



UCS Manager Configuration Best Practices (and Quick Start)

The screenshot displays the Cisco Unified Computing System Manager (UCS Manager) interface. The left sidebar shows a tree view of the configuration hierarchy, including Service Profiles, Policies, and Pools. The main pane shows the 'Service Profiles' tab with a table of profiles and their associated servers.

Name	Overall Status	Assoc State	Server
Service Profile ESX-1	ok	associated	sys/chassis-1/blade-1
Service Profile ESX-2	ok	associated	sys/chassis-1/blade-2
Service Profile ESX-3	ok	associated	sys/chassis-1/blade-3
Service Profile Linux-1	unassociated	unassociated	
Service Profile Linux-2	unassociated	unassociated	
Service Profile W2K8-1	unassociated	unassociated	
Service Profile W2K8-2	unassociated	unassociated	
Service Profile Windows_...	unassociated	unassociated	

Below the table, an 'Associative State' pie chart shows the distribution of profiles: 3 associated (green), 4 unassociated (grey), and 1 disassociated (yellow).

UCS TME White Paper
by

Jeff Silberman

Version 0.85

What You Will Learn

The UCS Manager (UCSM) presents a new paradigm for server management. Among the most relevant attributes of UCSM are the following:

- Foundation for stateless/utility computing model
- Policy-based management
- Full inventory, auto-discovery and device management
- Simple association of logical server/application images to physical server/blades

The focus of this White Paper will be on rapid deployment and on the most direct path to working in a stateless-server SAN-boot environment. In support of this utility-compute model, a number of concepts will be discussed relating to the UCSM elements that promote leverage, reuse and inheritance within the UCS management framework.

A highly distilled summary can be found in the “Ultimate Quick Start Guide”, in the Appendix.

First Time Setup of UCS Manager (UCSM)

Bringing up a UCS system for the first time is done through the serial console. A “First Time Wizard” will guide you through the standard questions (hostname, IP address, netmask, default gateway, etc).

Standalone or Clustered

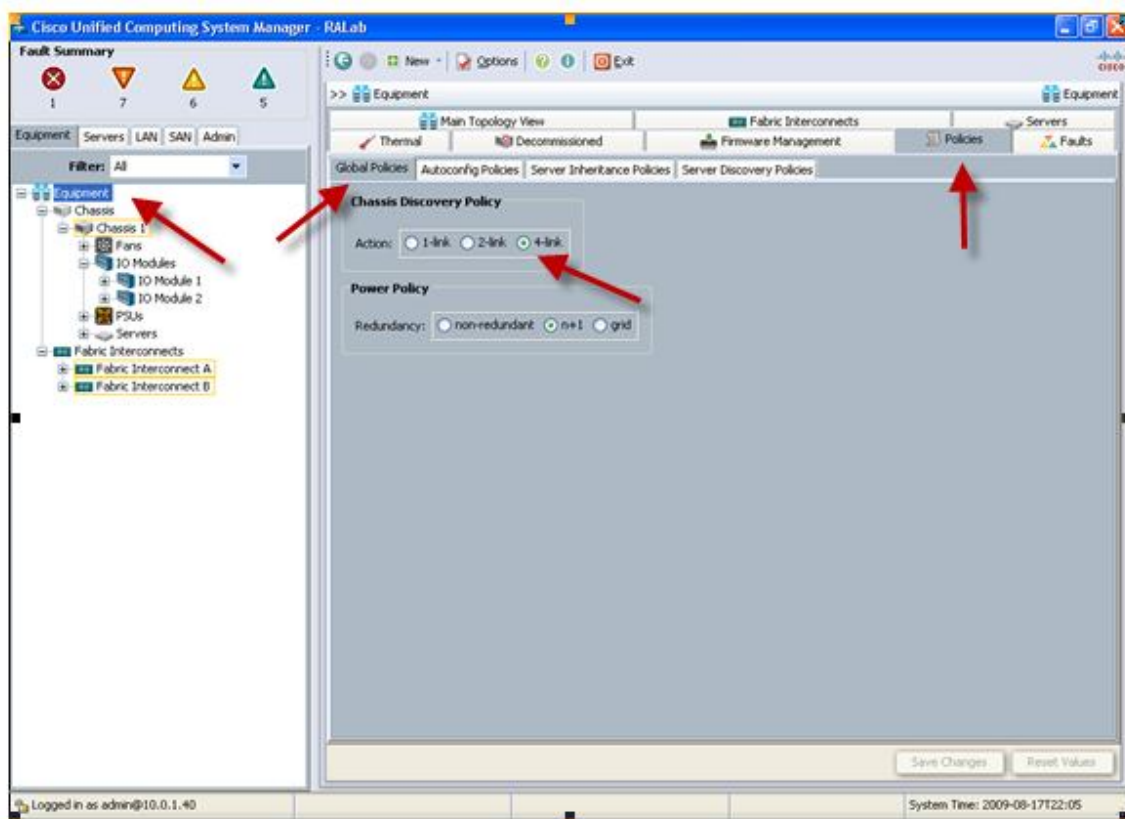
UCS systems can be standalone (single fabric-interconnect) or more typically clustered (two clustered fabric-interconnects). The clustered configuration will present an “A” side (primary) fabric interconnect and a “B” side (secondary) fabric interconnect for management/control purposes (the management plane is active/passive; the data plane is active/active). When configuring a clustered configuration, you will need to allocate 3 IP addresses in the management/admin subnet: 1 for each fabric interconnect, and one for the main “clustered-failover” interface that enables the same management IP address to be active, regardless of which system is acting as primary. Also, a “-A” and “-B” string are implicitly added to the end of hostname, and do not need to be specified explicitly.

After completing the “First Time Wizard” setup, UCSM can be easily managed from the UCSM GUI, by pointing a browser to the UCS management IP address.

Here are the steps that **must** be done, prior to general operations:

1. Chassis Discovery Policy

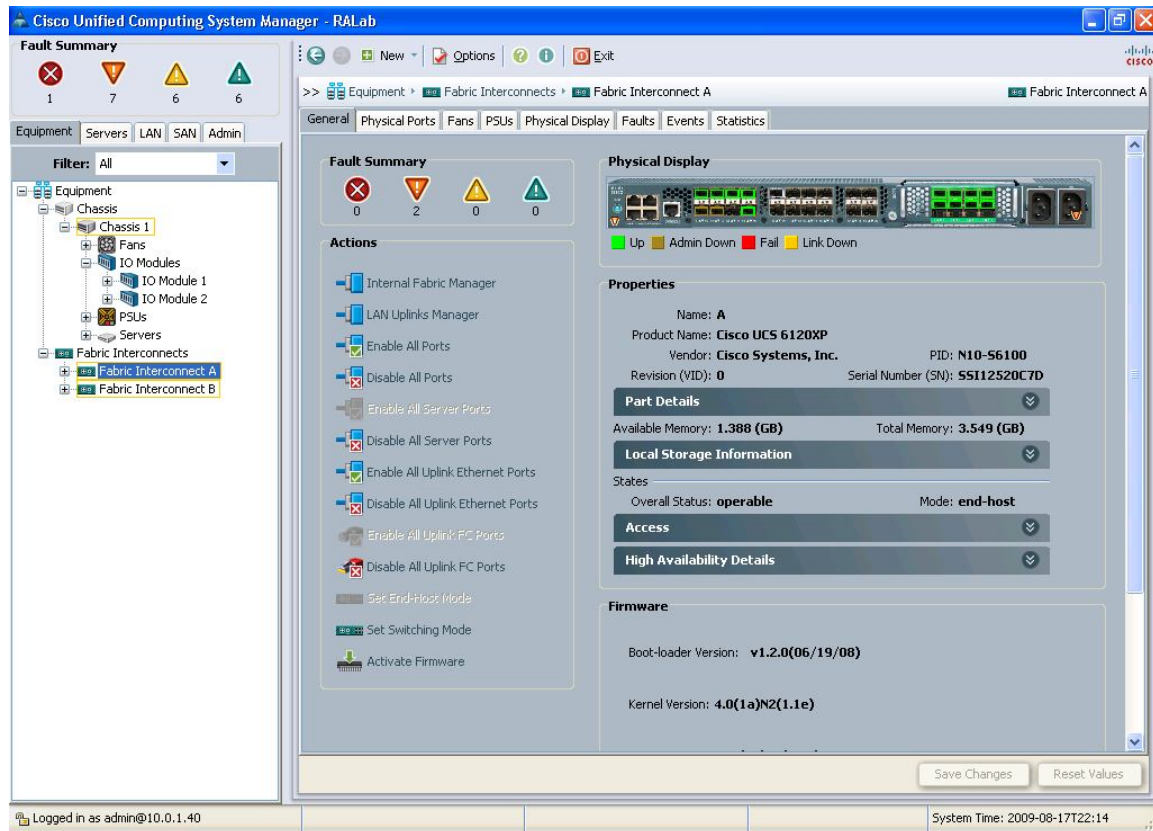
Chassis Discovery Policy specifies the number of connections between the Fabric Extender (FEX or IOM) and the Fabric interconnect (FI). The UCS Manager cannot intuit this value. Furthermore, this value is constant across all chassis within a given UCS Management domain. To set this from the Equipment tab, bring "Equipment" in to scope, and then follow "Policies->Global Policies" as below.



2. Enable Server and Uplink Ports

The "Fixed Module" supports two types of connections: those going "North" to the LAN are configured as "Uplink Ports"; those going "South" to the chassis FEX (or IOM) connections are configured as "Server Ports".

Individual ports can be configured/enabled by right-clicking over the desired ports, when the Fabric Interconnect is in scope, as below.



Setting the Chassis Discovery Policy plays a critical role in the ability of UCS to scale by adding chassis and servers. Once this policy is set, then whenever additional Server Ports are configured, the UCS Manager will automatically discover and inventory any attached equipment, without requiring manual intervention.

3. Create Management IP Addr Pool

Each physical blade is capable of supporting remote KVM and remote Media access (CD, DVD and USB). This is made possible by associating IP addresses for the cut-through interfaces that correspond to each blade's BMC. Typically, these addresses are configured on the same subnet as the management IP address of the UCSM. This pool is created from the Admin tab, by selecting "Management IP Pool".

4. Ensure uniform/established firmware versions


UCSM offers the ability to quickly/easily report on the F/W versions for all major H/W components. From the "Equipment" tab, bring "Equipment" in to scope and select "Firmware Management -> Installed Firmware", as shown below. In general, all Adapter Cards, BMC, IOM's, Fabric Interconnects, and UCS Manager should be at the same version.

The screenshot shows the Cisco Unified Computing System Manager (UCSM) interface. The left sidebar displays a tree view of the system components, with 'Equipment' selected. The main pane shows the 'Installed Firmware' tab, which displays a table of firmware versions for various components. The table columns are Name, Running Ver..., Startup Version, Backup Ver..., Update Sta..., and Activate Sta....

Name	Running Ver...	Startup Version	Backup Ver...	Update Sta...	Activate Sta...
UCS Manager	1.0(1e)	1.0(1e)	N/A	N/A	ready
Chassis 1					
IO Modules					
IO Module 1	1.0(1e)	1.0(1e)	1.0(0.172)	ready	ready
IO Module 2	1.0(1e)	1.0(1e)	1.0(0.172)	ready	ready
Servers					
Server 1					
Interface Cards					
Interface C.1.0(1e)	1.0(1e)		1.0(0.172)	ready	ready
BIOS		55500.86B.01.00.0036-191.061320...	N/A	N/A	ready
BMC Controller	1.0(1e)	1.0(1e)	1.0(1e)	ready	ready
Server 2					
Interface Cards					
Interface C.1.0(1e)	1.0(1e)		1.0(0.172)	ready	ready
BIOS		55500.86B.01.00.0036-191.061320...	N/A	N/A	ready
BMC Controller	1.0(1e)	1.0(1e)	1.0(1e)	ready	ready
Server 3					
Interface Cards					
Interface C.1.0(1e)	1.0(1e)		1.0(0.172)	ready	ready
BIOS		55500.86B.01.00.0036-191.061320...	N/A	N/A	ready
BMC Controller	1.0(1e)	1.0(1e)	1.0(1e)	ready	ready
Server 4					
Interface Cards					
Interface C.1.0(1e)	1.0(1e)		1.0(0.172)	ready	ready

Towards Utility Computing with Stateless Servers

For datacenters wishing to develop a utility computing model, UCS provides the infrastructure to support stateless server environments. Utility computing, in this regard, allows for logical servers (OS/applications) to be run on any physical server or blade, in a way that encapsulates traditional H/W identifiers (WWN's, MAC's, UUID's) as an integral part of the logical server identity --- or UCS "Service Profile".



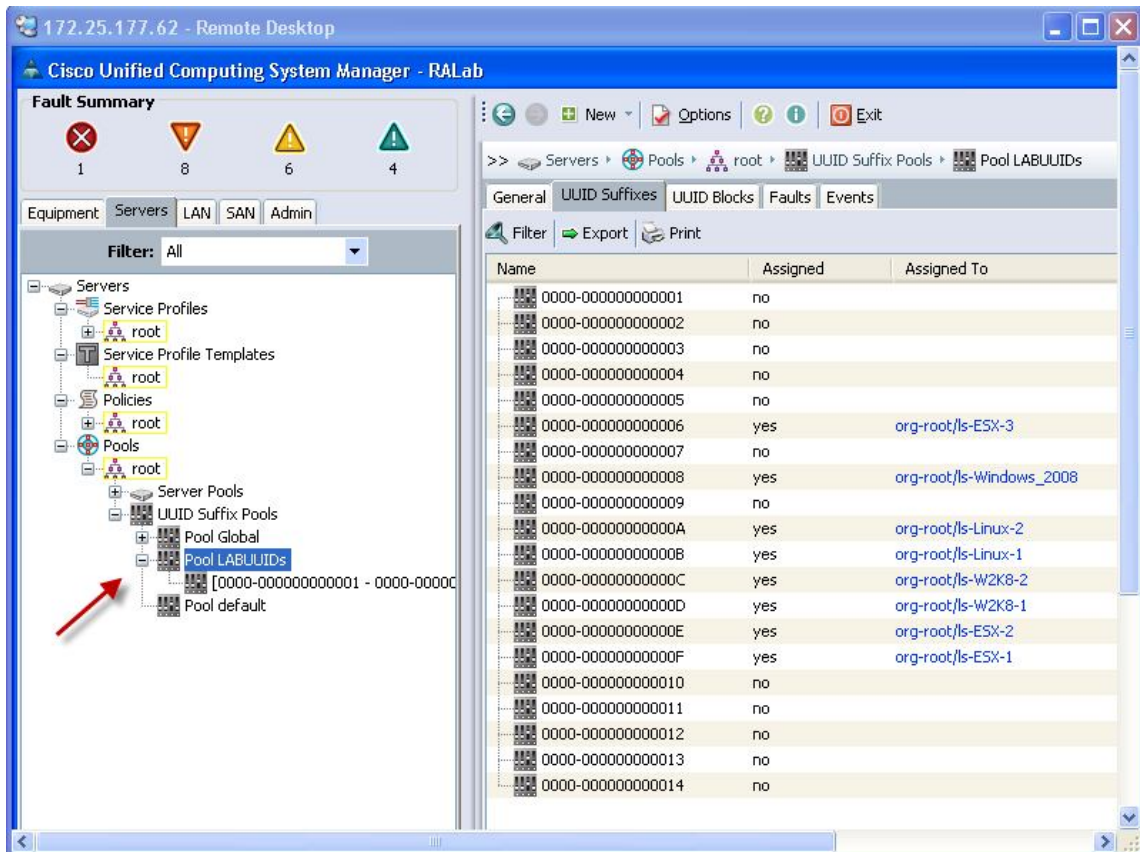
The foundation for utility computing and service-profiles are the building-block elements ("pools" and "policies") that can then be captured for reuse in a "template". Furthermore, VNIC and VHBA templates can be referenced for use in higher-level "service-profile templates". And to make rapid-provisioning as easy as possible, "service-profile templates" can then be used to rapidly instantiate actual service-profiles (possibly even automatically associating those instantiated service-profiles to actual physical servers/blades).

Pools

Pools are the base building block for providing unique identification of hardware resources. As the basis for the utility compute model, they allow "service-profiles" to be associated to any blade, while still providing the exact same ID/presentation to the upstream LAN/SAN. There are three sets of Pools used as part of Best Practices in UCS:

- **UUID Pools** --- Provides ID's, similar to a serial number or service-tag, and used most prominently by VMware
- **WWNN / WWPN Pools** --- Provides unique ID's for Fibre-Channel resources on a server (FC "nodes" and "ports").
- **MAC Pools** --- Provides unique ID's for virtual network interface ports.

These pools are all functionally organized, with UUID Pools maintained from the “Server” tab, WWNN/WWPN Pools, maintained from the “SAN” tab, and MAC Pools maintained from the “LAN” tab.



Management Consideration

UCS management domains may co-exist along with many other UCS domains and/or non-UCS servers, all of which may have their own respective sets of unique hardware identifiers and/or pools. Having duplicate WWN's and MACs presented to the LAN/SAN can be a major source of admin/operator problems. **Therefore:** Maintain a **single/centralized catalog** of unique IDs (MACs/WWNNs/WWPNs/UUIDs). Such a catalog must be defined and maintained at a central layer, well-above all individual UCS and non-UCS management domains. One effective strategy would include embedding unique “UCS domain IDs” as high-order bytes within WWNN/WWPN/MAC/UUID identifier ranges.

Best Practice : Pools

Define and use Pools as a standard practice. Make sure that:

- UUID pools get referenced when creating "service-profiles"
- MAC pools get referenced when creating VNICs
- WWNN pools get referenced when creating "service-profiles"
- WWPN pools get referenced when creating VHBAs.

Similarly, pools should also be referenced when creating any corresponding template objects (VNICs, VHBAs, service-profiles).

For basic environments without multi-tenancy concerns, populating and using the respective "default" pools may be sufficient for most needs.

Server Pools

Server Pools provide a way to partition/segregate physical blades in to different groupings/sets. The grouping criteria are left to the administrator. Possible criteria may focus on either physical server capability (CPU speed, memory size, local disks or not), logical business organizations (e.g. "Marketing", "Finance", etc.), specific customers being hosted, or specific geographies being served. Server Pools can also be built from individual/arbitrary blades. Once Server Pools have been defined, you can then use them to associate individual service-profiles or to rapidly instantiate multiple service-profiles from a service-profile template.

Policies

Policies provide a key mechanism for making the server environment "rule-driven". Defining and using the rich set of policies available in UCS enables administrators/operators to have fine-grained control, and to allow for greater degrees of automation. Below are the most common policies that should be used regularly.

Boot Policy determines boot devices/methods and boot order.

Suggested Practices:

- Have CDROM as the first in the boot order, for emergencies/recovery
- For SAN boot, define separate boot policies for each storage array that would be serving boot LUNs.
- For Network/PXE boot, define LAN/VNIC boot last in the boot order, following either SAN or Local boot.

Host F/W Policy is a way of associating qualified/well-known versions of BIOS/Adapter-ROM/Local-Disk-controller with “service-profiles”. Best Practice is to create one, based on the latest packages that correspond to the general UCS software release/version, and to reference that Host F/W Policy for all “service-profiles” and templates created.

Local Disk Policy specifies if/how to configure any local disks on the blade. Best Practice would be to specify “No Local Storage” for SAN boot environments, thereby precluding any problems at “service-profile” association time, when local disks might be present.

Scrub Policy determines what to do with local disks and BIOS upon “service-profile” disassociation. The default is no scrubbing. Best Practice is to set this to scrub the local disk, especially for Service Provider oriented customers, and environments where LAN/PXE install to local disk would be used.

Templates

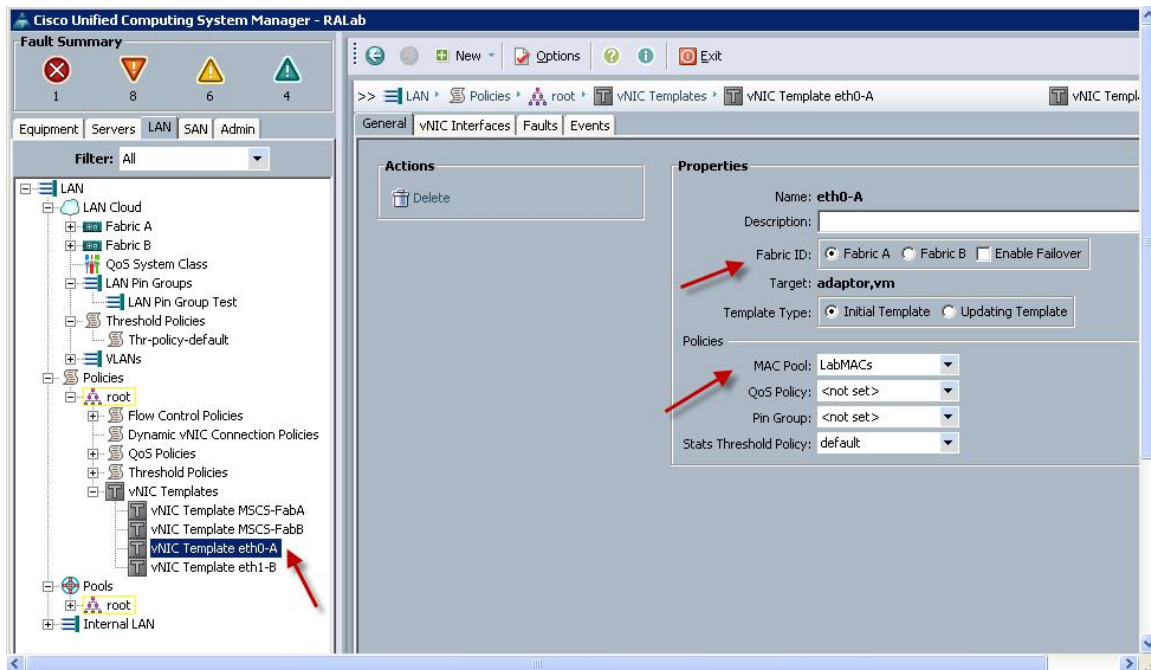
UCS Manager provides “templates” for the primary objects (VNICs, VHBAs and service-profiles), to facilitate reuse and rapid-deployment. Properties, attributes and policies can be defined at the template level, allowing for rapid instantiation and provisioning.

Best Practices:

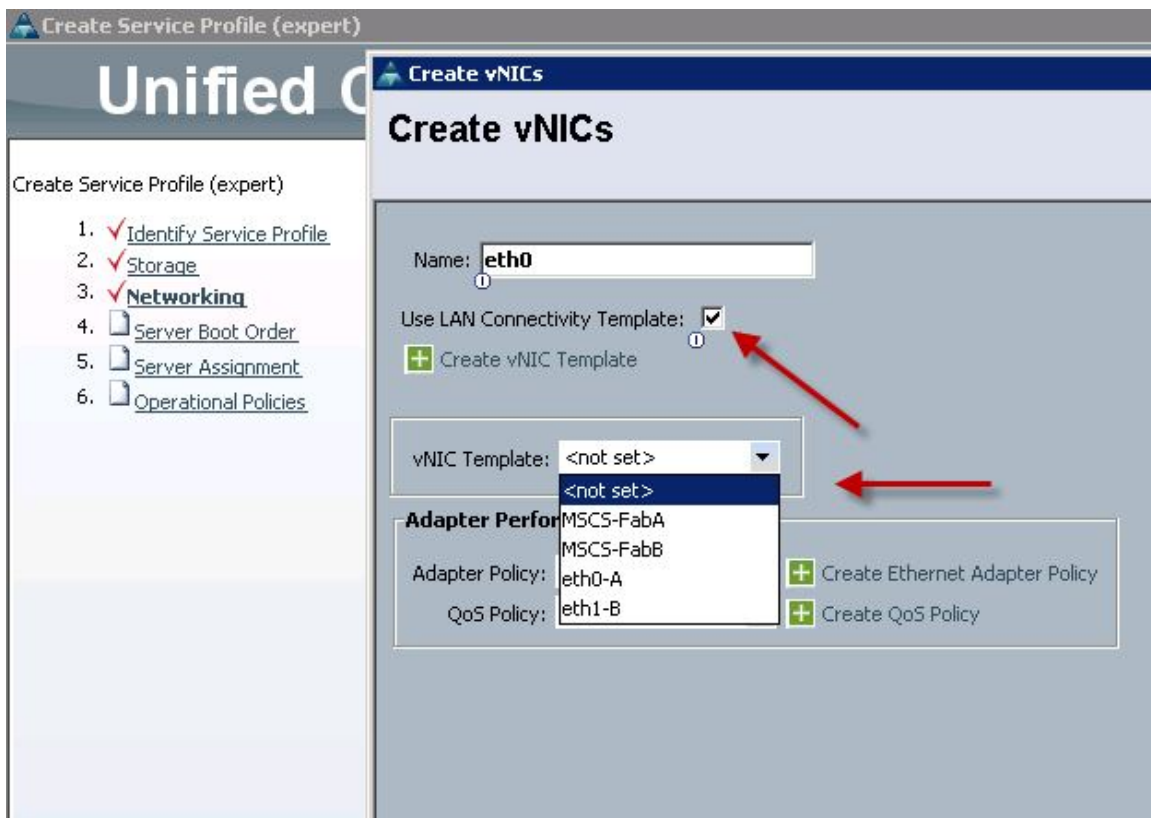
- Use “expert mode” when creating service-profile templates, so as to achieve the optimum level of control/definition in the utility-compute model.
- When creating “templates”, draw from the subordinate “pools” and “policies” that have been previously defined.

VNIC and VHBA Templates

VNIC/VHBA resources are always associated with specific fabric interconnect (A-side or B-side). A typical service-profile would have at least 2 VNICs/VHBAs --- one bound to each side. VNIC (or VHBA) templates can be used to encapsulate both the MAC Pool (or WWPN Pool) association as well as the Fabric ID.



Once a vNIC/VHBA template is defined, it can be referenced through "expert-mode" service-profile creation by clicking the "Use LAN (or SAN) Connectivity Template", as show below.



When creating VNIC templates, be sure to include all possible VLAN mappings, since these typically can not be modified.

For VHBA templates, only one VSAN to VHBA mapping is allowed.

Service-Profile Templates

Service-Profile Templates provide the umbrella for tying together all the various subordinate elements (Pools, Policies, VNIC/VHBA Templates), mentioned previously. Once created, the service-profile template allows for easy/rapid instantiation and server provisioning. There two types of Templates : "Initial" (the default) and "Updating". After a template has been created and instantiated as a service-profile, the template type governs subsequent update/edit behavior and capabilities. Service-profiles created from an "initial" template would need to "unbind" from the template, if specific local changes are required to the service-profile. Correspondingly, initial templates can not be edited while any service-profiles are currently bound to them.

A Word of Caution: Updating Templates

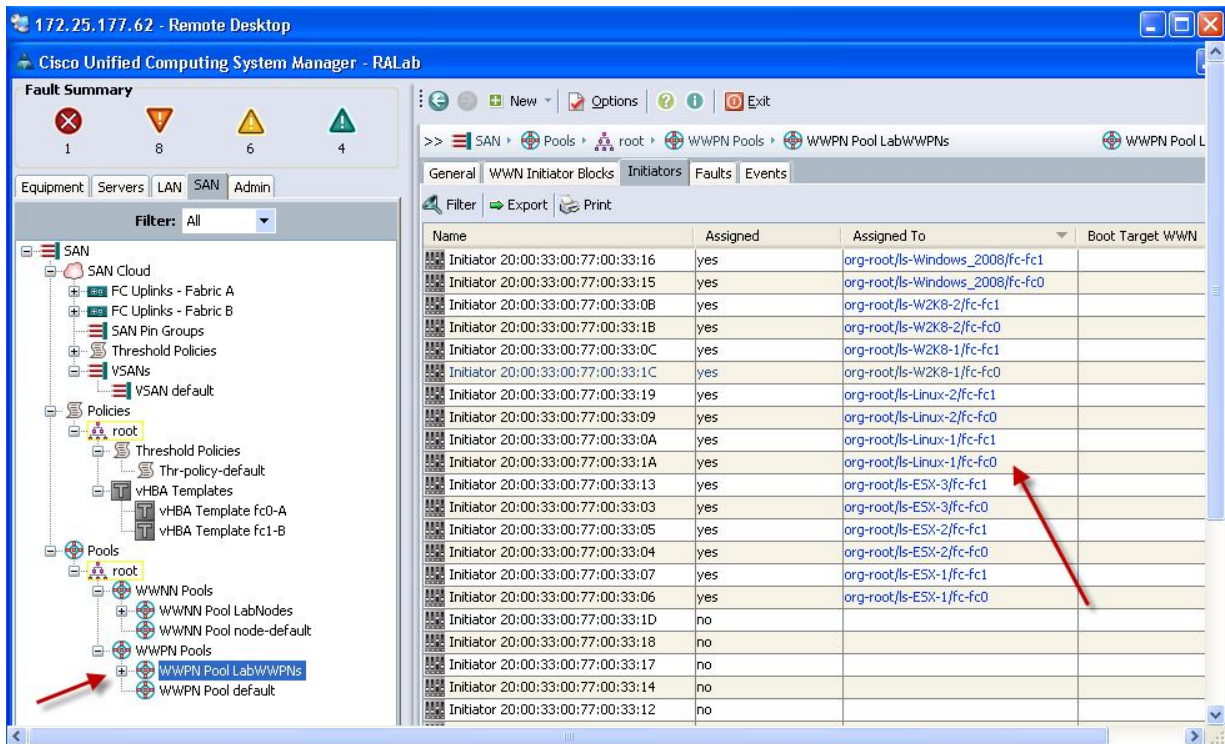
UCS Templates have a very powerful property called "Updating Templates". Updating Templates allow changes in the template, such as pools or policies, to be propagated **immediately** to any higher-level template or service-profile (instantiated or not). **However** : changes reflected to instantiated service-profile via Updating Templates may well cause a service interruption and/or server reboot. Therefore, Updating Templates should be used with the **greatest level of awareness and caution**. Updating Templates can be a tremendous time-saving asset during a scheduled maintenance window. However, Updating Templates can also have disastrous results when updating during normal operations.

SAN install/boot

SAN-install and SAN-boot are both major components in delivering the stateless-server utility model. As powerful as SAN-boot may be, there are many considerations for preparing a reliable/consistent/repeatable SAN-boot environment. Following are some common Best Practices.

Focus on WWPN as "the key"

When performing SAN configuration, the WWPN of the service-profile is the most significant "key" for interfacing with the SAN switch (zoning) and SAN storage array (LUN masking). The picture below illustrates how to see the current mapping between WWPN pool elements and the service-profiles to which they are assigned.



Use single-initiator based zoning --- not "open zoning".

Ensure that all used storage array ports and the VHBA WWPNs are on the same VSAN and are zoned together. Use the "single-initiator" based zoning as a best practice --- avoid placing all VHBA WWPNs and storage array ports in a single "open zone".

Use the FC switch "name server" to verify connectivity

After a service-profile has associated to a blade, viewing the SAN "name server" is a way to verify proper connectivity between the UCS 6100 fabric interconnect and the SAN in general. On Cisco MDS switches, this is done with "show flogi database" from the CLI. All ports that have logged on to the FC fabric will be visible. If intended VHBA and WWPNs are not present in this table, then the problem points towards a UCS configuration problem (service-profile did not associate properly; intended VHBA/WWPN was not referenced by the service-profile; ...)

Avoid mixing heterogeneous storage array types

Mixing storage array types in a SAN-boot environment can be problematic, especially if the storage arrays are not capable of providing precise "LUN mapping" --- specifying which LUN number gets presented to the host. Servers will only boot from LUN 0. Results may be indeterminate if multiple storage arrays present multiple "LUN 0"s.

When things go wrong ...

... look from these valuable vantage points, to assess the direction of the problem:

- FC Name Server
- Storage Array
- HBA Option ROM (Scan FC devices)
- Is the upstream SAN switch enabled/configured for NPIV?
- Are local disks plugged in to the blade?

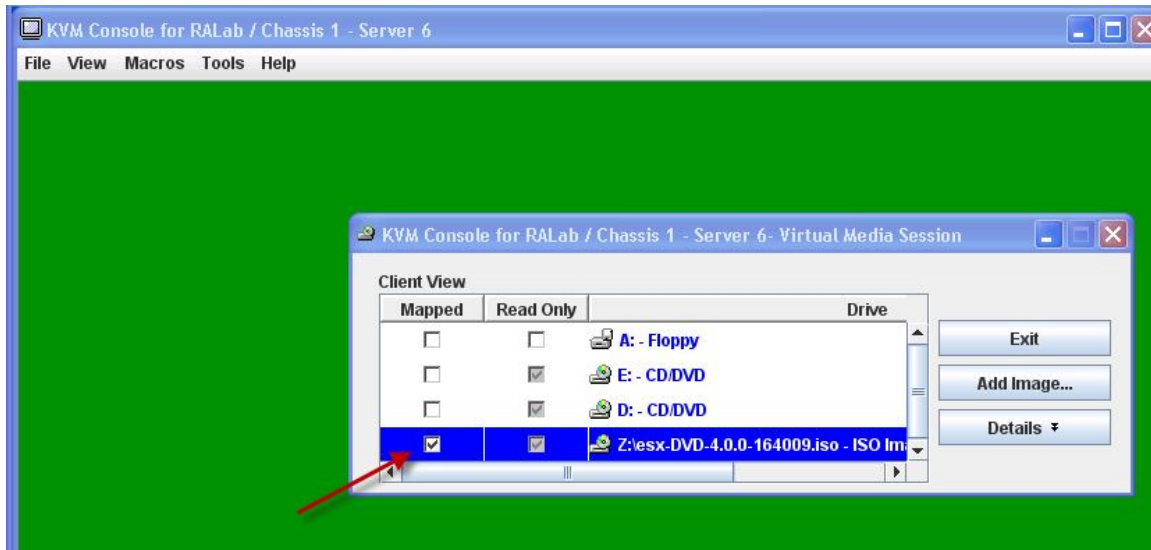
Observation	Likely Problem Causes
VHBA or Storage Array not present in the FC Name Server	<ul style="list-style-type: none">- Physical Cabling- Ports on 6100 or SAN switch are not "enabled"
VHBA visible from the Name Server but not from the Storage Array	<ul style="list-style-type: none">- Zone or VSAN misconfigurations- Storage array not properly cabled or configured in the Zone/VSAN.
Storage Array visible from HBA Option ROM, but LUNs are not present or show up as "LUNZ" for EMC.	<ul style="list-style-type: none">- LUN masking is not complete on the storage array
VHBA/Host shows as "registered" on a Clariion, but other hosts with same WWPN appear as logged in	<ul style="list-style-type: none">- In multipath environments, all paths must be explicitly registered for EMC Clariion
SAN install proceeded correctly, but system will not subsequently boot	<ul style="list-style-type: none">- Local disks are plugged in to blade- Advanced storage options may be needed, during installation phase

Imaging and Provisioning

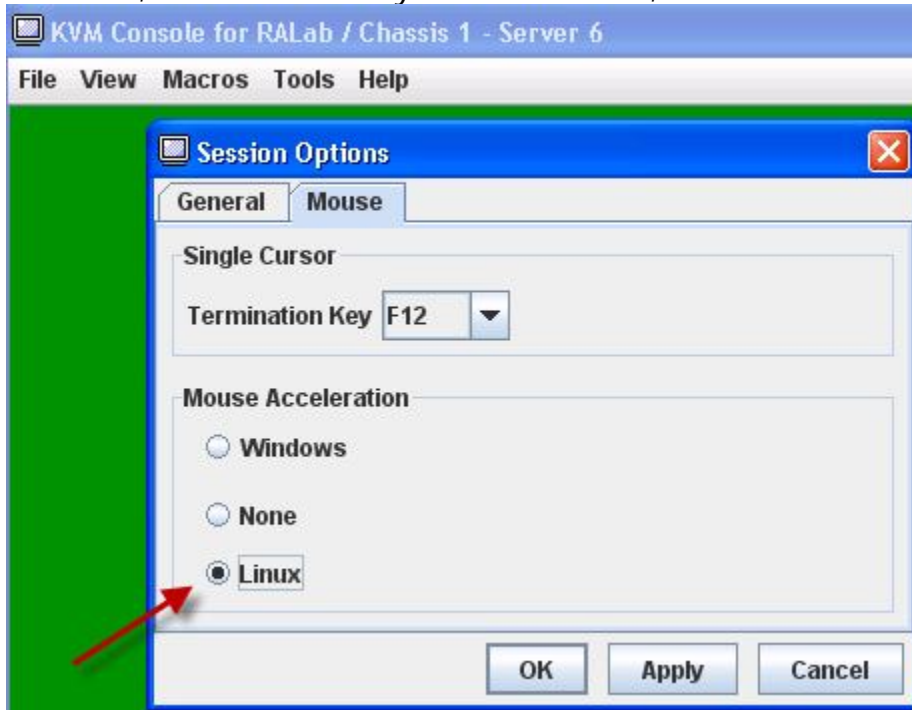
Once all the UCS Manager building blocks are understood, the actual provisioning and imaging of servers becomes one of the simpler tasks. UCS provides a "Remote Media" capability that allows any network-connected ISO image or physical CDROM to be mapped as the local CD drive for any physical blade.

To have a network-connected ISO image mapped as the local CD for a blade, open the KVM Console for a given blade, and select "Launch Virtual Media" from the "Tools" menu. Click "Add Image" and navigate

to the desired ISO image file. Once selected, be sure to click the “Mapped” check-box. Do not “Exit” Virtual Media, until the file is no longer needed. Once the Virtual Media is mapped to the blade, then “Change Service Profile Association” to the desired blade, and begin the OS installation.



Another convenient feature of the KVM Console is the ability to change the mouse acceleration, based on OS type. From the “Tools” menu, select “Session Options”, and click on the “Mouse” tab. For Linux and Vmware, select “Linux” style acceleration, as shown below.



Isolation

UCS addresses isolation requirements with the following primary components:

- VLANs
- VSANs
- Pin Groups

VLANs provide the foundation for network-based isolation. North-bound VLANs can be declared in the “LAN” tab, under the “LAN Cloud” for any VLANs to which UCS will connect. Any declared VLAN can then be referenced when creating VNICs or VNIC templates. Multiple VLANs can be associated to a given VNIC.

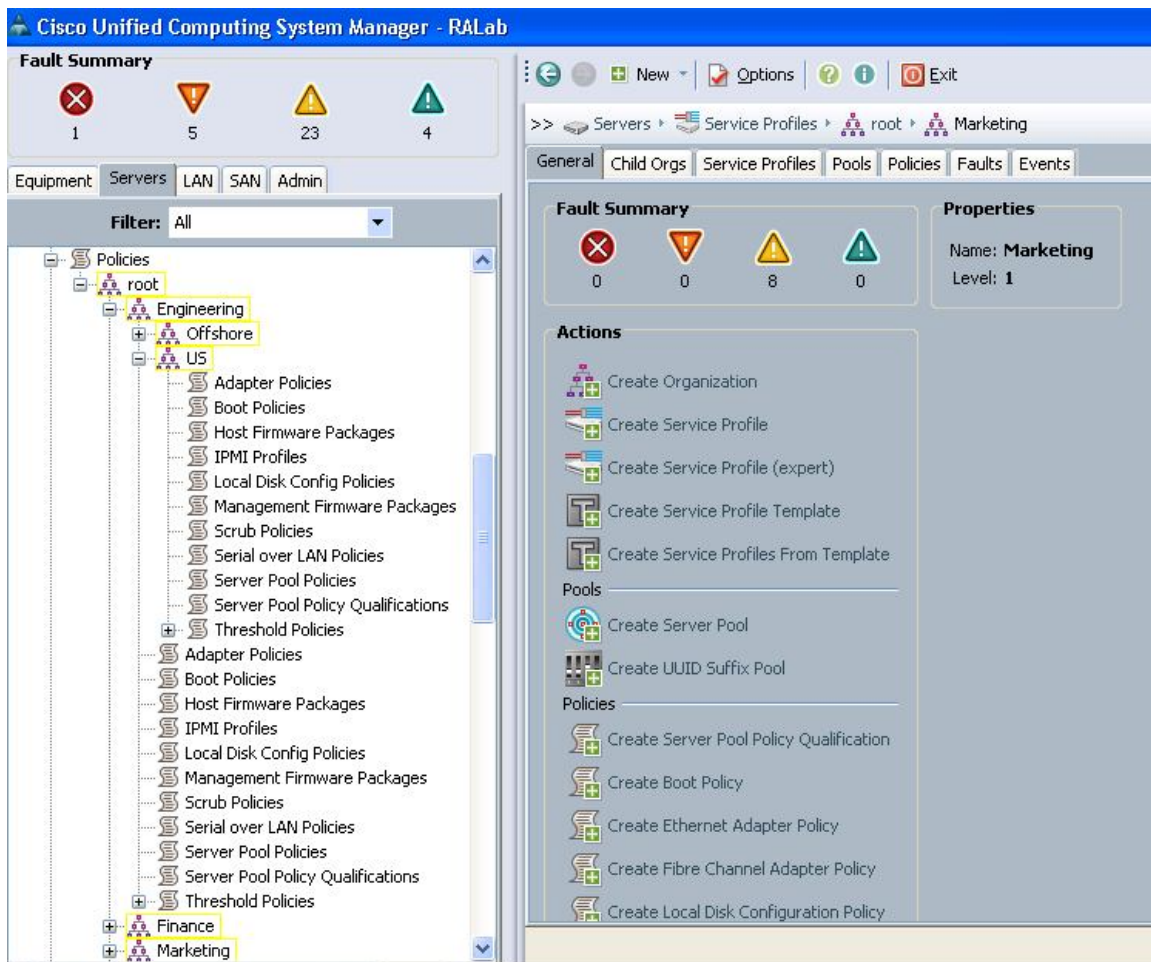
VSANs provide corresponding storage-based isolation capabilities. North-bound VSANs can be declared in the “SAN” tab, under the “SAN Cloud” for any VSANs to which UCS will connect. Any declared VSAN can then be referenced when creating VHBAs or VHBA templates. However unlike with VLANs, VHBAs can only associate with a single VSAN.

Pin Groups provide isolation to specific northbound interfaces for both network and storage traffic. Pin Groups are defined under “LAN Cloud” or “SAN Cloud” respectively as a collection of one or more uplink ports on either/both fabric interconnects. Once defined, Pin Groups can then be referenced as the target data path any given VNIC or VHBA (or template), to guarantee that all traffic associated with a given VNIC/VHBA will be isolated to pre-scribed physical uplink ports.

Organizations

The UCS Manager constitutes a single management domain, defined by the clustered pair of fabric interconnects and all the connected UCS chassis and blades. However, UCS administrators may want to create smaller management sub-domains, to partition resources logically, to scale management more effectively, to support multi-tenancy, or for whatever reasons. UCS allows management sub-domains to be created through the use of “Organizations”.

Organizations exist hierarchically, with “root” as the top-level default. Any sub-organizations created will have a corresponding set of Pools/Policies/Templates within their management scope, as can be seen below:



Two important properties apply to the management of hierarchical, “Organizations”: Inheritance and Override.

Inheritance applies to the set of Pools/Policies/Templates that exist in the direct hierarchical path of an organization. Consider the organizational schema highlighted above, with “Offshore” and “US” both as organizations under the “Engineering” organization. Any Pools/Policies/Templates that are defined at the “root” or “Engineering” levels are available to be used/inherited by any objects in the “Offshore” or “US” organizations. However, the inverse does not apply --- individual Pools/Policies/Templates defined in the “US” organization are local to “US” and cannot be used by parent nor peer organizations.

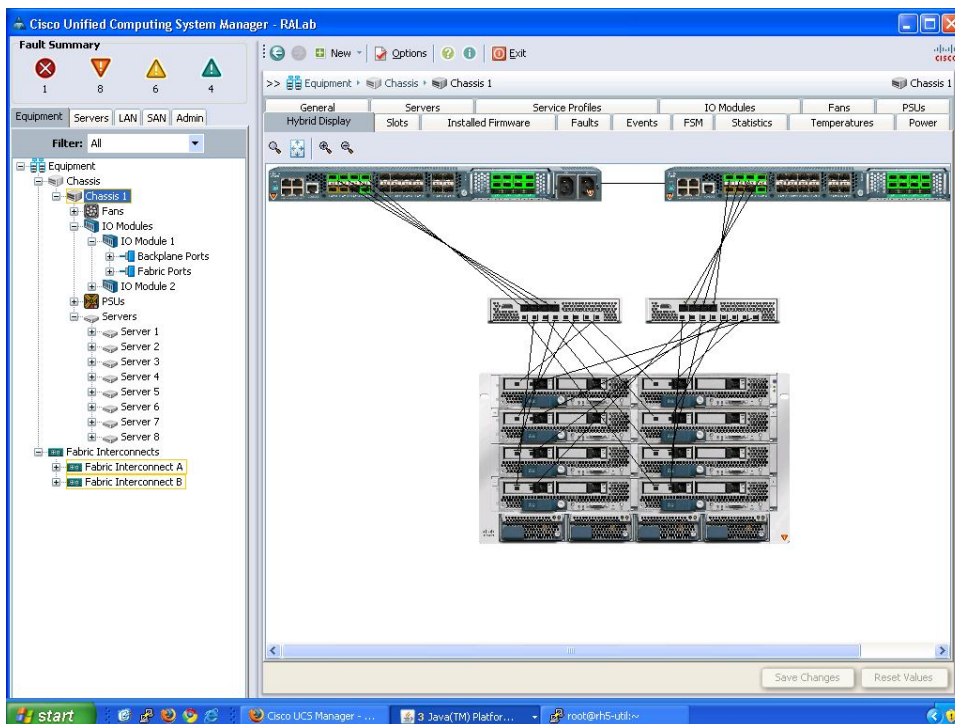
Override characteristics give certain autonomy to local organization administrators. One example is the ability to create local/private Pools/Policies/Templates, that cannot be shared by peers nor parents, as described above. Another example would be the override of a

Pool/Policy/Template object that may have been created at a higher organizational level. For example, a VNIC template named “eth0-A” may have been created at the “Engineering” organization. The “US” organization could possibly create a VNIC template with the same “eth0-A” name and with properties that would override the higher-level template with attributes specific to the local organization’s needs.

While Pools could be maintained at local organizational levels, a Best Practice would be to have a single set of UUID/MAC/WWNN/WWPN Pools maintained exclusively at the “root” organization --- and created in close coordination with the data center’s site-wide catalog.

Conclusions

UCS Manager offers a plethora of powerful configuration and policy options. Many of these options can greatly assist in formalizing business rules as applied to compute/network/storage resources. As UCS-based environments scale in size, UCS Management provides many core features and functionality to facilitate ease of growth, policy-driven automation, and management simplicity.



Appendix: Ultimate UCS Quick Start

To get a UCS system up and running as quickly as possible, here is a Quick Start Guide, highlighting the minimal steps required to provide utility and to be operational.

- 1) Set hostname, IP addr, gateway, etc. from the Serial Console connection.
- 2) Set the Chassis Discovery Policy for the number of FEX->FI connections (1, 2 or 4).
- 3) Configure/Enable Server Ports; Configure/Enable Uplink Ports; Configure/Enable FC Ports.
- 4) Create Management IP Address Pool (typically same subnet as UCS Manager Admin IP)
- 5) Create "Host Firmware Policy" with packages from most recent UCS software bundle
- 6) Create UUID Pool; Create MAC Pool; Create WWNN Pool; Create WWPN Pool (or populate the corresponding "default" pools)
- 7) For SAN boot, create a unique "Boot Policy" for each storage array boot target.
- 8) Create VNIC templates ("eth0-A", "eth1-B"), that both draw from the above MAC Pool, and are associated with Fabric-A and Fabric-B respectively.
- 9) Create VHBA templates ("fc0-A", "fc1-B"), that both draw from the above WWPN Pool, and are associated with Fabric-A and Fabric-B respectively.
- 10) Create service-profile templates that draw from all earlier established pools, policies and templates, as appropriate.
- 11) Instantiate service-profile from template and associate service-profile to a given blade --- **OR** --- set service-profile template to associate with a particular Server Pool.
- 12) Configure PXE server, or map a bootable ISO image to the virtual-media CDROM drive to begin the OS installation.