org pools, they can be cleaned up and deleted.

As we covered in the rules section of sharing resources across organizations, it is important to realize that currently your port domain policy (and any of its embedded policies) will still need to exist in each separate organization. This, in turn, necessitates the domain profile existing in each separate organization. As previously stated, the ability to leverage shared orgs with a port domain policy is being considered for possible support in future enhancements to this capability, so more domain profile flexibility and scale may be coming to future enhancements with IMM.

Explanation – Policy Cloning Across Organizations

In IMM, we've had policy cloning since late May of 2023, but that initial implementation of policy cloning only allowed cloning of a policy within the same organization. Now, we can clone policies across organizations.

It is important to understand that Policy Cloning Across Organizations is a completely independent capability from that of Configuration Sharing Across Organizations (CSAO). In other words, you can clone any pool or policy (with the current exception to port policies and their embedded policies) across any two organizations. The organizations do not have to be "shared."

With Policy Cloning Across Organizations, administrators now have a quick way of replicating a policy and all its settings to another organization. This will certainly speed up buildouts and deployments. It will also allow more senior administrators to create and build out "golden" policies in one policy reference organization and, based upon need, to clone those golden policies to the appropriate working organization for SP-template/SP consumption. The policy cloning is secured with Cisco Intersight RBAC throughout.

Policy Cloning Across Organizations also offers the ability to "deep clone" any embedded policies (and pools) referenced by the parent cloned policy. For example, the LAN Connectivity Policy (LCP) has an embedded pool and, depending upon the use-case configuration, numerous policies and/or pools within it, as listed below:

- MAC Pool
- Ethernet Network Group Policy
- Ethernet Network Control Policy
- Ethernet QoS
- Ethernet Adapter Policy
- iSCSI Boot
- usNIC Adapter Policy
- VMMQ Adapter Policy

If you clone an LCP across organizations, then those embedded pools and policies will be cloned, as configured, to the destination organization as well. Pools will only clone the pool name itself, and you will always have to create the new ID blocks in the destination organization (to preclude any ID conflicts).

This same behavior would also apply to the SAN Connectivity Policy (SCP).

As you'll see in the next section, you will have a couple of options during this embedded policy cloning process, specifically addressing the situation of having a policy in the destination organization that has the **same name** as the embedded policy in the original organization. This might not be a common scenario, but with a large enough policy model and many organizations it is possible and something to be aware of. Also, if that named policy in the destination org is used unknowingly by the administrator, it is possible the destination policy, while having the same name, will have different configuration settings, so paying close attention to details is warranted in this case.

Example – Policy Cloning Across Organizations

Since the LCP is one of the more involved policies with an embedded pool and numerous policies, let's use that as an example for Policy Cloning Across Organizations.

Below, Figure 10 shows an example LCP and its embedded policies. To set some context, we are going to clone the LCP from the Boston organization to the NewYork organization. These two organizations are standalone; nothing is shared. Notice also that most of the polices use "Boston" in the name, since they originally resided in the Boston organization; however, I am purposely being generic with the Ethernet Network Group policy. That is because there are two of those policies with the same name "Eth-Network-Group," but the organizations are different, one policy is in Boston, and the other policy is in NewYork.

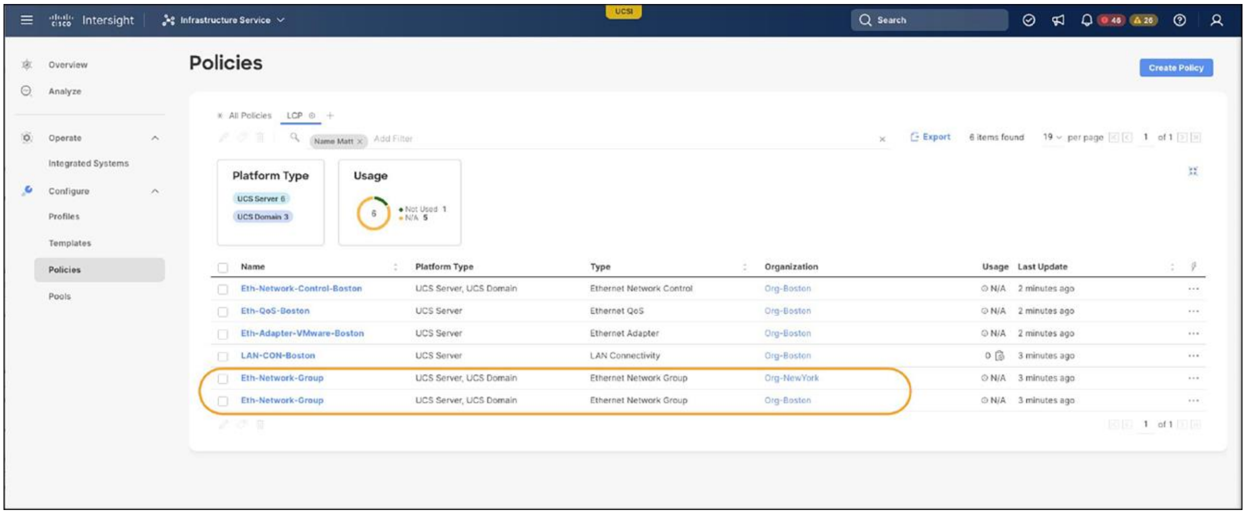

Policies

Platform Type: UCS Server 6, UCS Domain 3

Usage: 6 — Not Used 1, N/A 5

Name	Platform Type	Type	Organization	Usage	Last Update
Eth-Network-Control-Boston	UCS Server, UCS Domain	Ethernet Network Control	Org-Boston	N/A	2 minutes ago
Eth-QoS-Boston	UCS Server	Ethernet QoS	Org-Boston	N/A	2 minutes ago
Eth-Adapter-VMware-Boston	UCS Server	Ethernet Adapter	Org-Boston	N/A	2 minutes ago
LAN-CON-Boston	UCS Server	LAN Connectivity	Org-Boston	0	3 minutes ago
Eth-Network-Group	UCS Server, UCS Domain	Ethernet Network Group	Org-NewYork	N/A	3 minutes ago
Eth-Network-Group	UCS Server, UCS Domain	Ethernet Network Group	Org-Boston	N/A	3 minutes ago

Figure 10.
LCP with embedded policies

Viewing those two Ethernet Network Group policies, we can see that, whereas their names are the same, the configurations are different. The Eth-Network-Group policy in the Boston organization has VLANs 100-105 shown in Figure 11, and the Eth-Network-Group policy in the NewYork organization has VLANs 200-205 shown in Figure 12.

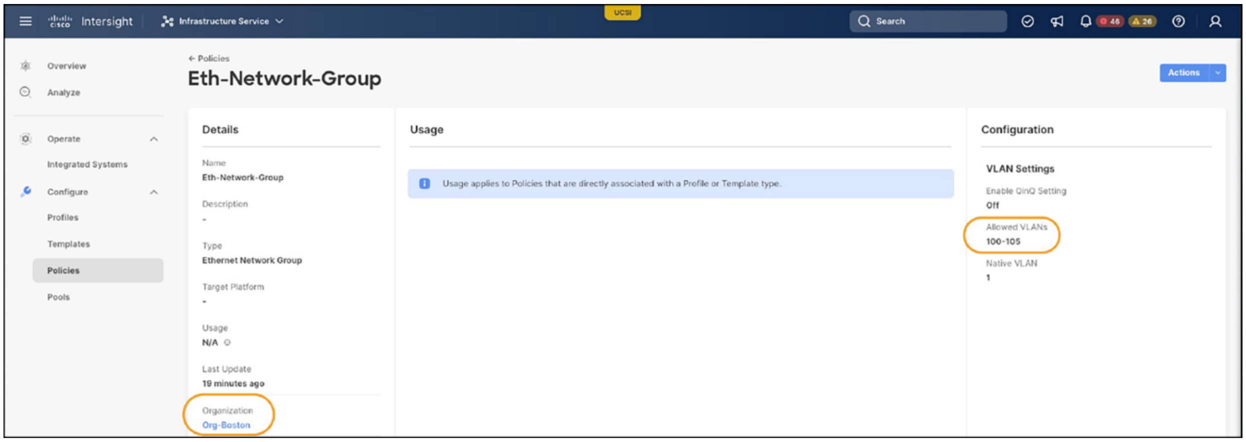

Eth-Network-Group

Details

Name
Eth-Network-Group

Description
-

Type
Ethernet Network Group

Target Platform
-

Usage
N/A

Last Update
19 minutes ago

Organization
Org-Boston

Usage

Usage applies to Policies that are directly associated with a Profile or Template type.

Configuration

VLAN Settings

Enable QinQ Setting
Off

Allowed VLANs
100-105

Native VLAN
1

Figure 11.
Eth-Network-Group policy – Boston organization

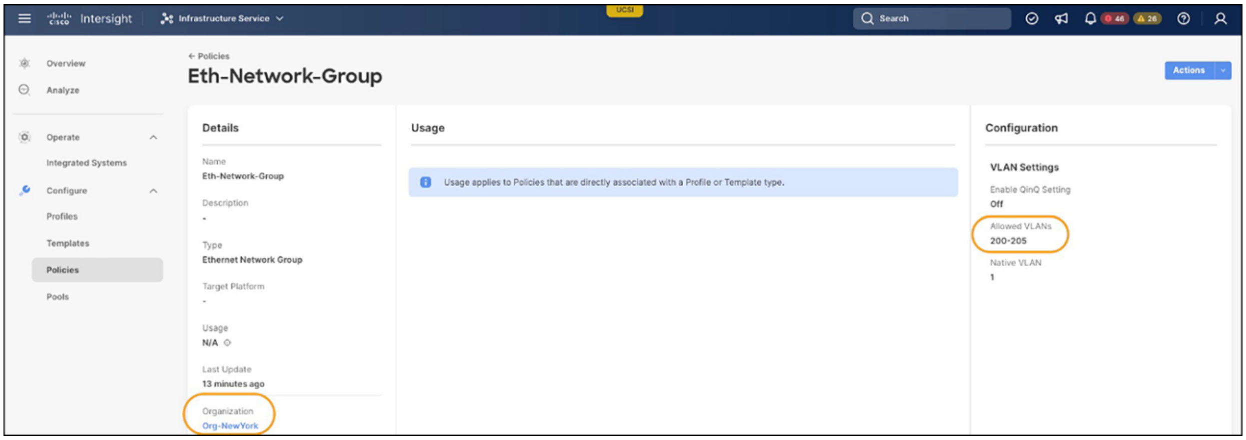

Eth-Network-Group

Details

Name: Eth-Network-Group
Description: -
Type: Ethernet Network Group
Target Platform: -
Usage: N/A
Last Update: 13 minutes ago
Organization: Org-NewYork

Usage

Usage applies to Policies that are directly associated with a Profile or Template type.

Configuration

VLAN Settings
Enable QinQ Setting: Off
Allowed VLANs: 200-205
Native VLAN: 1

Figure 12.
Eth-Network-Group policy – NewYork Organization

Now, let’s clone the LAN-CON-Boston LCP, as shown in Figures 13 and 14.

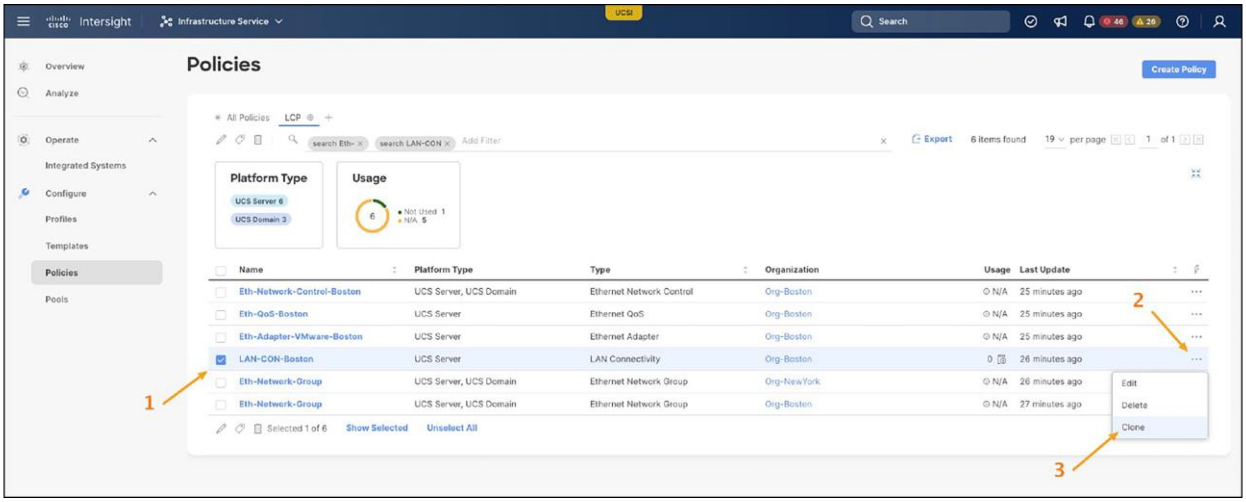

Name	Platform Type	Type	Organization	Usage	Last Update
□ Eth-Network-Control-Boston	UCS Server, UCS Domain	Ethernet Network Control	Org-Boston	N/A	25 minutes ago
□ Eth-QoS-Boston	UCS Server	Ethernet QoS	Org-Boston	N/A	25 minutes ago
□ Eth-Adapter-VMware-Boston	UCS Server	Ethernet Adapter	Org-Boston	N/A	25 minutes ago
☑ LAN-CON-Boston	UCS Server	LAN Connectivity	Org-Boston	0	26 minutes ago
□ Eth-Network-Group	UCS Server, UCS Domain	Ethernet Network Group	Org-NewYork	N/A	26 minutes ago
□ Eth-Network-Group	UCS Server, UCS Domain	Ethernet Network Group	Org-Boston	N/A	27 minutes ago

Figure 13.
Cloning LCP policy – Selection

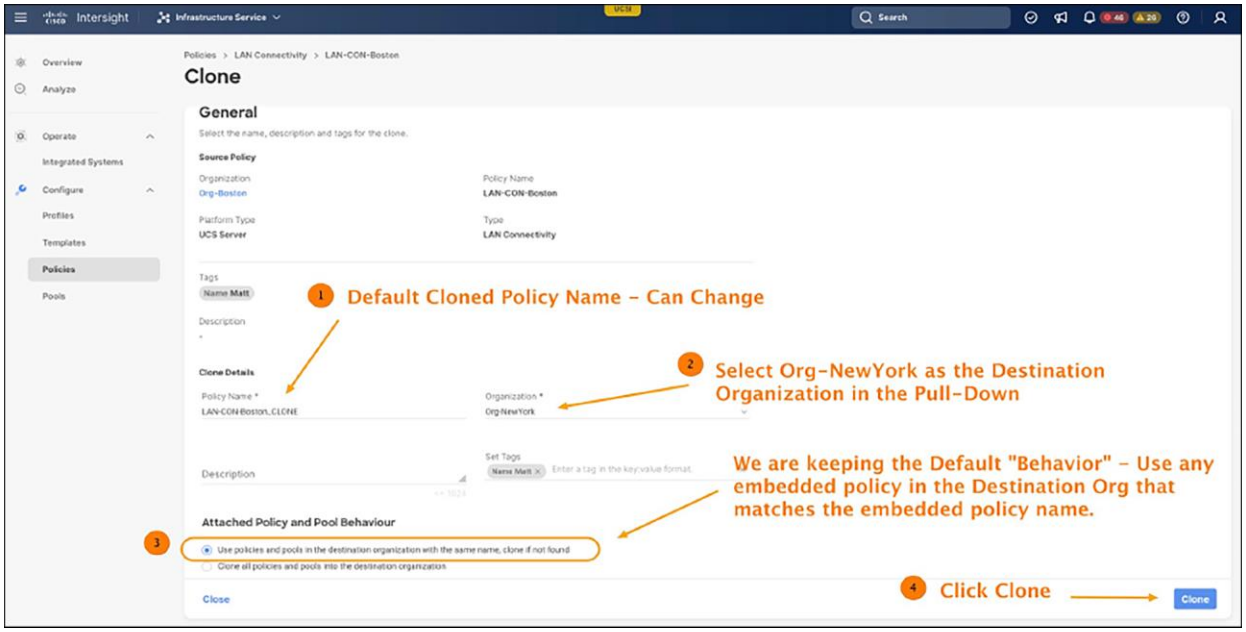

Clone

General

Select the name, description and tags for the clone.

Source Policy

Organization: Org-Boston

Policy Name: LAN-CON-Boston

Platform Type: UCS Server

Type: LAN Connectivity

Tags: Name Matt

Description: -

Clone Details

Policy Name *: LAN-CON-Boston_CLONE

Organization *: Org-NewYork

Description

Set Tags: Name Matt

Attached Policy and Pool Behaviour

● Use policies and pools in the destination organization with the same name, clone if not found

○ Clone all policies and pools into the destination organization

1 Default Cloned Policy Name – Can Change

2 Select Org-NewYork as the Destination Organization in the Pull-Down

3 We are keeping the Default "Behavior" – Use any embedded policy in the Destination Org that matches the embedded policy name.

4 Click Clone

Figure 14.
Cloning LCP policy - Configuration

Upon successfully cloning the LAN-CON-Boston LCP from the Org-Boston to Org-NewYork, we can see that our policy count has doubled. Those embedded policies were cloned to Org-NewYork, as shown in Figure 15.

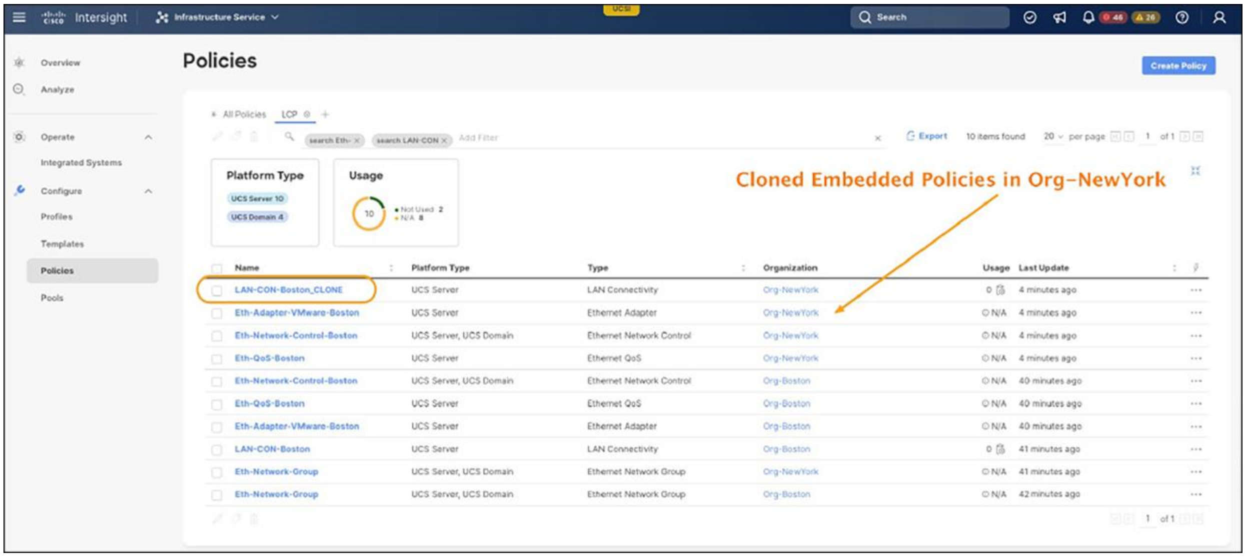

Policies

Cloned Embedded Policies in Org-NewYork

Name	Platform Type	Type	Organization	Usage	Last Update
LAN-CON-Boston_CLONE	UCS Server	LAN Connectivity	Org-NewYork	0	4 minutes ago
Eth-Adapter-VMware-Boston	UCS Server	Ethernet Adapter	Org-NewYork	N/A	4 minutes ago
Eth-Network-Control-Boston	UCS Server, UCS Domain	Ethernet Network Control	Org-NewYork	N/A	4 minutes ago
Eth-QoS-Boston	UCS Server	Ethernet QoS	Org-NewYork	N/A	4 minutes ago
Eth-Network-Control-Boston	UCS Server, UCS Domain	Ethernet Network Control	Org-Boston	N/A	40 minutes ago
Eth-QoS-Boston	UCS Server	Ethernet QoS	Org-Boston	N/A	40 minutes ago
Eth-Adapter-VMware-Boston	UCS Server	Ethernet Adapter	Org-Boston	N/A	40 minutes ago
LAN-CON-Boston	UCS Server	LAN Connectivity	Org-Boston	0	41 minutes ago
Eth-Network-Group	UCS Server, UCS Domain	Ethernet Network Group	Org-NewYork	N/A	41 minutes ago
Eth-Network-Group	UCS Server, UCS Domain	Ethernet Network Group	Org-Boston	N/A	42 minutes ago

Figure 15.
Cloned LCP policy and cloned embedded policies

However, we need to explore the newly cloned LAN-CON-Boston_CLONE LCP further. If we edit and view the LAN-CON-Boston_CLONE Policy and specifically edit one of the vNICs (vNIC0), we can see that the Ethernet Network Control Policy (Eth-Network-Group) used the existing Org-NewYork version of the policy since the policy existed with the same name, and we elected to use any embedded policies with the same name (Default Behavior) during the cloning process. See Figure 16.

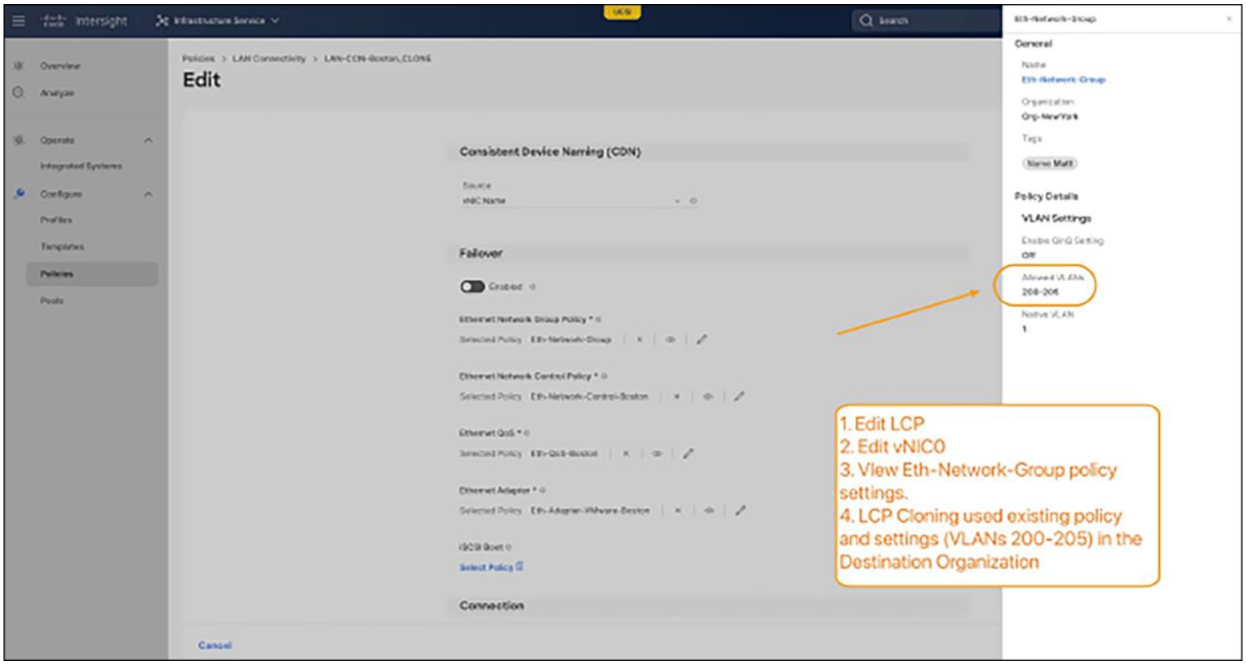

Policies > LAN Connectivity > LAN-CCN-Boston_CLONE

Edit

Consistent Device Naming (CDN)

Source
vNIC Name

Failover

Enabled

Ethernet Network Group Policy *
Selected Policy Eth-Network-Group

Ethernet Network Control Policy *
Selected Policy Eth-Network-Control-Boston

Ethernet QoS *
Selected Policy Eth-QoS-Boston

Ethernet Adapter *
Selected Policy Eth-Adapter-VMware-Boston

iSCSI Boot
Select Policy

Connection

Eth-Network-Group

General

Name
Eth-Network-Group

Organization
Org-NewYork

Tags

Policy Details

VLAN Settings

Enable QinQ Setting
Off

Allowed VLANs
200-205

Native VLAN
1

1. Edit LCP
2. Edit vNIC0
3. View Eth-Network-Group policy settings.
4. LCP Cloning used existing policy and settings (VLANs 200-205) in the Destination Organization

Figure 16.
Verifying policy substitution with same-name behavior

I wanted to purposely call out the behavior of policy substitution if policies have the same name. While this capability can be used to benefit, it also can lead to some undesired results if an administrator does not pay attention to details. Of course, you can always select the alternative cloning behavior and simply clone all embedded pools and policies from the source organization to the destination organization.

Rules - Policy cloning

- The clone option is available on both policy list view and policy detail's view.
- Policy cloning within the same organization and across organizations is supported.
- For policies that reference embedded sub-policies, it is important to consider existing policy names in the destination org when cloning policies across organizations.
- For policies without any embedded policies, or pools, then the policy is cloned to the destination organization with the same policy name.
- **Deep-clone behavior** – For policies with embedded policies, or pools, (LAN/SAN connectivity policies, for example), the user can choose one of the behaviors listed below to deep clone those policy references.

1. **Use policies and pools in the destination organization with the same name; clone them if not they are not found** – If any policies in the destination org have a name that **exactly matches (case-sensitive)** the embedded policies referenced in the original source org cloned policy, then those embedded policies in the destination org will substitute for the intended embedded policies in the source org (keep in mind that the embedded policies with the same name could have different settings – be aware of this). If there is no policy in the destination organization that matches any embedded policy by name, then the intended source org's embedded policies will be cloned with the same name. Pool references will always be matched by name. If not found, the pool will be cloned in the destination organization, **but no identifier blocks will be created. (To prevent any unintended ID conflicts, simply add new, unique blocks to the cloned pool reference).**
2. **Clone all policies and pools into the destination organization** – New policies will be created for the source and all attached policies. If a policy of the same name and type already exists in the destination organization or any organization with which it shares policies, a clone will be created with the provided suffix added to the name. Pool references will always be matched by name. If not found, the pool will be cloned in the destination organization, **but no identifier blocks will be created**.
3. **The deep-clone options do not appear in the UX for those policies with no embedded pools or policies.** ****
    - Each cloned policy results in its own autonomous policy with its own set of configurations.
    - Currently, you can clone only one policy instance at a time.
    - Cloning a policy is not applicable for Kubernetes clusters.
    - Cisco HyperFlex® clusters cannot be cloned across multiple organizations.

Conclusion

Cisco Intersight Managed Mode (IMM) is ever evolving and becoming more and more powerful to manage Cisco UCS servers at scale. The original design goals are steadfast and intact, which is to provide more capability, simplification, visibility, and flexibility as compared to Cisco UCS Manager and Cisco UCS Central Software.

As discussed in this guide, adding the capabilities of IMM Configuration Sharing Across Organizations and Cloning Across Organizations greatly enhances the design capabilities of large-scale policy models with IMM and creates better operational efficiencies throughout.

Building shared organizational models directly reduces the number of touchpoints with policy creation and editing, thus reducing policy sprawl. Now, you can take that golden policy that is configured properly in a shared organization and with a few clicks, share that organization and that policy with another consuming organization for immediate use. You can also leverage different layers of RBAC and administrative control by having more senior administrators create and edit the policies and share them with more junior administrators to consume those policies in their server profiles.

With policy cloning, this capability does not require organizational sharing and is best used when duplicate policies are quickly needed in different organizations, perhaps with different users and roles having access and control over those policies in those separate organizations. A similar policy can be cloned across organizations and then edited as required in the destination organization, versus creating the policy from scratch.

As with the majority of IMM features and functionalities, these simply get pushed to the cloud platform where they appear and can be used. There is no need to upgrade any infrastructure or blade firmware.

These capabilities, coupled with the underlying IMM architecture that allows the nondisruptive flexibility to morph existing siloed organizational models into at-scale, shared organizational models, are powerful and agile, and are a testament to Cisco Intersight.

Additional resources

Configuration Sharing Across Organizations

Cloning a Policy Across Organizations

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 **https://www.cisco.com/go/offices**.

Printed in USA C11-3956423-00 10/23