Cisco IT Migration to ACI

This is the third white paper in a series of case studies that explain how Cisco IT deployed ACI to deliver improved business performance. These in-depth case studies cover the Cisco IT ACI data center design, migration to ACI, network security, the ACI NetApp storage area network deployment, virtualization with AVS, UCS, KVM, and VMware, and server load balancing. These white papers will enable field engineers and customer IT architects to assess the product, plan deployments, and exploit its application centric properties to flexibly deploy and manage robust highly scalable integrated data center and network resources.

Contributors to this white paper from the Cisco IT Migration Team include: Anitha Parimi, Principal Engineer, Ishan Mehta, Network Engineer, Benoit Mendy, Compute Engineer, Gary Coburn, Compute Engineer, Nikhil Mitra, Storage Engineer, Karlo Atienza, Network Engineer (Hadoop), Bidhu Das, DBA Architect (Hadoop), Jeff Lehman, Program Manager.

v1.1, June 15, 2020 – updated with copy edits for clarity.

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

Table of Contents

Cisco IT Migration to ACI

Cisco IT's objective is to migrate to interconnected Application Centric Infrastructure (ACI) fabric datacenters that integrate datacenter operations more tightly. This unity of datacenter operations is at the heart of the value Cisco IT expects to derive from ACI. Cisco is extending ACI to its datacenters worldwide. These deployments are seamless and efficient to complete because ACI employs a common policy model in each datacenter.

Cisco IT sees the common policy model of ACI as the key to a fast IT environment.

Cisco IT ACI Fast IT

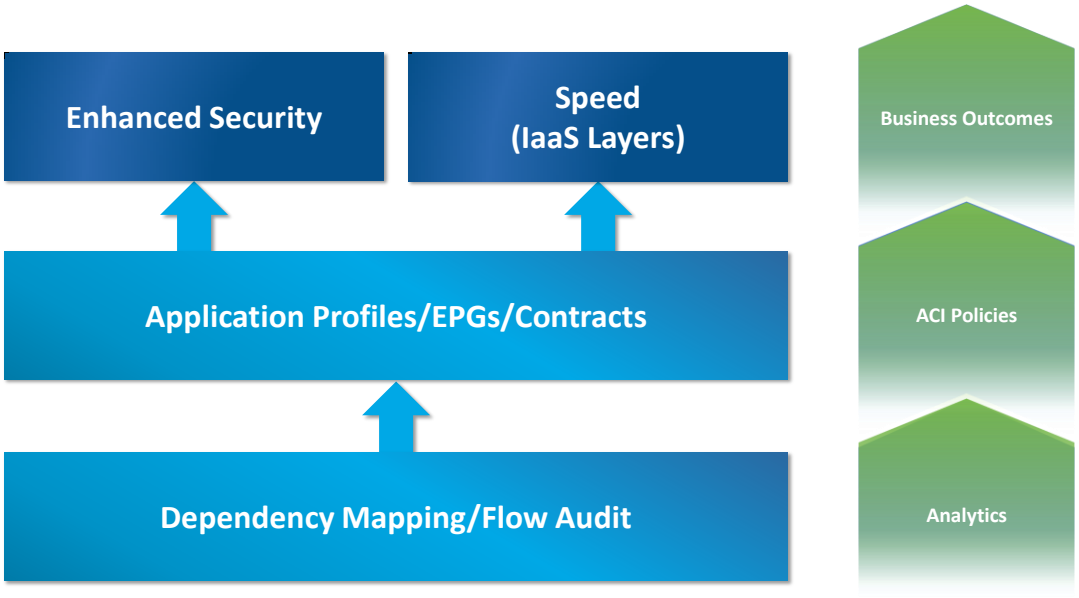

The Cisco IT vision of a fast IT environment enables the following:

- Speed and Agility: Continuous delivery of adaptive infrastructure enables fast IT.
- Enhanced User Experience: Reduced downtimes for maintenance, resiliency for unplanned outages, and automation for self-service provisioning of IT resources.
- Adaptive Security: The ACI policy model is extensible across the Cisco IT hybrid clouds while enabling scalable data centers with a higher level of network security, and simplified management.

Cisco IT projects that the ACI fast IT environment will yield dramatic improvements in business performance.

ACI Fast IT Business Performance Improvements

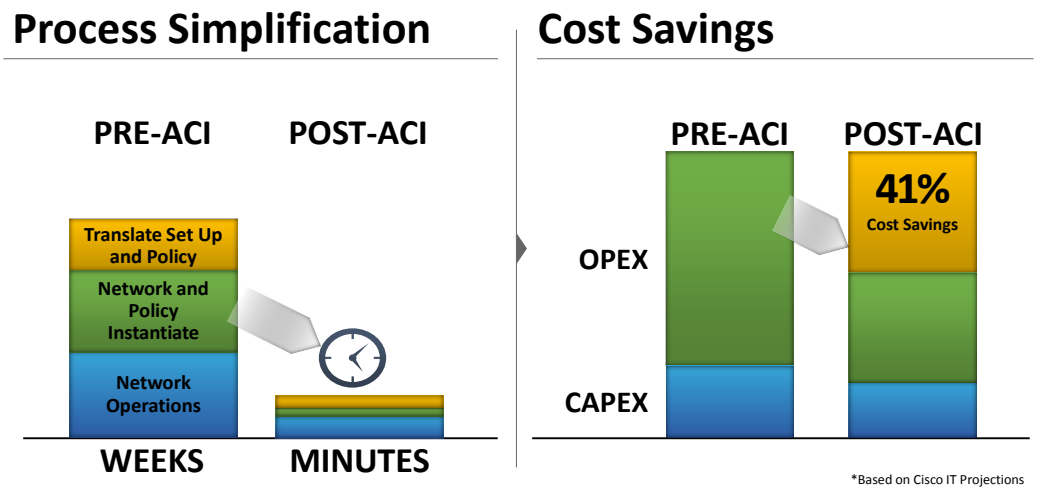

The journey to migrate the Cisco IT traditional data center to the ACI fabric starts with understanding the known foundations of network, compute, and storage infrastructure that support business applications. The challenge of understanding what exactly the business applications are consuming from the data center infrastructure foundations is the beginning and end of this migration journey.

As Cisco IT Principal Engineer Anitha Parimi says, "While the journey is still progressing, the large-scale production data center deployments to date have transformed our business performance. We gracefully migrated Hadoop to ACI and performed large scale in-place upgrades with minimal down time." This white paper will show exactly how the ongoing Cisco IT migration to ACI is a transformative journey.

Cisco Data Center Scale

The scale of the Cisco IT data center deployments presents both migration opportunities and challenges. The Cisco IT organization operates multiple business application and engineering development data centers distributed around the world.

Cisco IT Worldwide Data Centers

Cisco IT supports 141,000 employees (71,000 regular employees and 70,000 contractors) in 583 offices across more than 100 countries. The data centers occupy more than 269,000 sq. ft. of floor space and draw 30.1 MW of UPS power. More than 11,000 Cisco Unified Computing System™ (Cisco UCS®) blades are deployed with 92% of the servers virtualized in new data centers.

The infrastructure for the core business data centers (DC) is big. For example, the Allen, Texas DC includes 856 network devices that support 2300 traditional and private-cloud applications, run 8000 virtual machines, include 1700 Cisco Unified Computing System™ (Cisco UCS®) blades and 710 bare metal servers, with 14.5PB of NAS storage and 12PB of SAN storage. Cisco is driven to migrate to ACI because, as its data centers grow, quick and agile application deployment becomes increasingly challenging.

Migration Process

Cisco IT's goals for migrating to a new data center architecture are to enable provisioning infrastructure resources without any delay and to enable secure applications running in the enterprise to scale gracefully.

High Level Migration Process

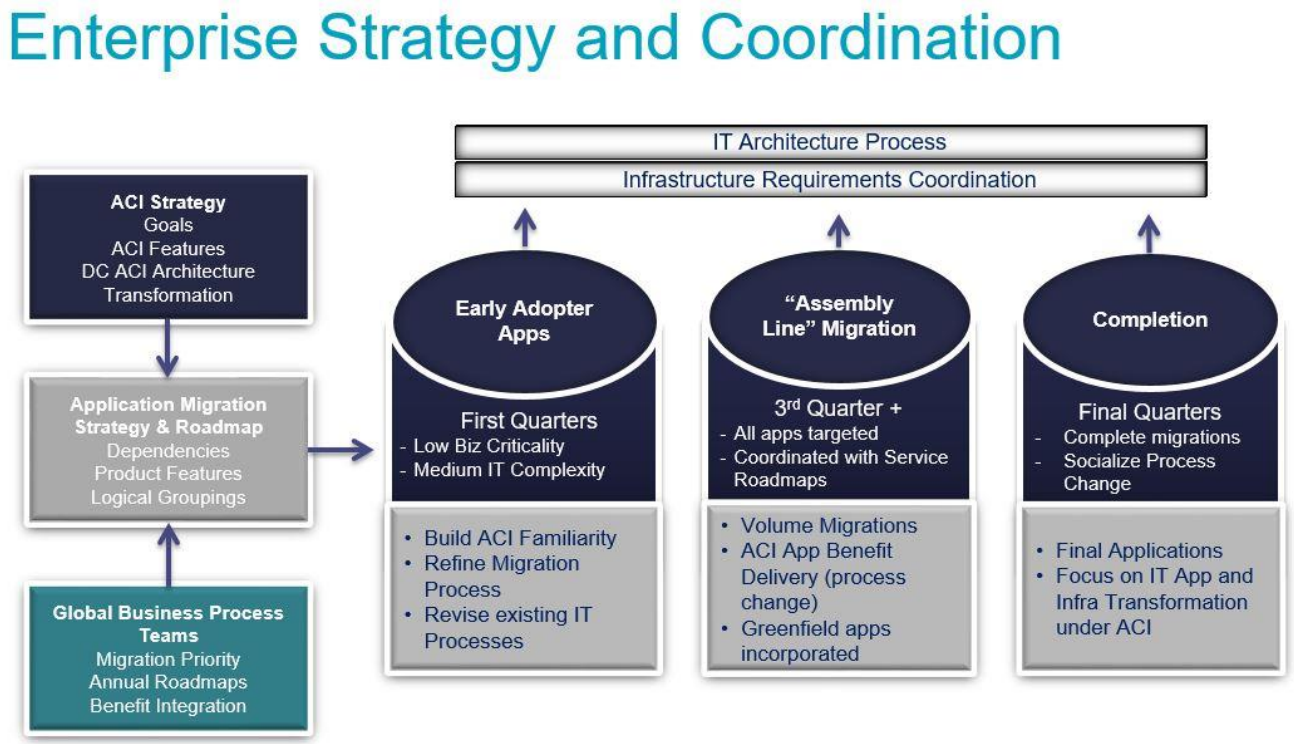

The high-level migration process Cisco IT adopted includes familiar planning cycles, as well as design and migration tasks that enable capturing the unique value of application profiles in the ACI policy driven model.

Migration Roadmap

Phase one of the multi-year Cisco IT migration roadmap started with running the Nexus 9000 series switches in standalone mode, then changing over to run those switches in ACI mode.

Cisco IT ACI Migration Roadmap

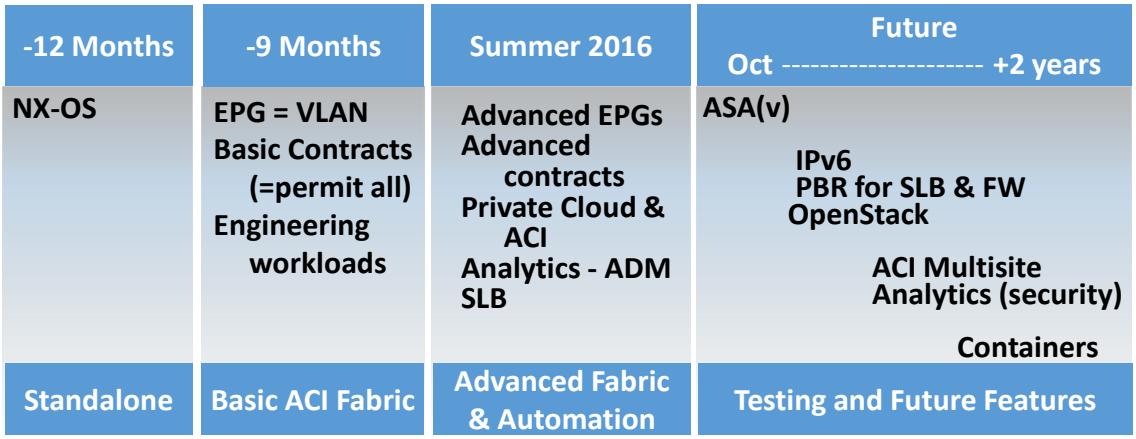

This roadmap enabled the following key outcomes:

- Seamlessly integrate applications across the existing network and the ACI fabric.
- Use a policy model to define applications.
- Lower network provisioning time by 40% by automating the provisioning steps.
- The complete real time application dependency model helps reduce incidents.

Basic ACI Fabric with Allow-all EPG Communications

The endpoint group (EPG) is the most important object in the policy model. Endpoints are devices that are connected to the network directly (like servers) or indirectly (like web clients). Knowing that 80% of the workloads in its traditional data centers don't have ACLs, Cisco IT decided to bring up ACI using "allow-all" endpoint group to endpoint group (EPG to EPG) communications in the fabric, while using their existing security infrastructure outside the fabric.

The Cisco IT standard ACI DC fabric has four spine switches, one pair of border leaf switches for external connectivity, two or more pairs of leaf switches for end point connectivity, and the minimum supported number of three APIC controllers. The scale out

capacity is 288 leaf switches with up to 12 40GB links between each spine and leaf switch. For details about the Cisco IT standard ACI design, see the Cisco IT ACI Design white paper.

Cisco IT Standard ACI Data Center Fabric Today

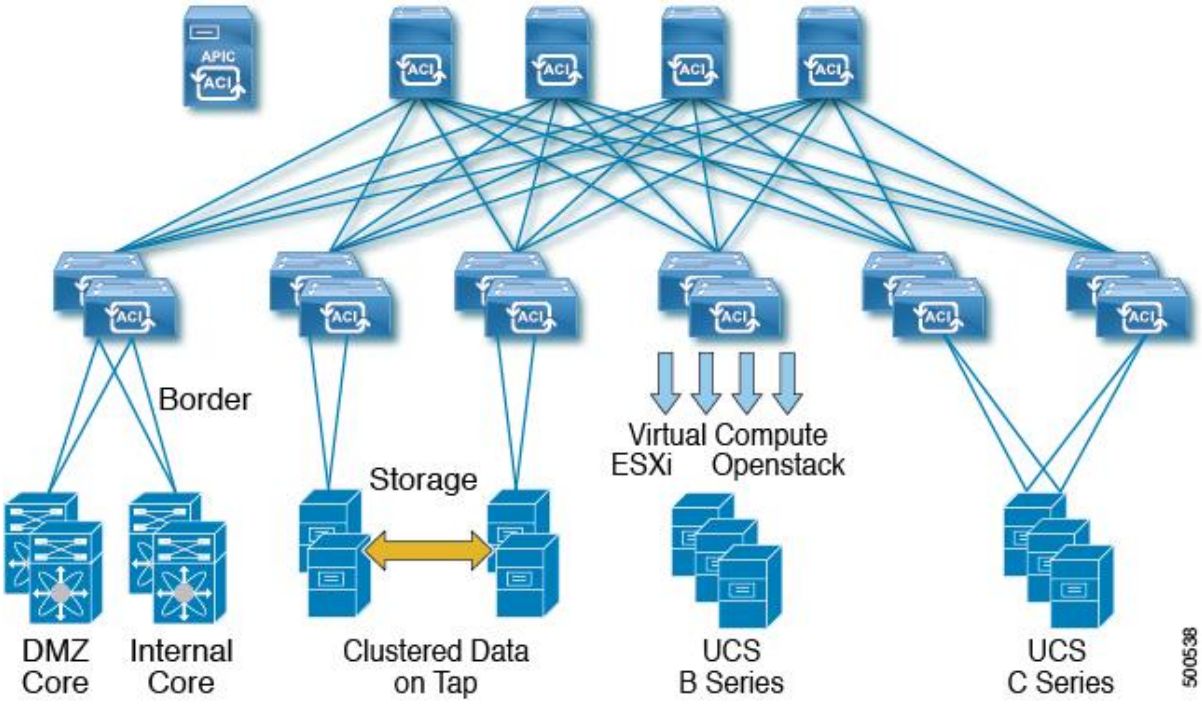

Cisco IT uses a standard ACI fabric topology in production data centers such as those in Research Triangle Park, North Carolina, Richardson, Texas, and Allen, Texas.

Cisco IT is evaluating the ACI v2.0 multipod feature for potential future deployment. Multipod uses a single APIC cluster for all the pods; all the pods act as a single fabric.

ACI Multipod – Evaluation for Future Enhancement

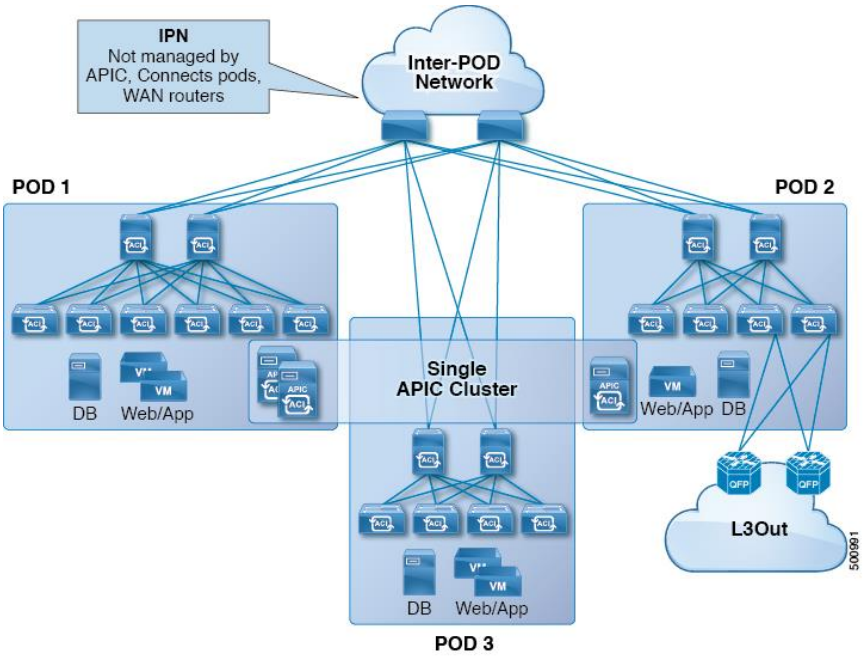

Multipod enables provisioning multiple pods per floor, building, or region that provide Layer 3 connectivity between pods.

Advanced ACI Fabric with Allowed List Contract EPG Communications

Now, Cisco IT is proceeding with advanced ACI fabric and automation deployments. Starting in the summer of 2016, Cisco IT began to use Tetration Analytics for performing application dependency mapping, deploying advanced ACI allowed list security policies, ACI based server load balancing, and hybrid public and private cloud workloads.

Migration Steps

The migration steps include one-time tasks to bring up the ACI data center, and iterations of tasks to migrate applications.

Migration Steps

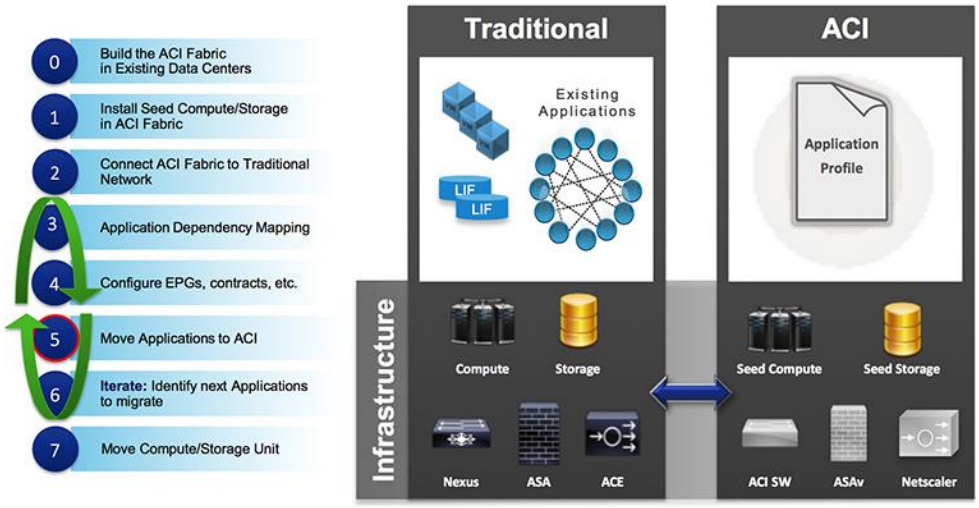

The Cisco IT migration to the ACI started with a build-out of the ACI fabric in parallel with the traditional network in the datacenter.

Build-out in Parallel of ACI Fabric with Existing Data Center Resources

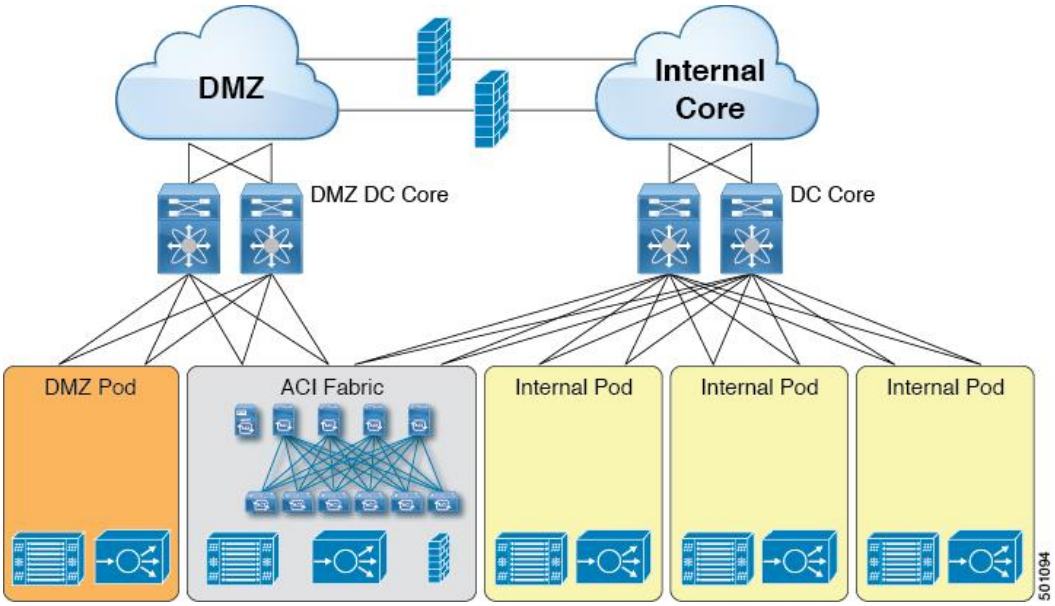

The initial basic ACI deployment enabled the ACI fabric to coexist with the existing data center infrastructure, while keeping the firewall functions provisioned outside of the ACI fabric. In subsequent phases, Cisco IT used ACI policy-based routing policies to dynamically deploy L4-L7 service graphs that place multiple virtual ASA firewall appliances inside the ACI fabric. This solution provides simple automation that enables smaller firewalls that can be deployed per application.

One-time Task - Provision Multi-Tennant Infrastructure for the Organization

Cisco IT provisions core services as ACI tenants that support business operations services. Cisco IT designed ACI tenants to align with their IT services framework.

Cisco IT Mapped Multi-Tenant Infrastructure for the Enterprise Organization

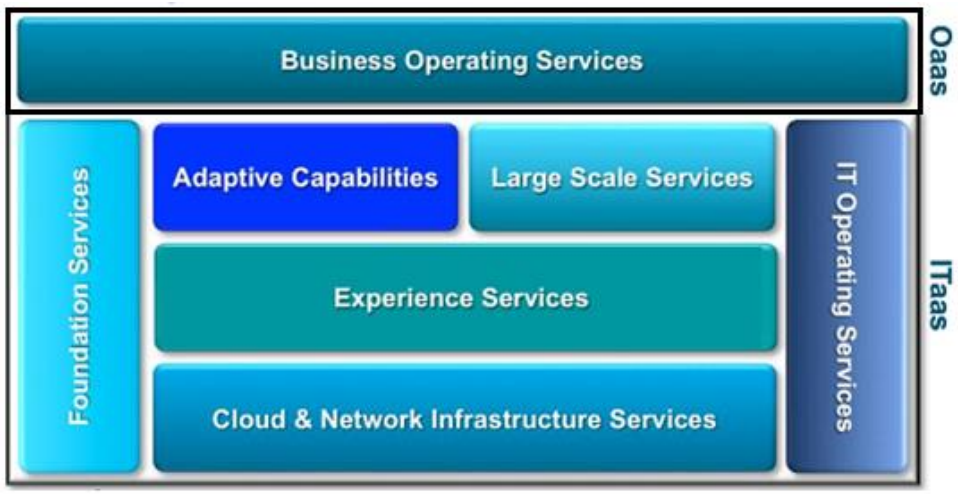

Enterprise services are grouped into these tenants:

- Foundational Services: access management, application design and development, IT quality, IT support and operations, enterprise change management, product and software development and management
- Adaptive Capabilities: Cisco commerce, software and subscription, channel partner program and sales, coverage credit and compensation
- Large Scale Services: order to cash, xRM and customer care, core finance, HR, and WPR
- Experience Services: knowledge and content management, collaboration, employee productivity and devices
- Cloud & Network Infrastructure Services: database and middleware services, infrastructure services, network services
- IT Operating Services: enterprise architecture, IT and business relationship management, IT business services

In addition, during the early phase of implementing ACI, Cisco IT built automation first to enable business application developers to quickly self-service deploy proof of concept applications within their own tenants.

Cisco IT ACI Tenant Design

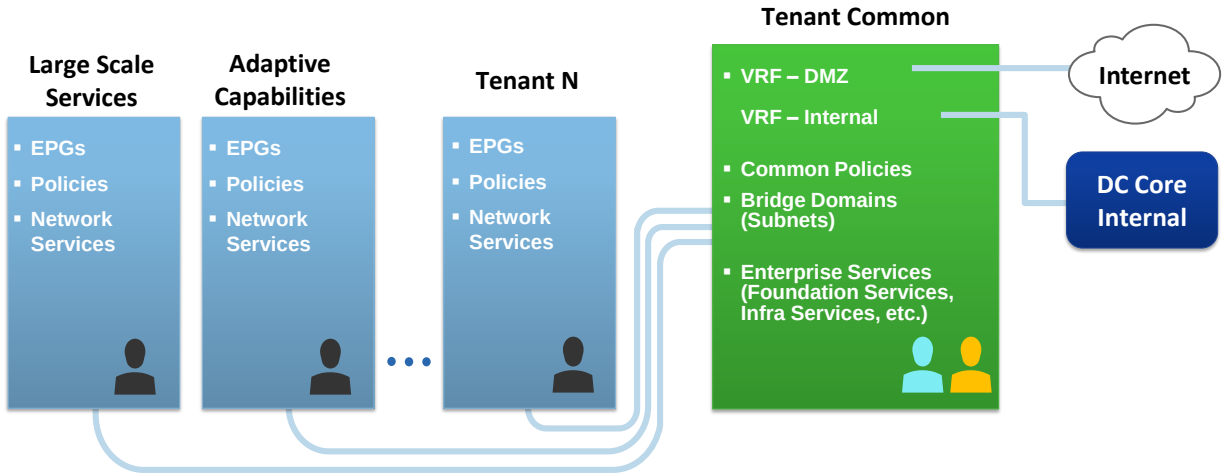

Cisco IT enables application developers to select VM/memory/CPU resources that are isolated within a tenant. If additional network access is needed (such as Web servers, network protocols/ports), the network services team handles such requests.

One-time Task - Provision Shared Services Contracts

Contracts specify what is allowed in the ACI fabric. They can be reused whenever a new service or EPG needs access to standard data center functions. Compared with ACLs, contracts are much simpler to manage, and remove the need for most ACLs.

Prior to ACI, 80% of Cisco IT's ACL entries were set up enabling communication to shared infrastructure services and shared application middleware services. Cisco IT opted to present these as contracts in the ACI tenant common. They are easily consumed by any EPG in ACI. Doing so eliminates the significant time spent previously managing ACLs.

Cisco IT Shared Services Contracts Architecture

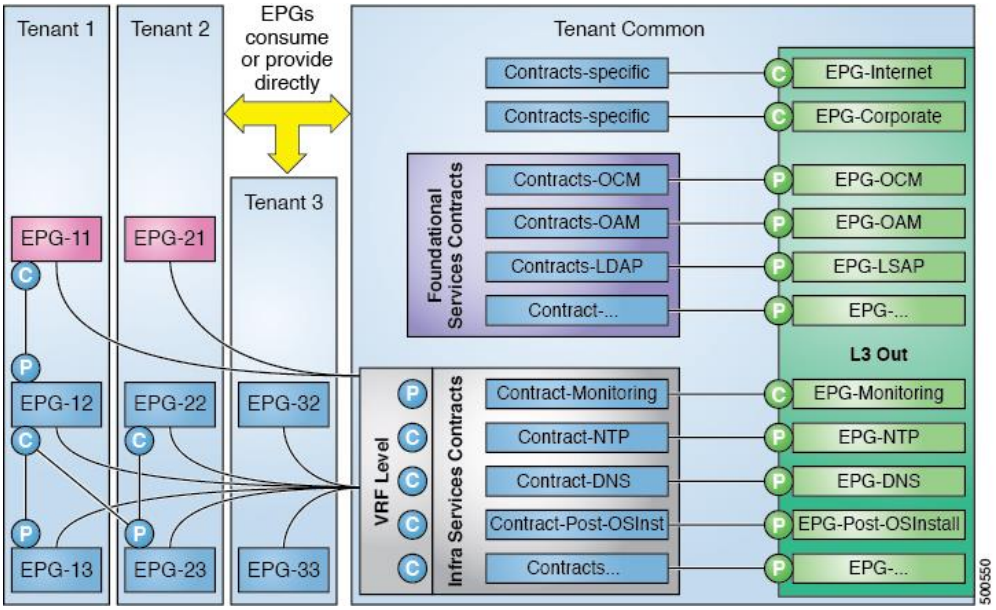

Clients can select shared foundational services (middleware) contracts as needed but they can't configure them. Foundation services include the following:

- Active Directory
- LDAP
- Authentication Systems
- Oracle Connection Manager
- Messaging middleware
- Web Security Gateway

Shared infrastructure services contracts are applied to all EPGs. Clients can see them but can't configure them. Shared infrastructure services include the following:

- DNS
- NTP
- Security and Performance Monitoring Systems
- Syslog
- Puppet

Standard network management protocols such as ICMP or SSH and are set up once in individual contracts that are reused across the fabric. For example, a single vzAny contract specifies that ICMP is allowed.

Contract Example: One ICMP Contract for Many L3 Out Connections

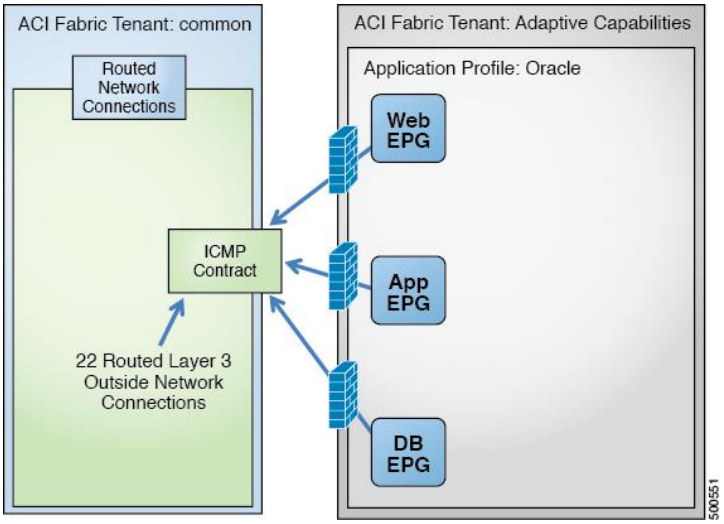

From then on, that single vzAny contract is reused for all routed connections automatically by virtue of the host being a member of a logical group that must comply with the rules of the contract. By dynamically applying contract rules to all EPGs in a context (VRF), vzAny automates the process of configuring EPG contract relationships. Whenever a new EPG is added to a context (VRF), vzAny contract rules automatically apply. The vzAny one-to-all EPG relationship is the most efficient way of enabling the ACI leaf switches to apply contract rules to all EPGs in a context (VRF).

One-time Task: Install Seed Compute and Storage in the ACI Fabric

Cisco IT started with building seed compute on the ACI fabric. Then, depending on the application, some applications in dedicated subnets were moved over a Layer 2 connection from the enterprise core. Other applications were moved with a change of the application IP address. The new fabric has an ASA firewall and a NetScaler load balancer.

IP NAS storage moved to the fabric next. With compute and storage on the fabric, the required components were available for migrating applications. Load balancer configurations that applications require migrate to the fabric connected load balancer. The process of dependency mapping, defining the application profile EPGs, and moving the respective compute, storage, load balancing and firewall configurations iterates as each application migrates to ACI. As compute and storage from the traditional network free up, they are repurposed to the ACI fabric for subsequent application migrations.

One-time Task – Integrate ACI with Traditional Networking to Enable Migration

In an existing data center, Cisco IT used the following steps to bring up the ACI fabric with connectivity to the existing traditional network.

Cisco IT Migration Physical Networking Topology

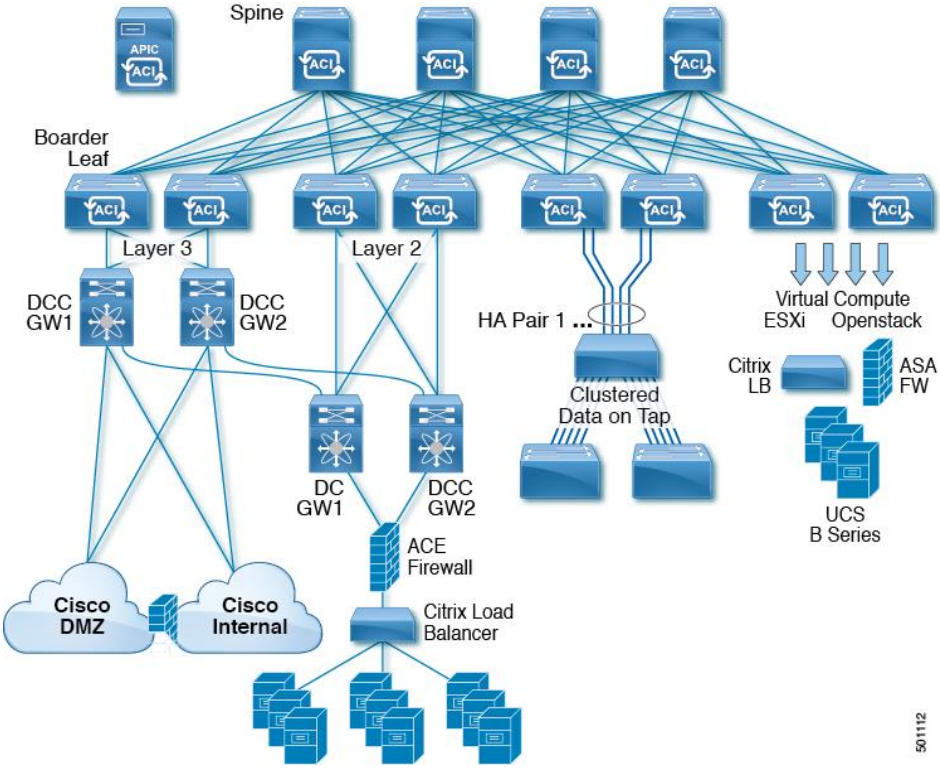

Cisco IT performed application network migration according to one of these two scenarios:

- Migration retaining the IP – L2 Extension. The layer 2 connection between the existing data center resources to the ACI fabric enabled moving the workloads of subnets. For large workloads, while portions of the workloads were moved, connectivity between data center core gateways (DCCGW) and the data center gateways (DCG) enabled the portions split between the ACI fabric and the existing data center to communicate across the split.
- Migration change in IP address

Moves of application VMs to ACI (detailed below) use the same process as migrating VMs across data centers.

The steps for performing a migration retaining the IP address over an L2 extension segment are as follows:

1. L2 physical link extension between N7k and ACI border leaf. Double sided VPC is configured.
2. Identify the VLAN that needs to be migrated. Trunk the VLANs over the L2 extended link.
3. Prebuild Tenant/AP/EPG/Contracts for applications.
4. Cisco IT uses Hot Standby Router Protocol (HSRP) as a redundancy protocol. Virtual IP (VIP) is the default gateway for all the hosts within the VLAN.
5. During the cutover, the HSRP – VIP is removed from the N7k and configured as a bridge domain IP in the ACI fabric.
6. In addition to step 5, EIGRP routing for that SVI is removed from N7k and SVI is shut. The route is advertised from the ACI fabric.
7. Steps 5-6 cause network unavailability for up to 60 sec. Verify the cutover IP, availability of all the hosts within that subnet and reachability via the ACI route.
8. Partner with compute and begin the VM migration from traditional to ACI infrastructure.
9. Clients to verify

The steps for performing a migration with an IP address change are as follows:

1. Identify the hosts in GP subnet which needs to be migrated.
2. Run the scripts that check ACL dependency and SLB configuration setup.
3. Add the ACL lines with new IP address and SLB configuration with new IP.
4. Prebuild the Tenant/AP/BD/Subnet/EPG/Contracts.
5. Work with compute on migrating the hosts and re-IP.
6. Clients to verify.

The Cisco IT Citrix NetScaler server load balancer standard configuration in the ACI tenant common includes the following:

- NetScaler has a dedicated management subnet. An ACI bridge domain (BD) is configured under tenant common. The bridge domain name “slb-mgmt-01-bd”, uses a different MAC address then the ACI pervasive gateway.

- The server load balancer management application profile EPGs are configured in the ACI tenant common.

- Each NetScaler has 10 VNIC cards and always maps Network – Adapter 1 (1/0 on the NetScaler) to the management EPG.
- As a best practice, interfaces that are not in use on NetScalers are shut down. They are enabled only when they are to be mapped to the load balanced EPG.

Dependency Mapping for Each Application Migration

Working with the Cisco application developer owners, the security team, and the networking teams, the migration team assembled the 'tribal knowledge' into a best effort definition of application dependencies. This information enabled the initial phase of migrating applications to the basic ACI fabric.

Untangling Application Dependency

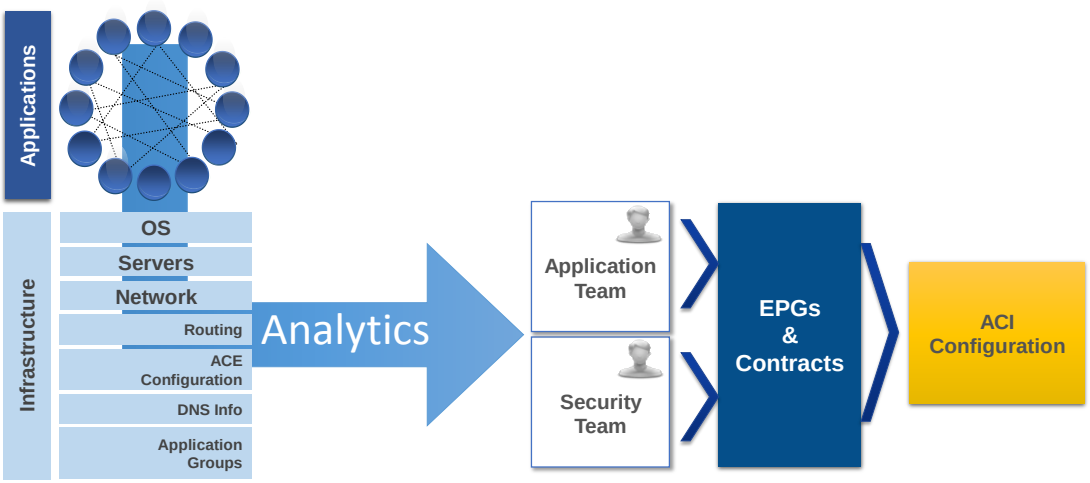

While this information was significant, the applications Cisco IT migrates during the ongoing advanced fabric phase require a more thorough process to assure that there are no gaps caused by insufficient visibility into the datacenter environment.

For the advanced fabric phase of its migration to ACI, Cisco IT is using Tetration Analytics to identify exactly how applications consume data center resources. Tetration derives deep telemetry from lightweight software sensors that run on servers and built-in hardware sensors in the Nexus 9K platform. It delivers real-time analytics to achieve actionable insights by searching billions of records in seconds. Tetration is capable of

processing millions of flows per second with the capacity to retain and replay billions of flow records without aggregation. This enables validating the information the various stakeholders provide, identifying gaps in that information, and automatically grouping the application dependent system components into logical units that map into ACI application profile EPGs.

ACI Plus Tetration Analytics Tightens Security

Security is a priority for Cisco, even as the explosion of east-west traffic in recent years has greatly expanded the attack surface. To bolster security, Cisco IT chose to implement the ACI "zero-trust operations" that include the use of allowed list policies. This operating model changes traditional network default communication permissions between applications from "permit any" to the ACI default of "permit none" unless explicitly allowed. This prevents attacks from propagating across applications, tenants, and data. Compliance can be validated quickly by comparing actual traffic flows with ACI allowed list policies.

Cisco IT phased in the full implementation of the allowed list security model. Cisco IT first moved applications to a basic ACI fabric deployment with ‘allow-all’ contracts because manual analysis of application flows was difficult and the because of the risk of missing flows. All the existing security infrastructure outside the ACI fabric still applied to these phase 1 basic ACI fabric application flows. In the phase 1 basic ACI fabric portion of the migration roadmap, applications moved to ACI still benefit from the zero-trust environment due to the isolation ACI tenants, application profiles, and EPGs provide. Even in the “allow-all” mode of the phase 1 basic ACI fabric, communication cannot jump from tenant to tenant, from application profile to application profile, or from endpoint group to endpoint group without explicit permission granted in ACI contracts.

Starting in early 2016, Cisco IT began using Tetration to migrate applications to its ACI zero-trust security environment.

Tetration Analytics

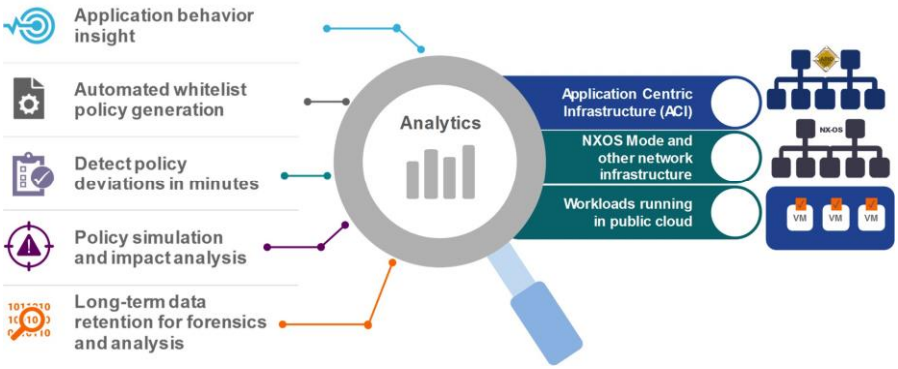

Cisco IT is now migrating applications using contracts to allow only what the applications need. According to an IDC white paper, Tetration and ACI support the following efficiencies:

- Defining and creating security policies for the allowed list security model: With Tetration, Cisco's IT staff spend far less time defining, creating, and validating packet flows before applying a security policy. Cisco expects this saves 2,600 hours of staff time per 100 applications (81% less staff time).
- Application grouping: Tetration automatically groups similar servers. Cisco expects this results in 150 hours of staff time savings per 100 applications (75% less staff time) in determining dependencies.
- Device configuration with orchestration: ACI implements allowed list security models without needing change requests and approvals to be generated. Cisco anticipates this avoids 900 hours of staff time per 100 applications (75% less staff time).

In total, these efficiencies with Tetration and ACI enable Cisco IT to save 3,650 hours of staff time per 100 applications as it maps application dependencies for thousands of applications, while enabling use of the ACI allowed list security model.

Migration Steps for Each Application Migration

The ACI application migration process starts with considering how an application architecture is overlaid on these ACI features: the application Profile (AP) which contains

endpoint group (EPGs), and contracts. For example, an EPG can contain hosts that have similar application functions and security requirements. Typical application EPGs correspond to web, application and database hosts. More complex applications could have multiple application EPGs. Contracts define the allowed network connections between the application EPGs or between an application EPG and an external network.

A central activity in the migration process is mapping the physical application architecture to EPGs and defining the contract requirements between EPGs.

Cisco IT ACI Application Migration Process Flow

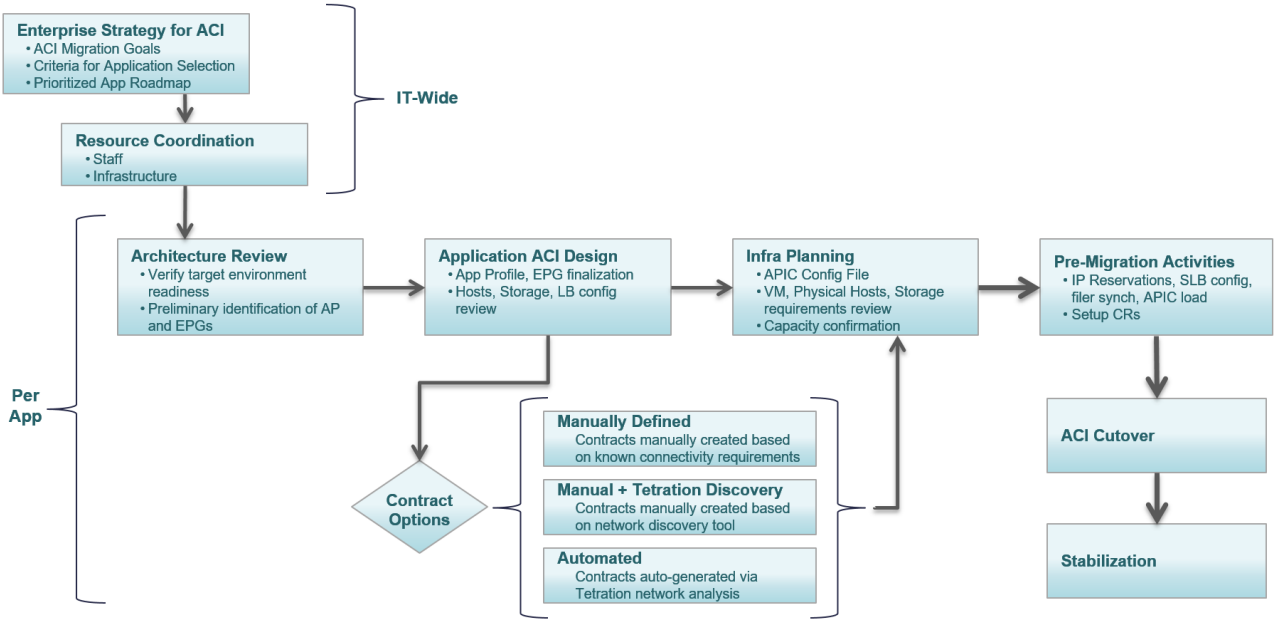

The application migration process starts with an architecture review that includes the following steps:

- Validate that the application is supported by available ACI capabilities
- Specify the deliverables
  - Architecture Diagrams for all lifecycles
  - Dataflow diagrams
  - Host and host configuration lists, storage requirements, load balancers
  - Identification of required ACI features and EPGs
  - Processes
- Engage the application team

    - Collects architecture and host details
    - Review meeting (ACI and app architects and PMs)
- Challenges
    - Key architect availability
    - Incompletely documented or changing architectures
    - Firewalls and ACLs

A key task in the application process is specifying the contracts that govern the communications between EPGs.

- Manual specification: suitable for applications of low complexity and migrations taking place early in an organization's learning cycle.
    - Manual definition of contracts based on documented application network connection requirements.
    - Risk mitigation: allow a suitable period for contract updates to incorporate requirements not discovered pre-migration.
- Manual with network connection discovery: suitable for applications of low to medium complexity. Manual contract definition for high complexity apps is possible but could require significant time for review and validation.
    - Manual definition of contracts based primarily on network diagnostic tool reporting of connections made with application hosts. Care must be taken to ensure diagnostic data is collected over periods of representative network traffic (quarter end, marketing campaigns, etc.)
    - Tetration Analytics used for connection discovery during Cisco migrations.
- Automated
    - Automated generation of Contracts via Tetration Analytics. Supporting data must be collected over periods of representative network traffic.
    - Suitable for applications of all complexity
    - Contracts should be reviewed by Network and Application teams.

Cisco IT ACI Migration Responsibility, Accountability, Consulting, Informed Matrix

Activity	ACI PMO PM	Infra Implementation Engineers	Network Functional Design SME	App PM	App Architect /SME
App Dependencies Validation & Update	Informed	Consulted	Consulted	Accountable	Responsible
Hosts / Storage / DB Documentation	Informed	Consulted	Informed	Accountable	Responsible
Infra Planning & Resource Alignment	Accountable	Consulted	Consulted	Consulted	Informed
App Profile / EPG / Contracts Design	Accountable	Consulted	Responsible	Informed	Consulted
ACI Design Approval	Accountable	Informed	Responsible	Informed	Consulted
Cutover Execution	Accountable	Responsible	Consulted	Responsible	Consulted
Application Sanity Testing	Informed	Informed	Informed	Accountable	Consulted
Migration Signoff	Accountable	Informed	Informed	Responsible	Consulted

The Cisco IT ACI migration project responsibility matrix maps the key high-level tasks with the corresponding roles and responsibilities.

The application cutover process included the following steps.

Cisco IT Application Cutover Overview

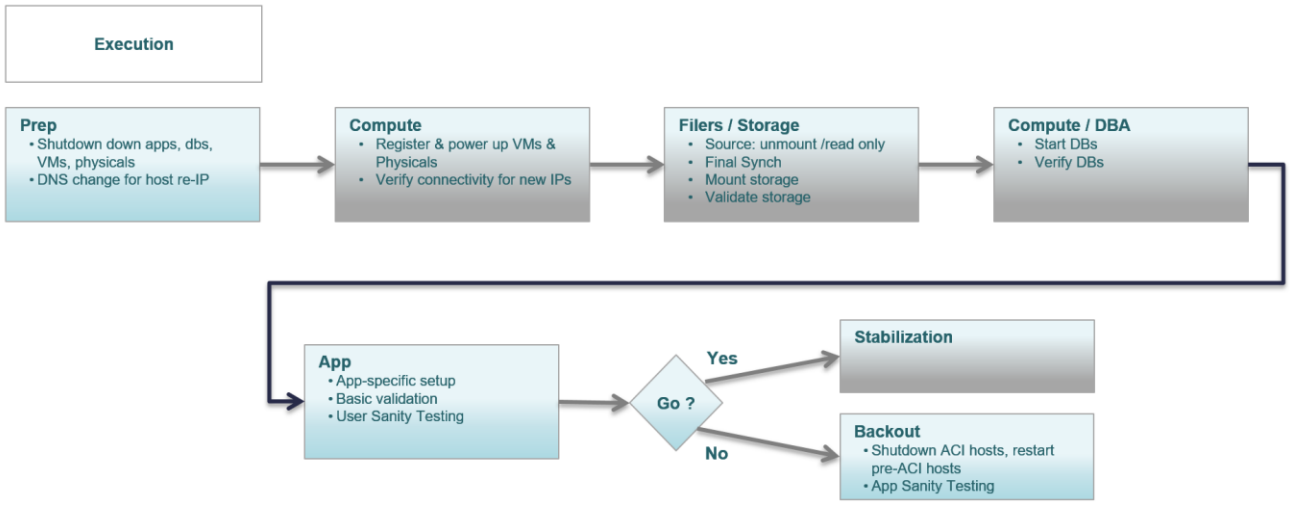

To provision network connectivity for applications migrating to ACI, Cisco IT first creates the ACI tenant (if it doesn't already exist), application profile, bridge domain and associated EPGs.

ACI Tenant Application Profile

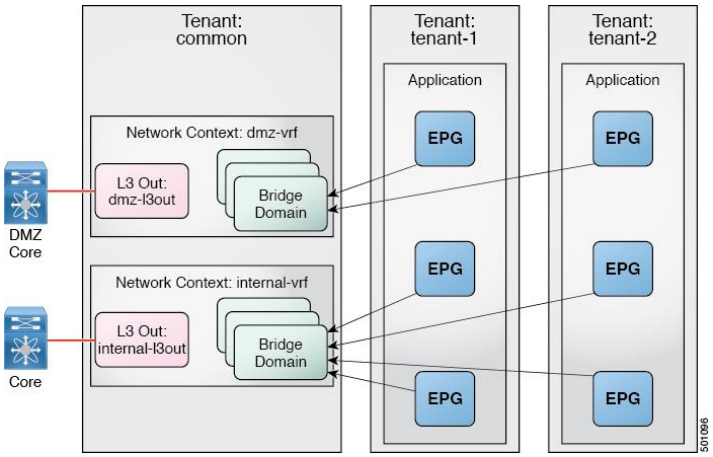

1. Create tenant, application profile, bridge domain, and EPG.
2. Add VLAN to vPC/PC interface.

SLB Migration with Hot Standby Router Protocol Swing

The server load balancer (SLB) high level migration steps include the following:

1. Identify current Cisco Application Control Engine (ACE) and Global Site Selector (GSS) configuration specifications, including VIPs, server farms, sticky method, L7 rules, SSL requirements, etc.

2. Using these specifications, prepare configuration templates in Citrix NetScaler OS format for the early phase basic ACI fabric deployment. For later phase deployments that implement ACI policy-based routing service graphs for automatic L4-L7 insertion, prepare configuration templates in XML for uploading to ACI.

SLB Prebuild

1. Create the ACI bridge domain without IP added.
2. Configure two contracts:
    a. One for all the clients to access the VIP, which will be provided by the VIP EPG and consumed by all the client EPGs & extnets.
    b. The second contract is between the VIP EPG and the Real Server EPG for monitoring.
3. Configure Application Profile (or, can use existing application profile).

a. Can use existing application profile in tenant common.
b. Configure two EPGs:
    - One for the VIPS on the ACE
    - The other for the VIPs on the NetScaler. VIPS on the ACE will be L2 extended VIA the double sided vPC between the border leaf and the POD GW
c. For mtiaas-mig-vip-egp, set up static binding with the vPC link and for mtiaas-int-vip-epg, associate it to the VMM domain so that the EPG name populates in vCenter via the AVS.
d. Both EPGs provide the contract "slb-clt-vip-contract". The mtiaas-int-vip-epg consumes the "slb-vip-real-contract" contract. This enables monitoring between VIP EPGs and REALS EPGs. If the VIP is on ACE and the servers are on ACI, the contract is consumed by mtiaas-mig–vip-epg as well.

SLB Cutover

1. Trunk VLAN [##] on the L2 extension.
2. Remove Hot Standby Router Protocol (HSRP) from the standby GW.
3. Remove HSRP/EIGRP from GW1, configure HSRP on a bridge domain and advertise the route over the DC core.

Citrix NetScaler Migration Steps

1. Re-check the ACE SSL and GW usage.
2. Create the summary sheet with a timeline and the corresponding VRFs, VIPs, and contacts.
3. Start pre-configuration for the NetScalers.
    a. Caution – to prevent an outage, be sure all the NetScaler VNICs are in a quarantine VLAN and only 0/1 is in the management port-profile.
    b. Disable the VIPs and service groups that are pre-configured for migration.
    c. Observe naming standards
    d. Configuration high level steps:
        - Add VLANs
        - Add Subnet IPs (SNIP)
        - Bind VLANs with subnet IPs (SNIP) and Interface
        - NAT configuration
        - ADD RSERVERs

- ADD Service groups
- ADD Monitors
- ADD LB Servers
- BIND LB Server with Service group
- SSL OFFLOAD Example
    1. Place certs + key in folder
    2. Add cert and key
    3. Create vserver as type SSL
    4. Bind cert to vserver

1. Trunk all the VLANs to the UCS cluster and create a trunk port-profile that allows all the VLANs configured on ACE.
2. Follow these steps during the Change Request (CR):
    a. Remove the policy from the interface on the ACE or remove VLAN from the SVCLC group if migrating the whole context at once.
    b. Associate the ACI port-profile with the NetScaler interface 1/1.
    c. Enable the VIPs and service groups.
    d. Clients to test.

Storage Migration for Each Application Migration

The high-level phases of automated storage migration include the following:

- Iterate small scale dry runs.
- Develop automation.
- Use automation to move high volumes of storage to ACI.

Cisco IT ACI Production NetApp cDOT Storage Cluster

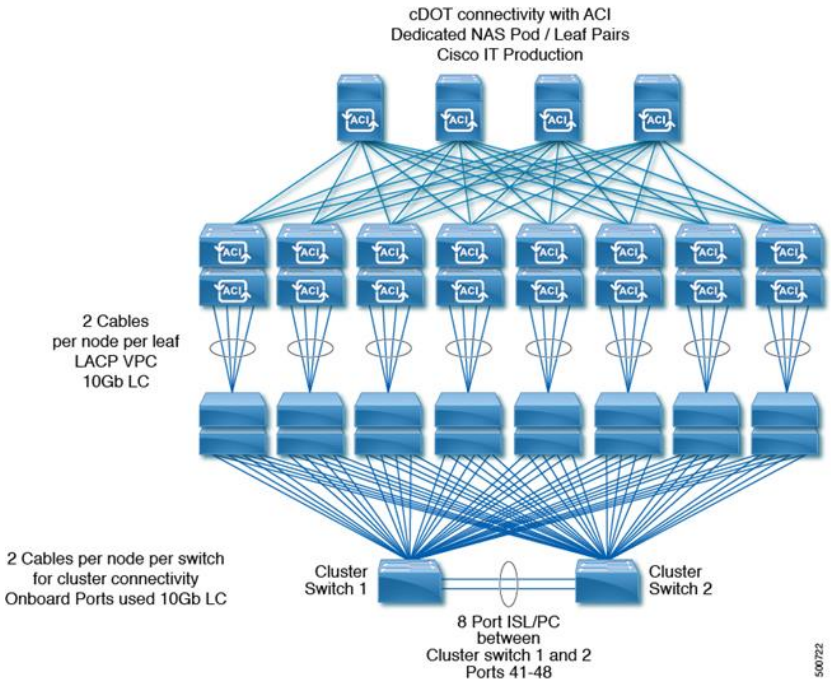

Seed New ACI based C-DoT cluster

1. Deploy new cDOT systems.
2. Transfer ESX NAS datastores using storage vmotion.
3. Migrate VM OS host NFS shares from 7-mode to CDOT systems using these migration tools:
   - One time / one-way Async Snapmirror from 7mode node to cDOT node
   - NDMPCOPY (file based if above cannot be used in a given case)
   - XCP from NetApp
4. Convert vacated existing 7-mode MTIaaS nodes to cDOT and join to cDOT cluster
5. Phased rollout – migrate/convert 7mode systems that still have 18-24 lease time left

ESX Server NFS Datastores for VMDKs

This is done using svmotion. Storage IMPs only need to set up and present MTIaaS cDOT /ACI NFS datastores. VIF hosting teams can then execute the svmotion.

Host OS NFS Shares - For Managed vs unmanaged VMs

- Managed: Use script to change /etc/fstab.
- Unmanaged: Send email to clients on what to change.

Migrate Host OS NFS Shares

1. Identify NFS host mounting vfiler:shares.
2. CITEIS team sends storage information for analysis to Storage IMP team
3. Multiple hosts mounting same share need to be included in the same CR/outage.
4. Once the source and destination 7mode and cDOT nodes are determined, complete these tasks:
   - Set up initial seed async SMs (at least a week before if possible, maybe longer for large source shares).
   - Schedule a SM update using the cron function in CDOT.
5. At the time of the outage, complete these tasks:
   - Unmount all VM HOST SHARES,
   - Edit host '/etc/fstab' file.
   - Shut down the client VM.
   - Do final SM src/dst sync.
   - Restart the clients.
   - Verify new NAS share access.

Block Storage / SAN Behind ACI

With ACI, Cisco IT did not see any need to change its SAN design.

Cisco IT ACI FCoE / SAN Storage

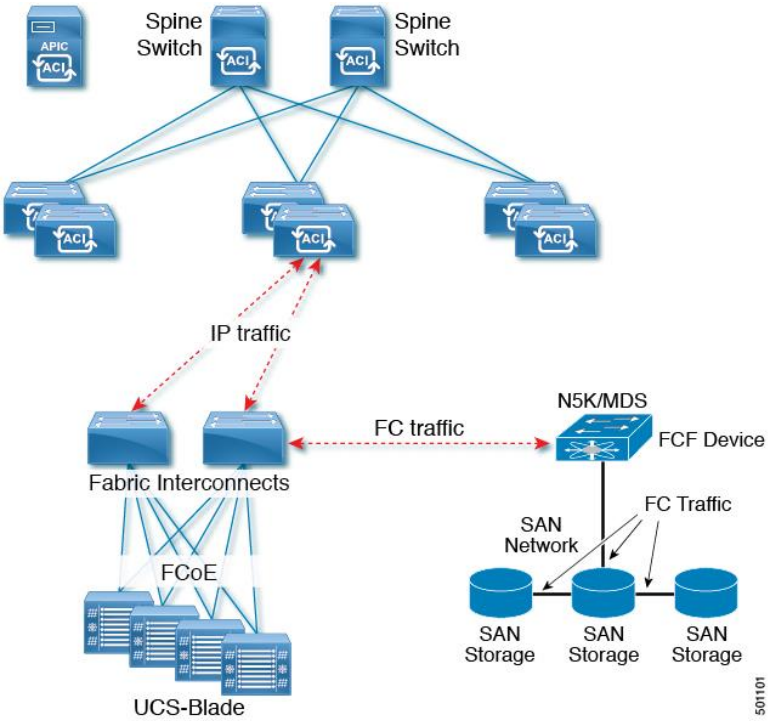

Cisco IT uses UCS Blade Servers with FCoE between UCS B-series servers and the UCS Fabric Interconnect. The UCS FI splits between IP traffic and Fiber Channel traffic. Cisco IT connects its UCS FI to the existing MDS infrastructure and makes the relevant paths available to the Block Storage Frames prior to the ACI migration.

VM Migration

ACI fabric virtual machine manager (VMM) domains enable an administrator to configure connectivity policies for virtual machine controllers.

ACI Data Center-Wide Mobility Domain

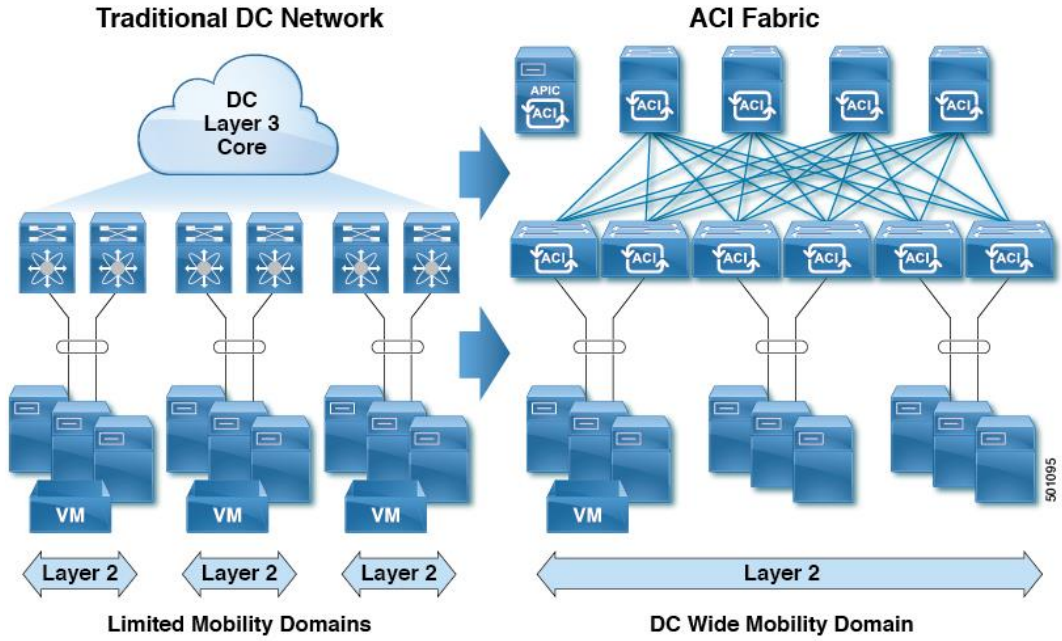

Unlike the limited mobility domains of traditional data center networks, ACI enables data center-wide mobility domains.

More Efficient Resource Pool Consumption

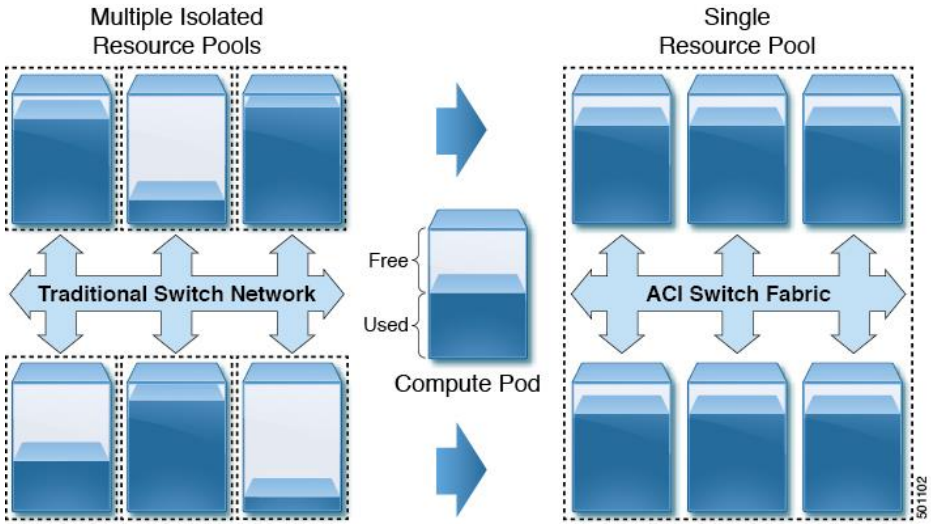

The ACI single data center-wide mobility domain enables more efficient consumption of compute resource pools.

In ACI, multiple VMM domains can coexist and interoperate.

ACI Multiple VM Controller Integration

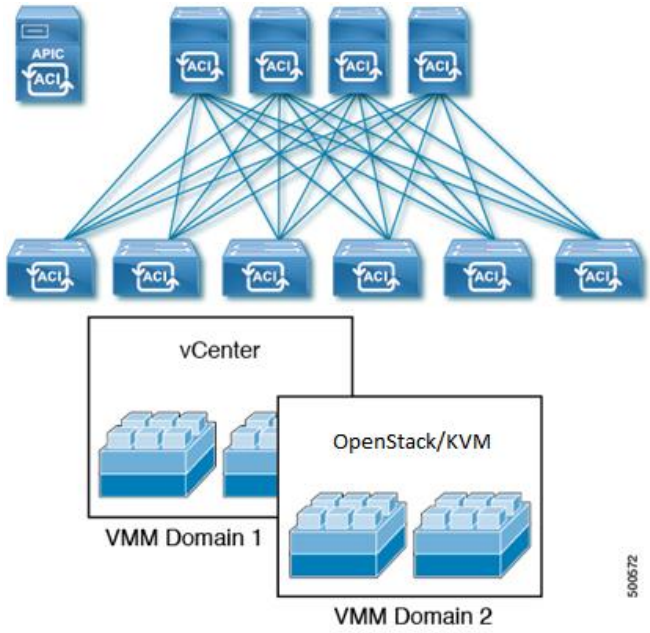

VMM domains provide the following:

- A common layer in the ACI fabric that enables scalable fault-tolerant support for multiple VM controller platforms.
- VMM support for multiple tenants within the ACI fabric.
- Automated static or dynamic VLAN allocations from specified VLAN pools.

VMM domains support micro segmentation that enables the following:

- Stateless white list network access security with line rate enforcement.
- Per-microsegment granularity of security automation through dynamic Layer 4 - Layer 7 service insertion and chaining.
- Hypervisor agnostic micro segmentation in a broad range of virtual switch environments.
- ACI policies that easily move problematic VMs into a quarantine security zone.

When combined with intra-EPG isolation for bare metal and VM endpoints, micro segmentation can provide policy driven automated complete endpoint isolation within application tiers.

Process for Each Application VM Migration

Cisco IT moves application VMs to ACI using the same process as migrating VMs across data centers. Cisco IT developed power shell scripts to automate VM moves.

Process for Each Application VM Migration

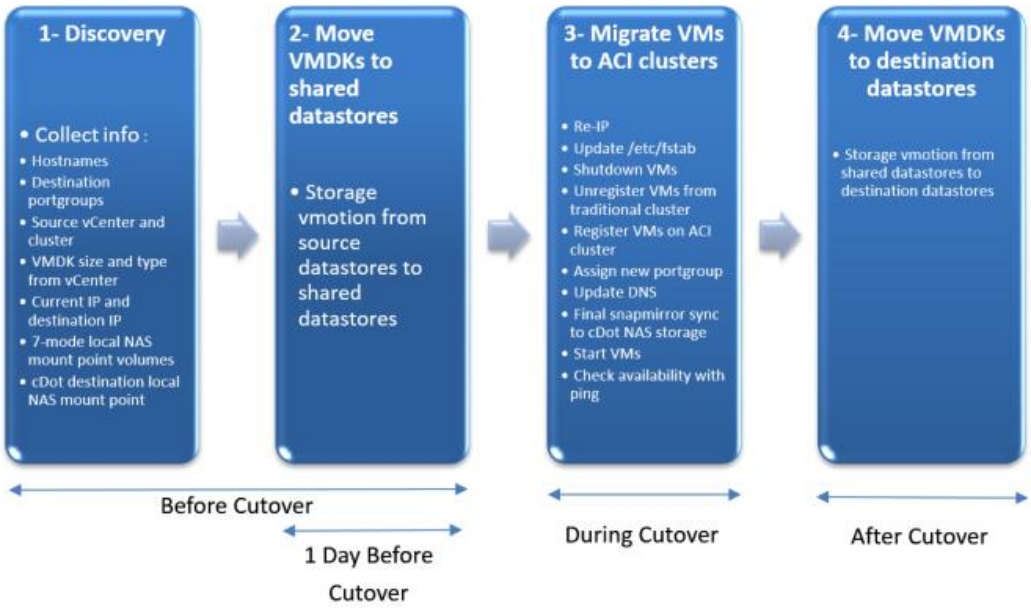

1. **Discovery Pre-Work before Cutover**

   This task is performed at any time before the Cutover change request (CR) migration. The goal of this step is to gather all the information that will be used during migration.

   During this step, information is gathered from sources illustrated below:

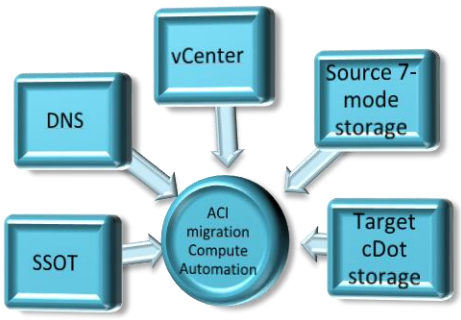

   Prepare SSOT (Single Source of Truth) with csv file or CMDB API.

Running the **1- Discovery** script gathers the following information that is used during the migration:

- OS, Lifecycle, Data Center, traditional and ACI cluster, ACI destination port group from SSOT
- VMDK size and storage type from vCenter
- New IPs from DNS by checking for example <hostname>-aci entry
- Mapping between Source 7-mode local NAS mount point volume and Target cDot local NAS mount point volume.
- This step also calculates total storage provisioned/used for each destination ACI cluster. If available ACI cluster datastores do not have enough space, additional storage is added on the ACI cluster datastores at this time.

2. **Storage vmotion to Move VMDKs to Shared Datastores**

CISCO IT deployed 3 * 25TB shared datastores mounted on traditional and ACI clusters. They are dedicated to ACI migrations and are used as temporary shared datastores to minimize downtime during a cutover window.

This step is done 1 day before Cutover. It moves online VMs to shared datastore. Running the **2- Move VMDKs to shared datastores** script starts VM storage vmotion from source datastores to shared datastores, automatically initiating several concurrent storage vmotions.

3. **Migrate VMs to ACI During Cutover**

This step moves the VMs to destination clusters with downtime during Cutover window. Running the **3- Migrate VMs to ACI clusters** script executes the migration of VMs to destination ACI clusters.

This script automates the following tasks:

- reIP the host with new IPs
- update /etc/fstab with cDot new
- Shutdown VMs gracefully
- Unregister VMs from traditional source cluster
- Register VMs to destination ACI cluster
- Assign new Port group to VMs
- Import VM permissions

- Update DNS update with new IPs
- Final snapmirror sync from 7-mode storage to cDot storage for all NAS local mount points
- Start VM
- Check availability with ping on destination ACI cluster

4. **After Change Request Cutover - Move VMDKs to destination Datastores**

This step performs storage vmotion from shared datastores to destination datastores. During this phase, VMs move from shared datastore to destination datastore cluster. According to the initial type of storage on the source cluster, a VM moves to a NAS or SAN datastore dedicated to the destination ACI cluster.

This phase starts when the application team gives their sign-off. This is done after Cutover, as this is a normal non-disruptive activity.

Running the **4- Move VMDKs to destination datastores** script executes this storage vmotion to destination datastore cluster.

Compute Migration Automated Tasks Checklist for Each Application Migration

TASKS Checklist for ACI Migration Automation
1. Run *1-Discovery* script to collect all information needed for the migration.
2. Add storage to ACI cluster datastores, if needed.
3. Initiate snapmirror from 7-mode storage to cDot storage for all local NAS mount points.
4. Run *2-Move VMDKs to shared datastores* script one day before change request to storage vmotion VMs from source datastore to shared datastores.
5. Run *3- Migrate VMs to ACI clusters* script during cutover to reIP, update /etc/fstab, shutdown VMs and start VMs migration from traditional to ACI clusters.
6. Update DNS with new IPs before starting VMs.
7. Launch final snapmirror sync before starting VMs.
8. Start VMs.
9. Ping VMs to check availability.
10. Hand over to application team to validate and sign-off.
11. After cutover, run *4- Move VMDKs to destination datastores* script to move VMs from shared datastores to destination datastores.

Cisco IT ACI Hadoop Migration Case Study

The logical view of Cisco IT pre-ACI Hadoop ecosystem is illustrated below.

Pre-ACI Hadoop Ecosystem

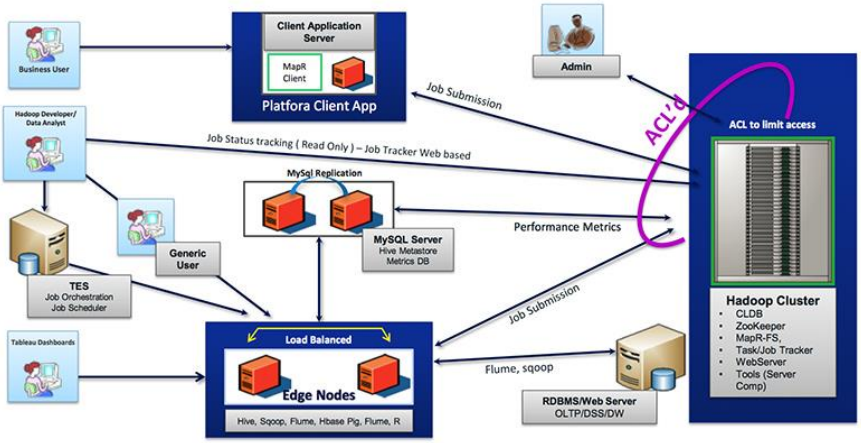

The physical view of Cisco IT pre-ACI Hadoop ecosystem below is a Cisco validated design.

Pre-ACI Hadoop Physical System

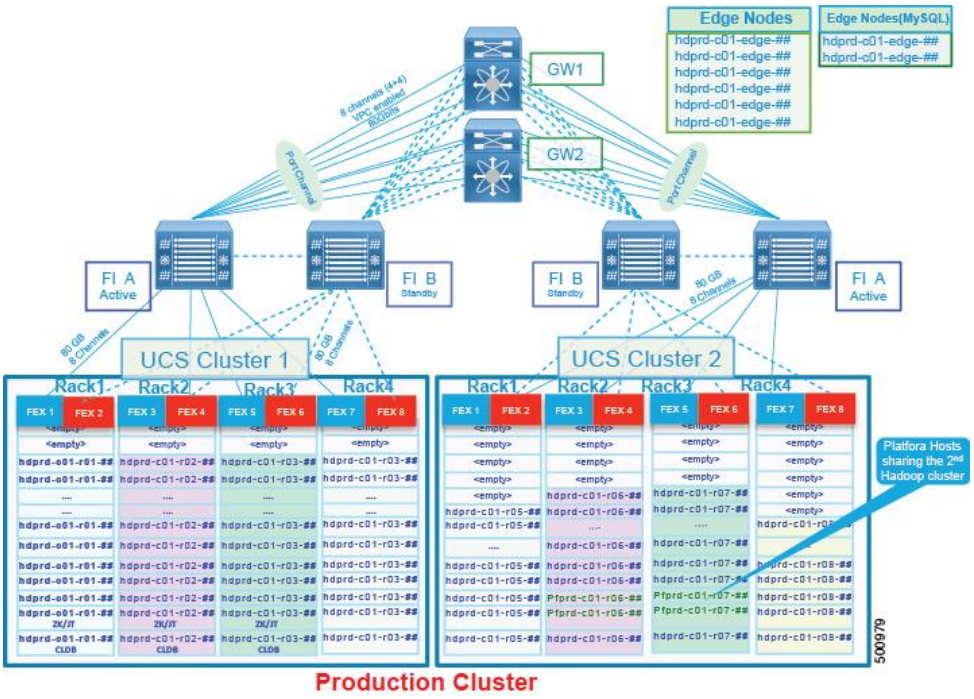

Identifying logical flows of data in the physical system is a challenge.

Tetration Validated/Enhanced the Hadoop Application Dependency Mapping

Tetration uses machine learning to present grouped flows that are labelled according to the naming conventions in the data center and the labels can be color coded. The various Hadoop cluster flows are pictured in the following color coded Tetration screen.

Cisco IT Hadoop Tetration Pre-ACI Application Dependency Map

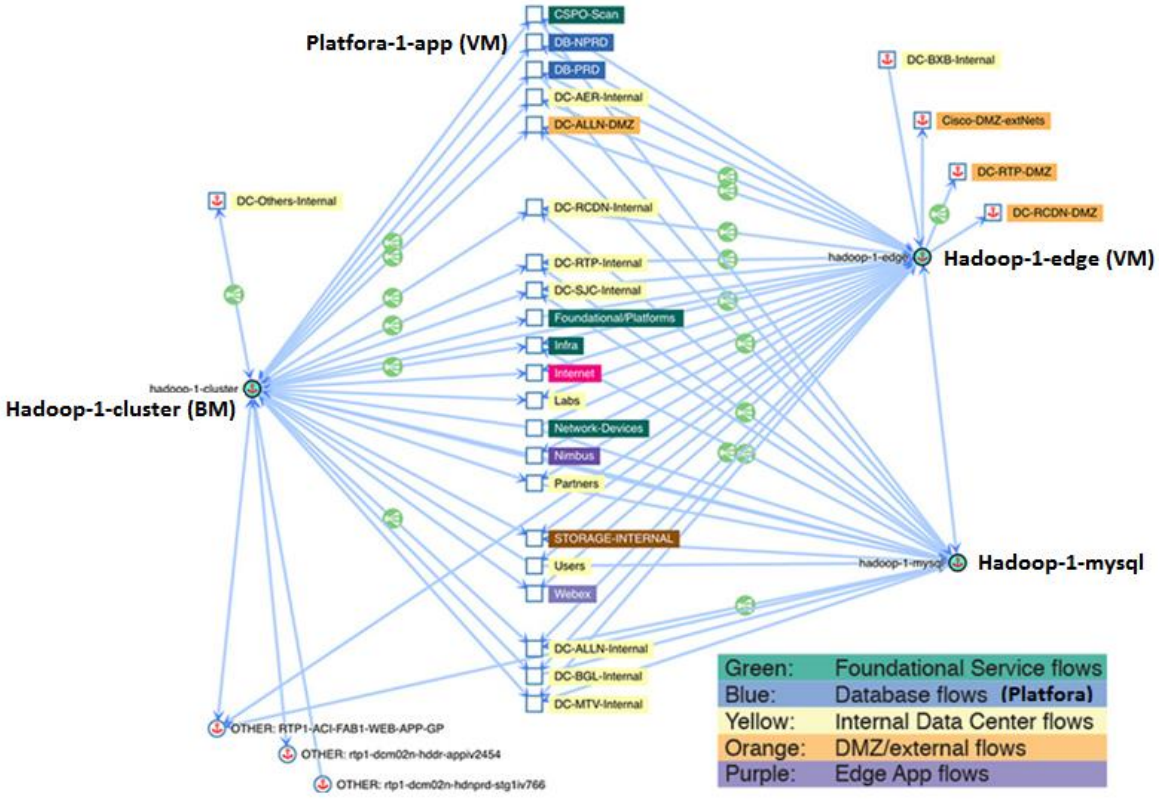

The Hadoop clusters are arrayed around the perimeter of the screen. The color-coded types of flows include the following:

- Green: Cisco IT foundational services, including LDAP, OAM, OCM, etc.
- Blue: Database Hadoop flows, including Platfora
- Yellow: Cisco enterprise internal
- Orange: DMZ/external flows
- Purple: Edge application flows

The following illustration is an example of where in the data center topology the internal and DMZ flows occur. Cisco IT uses two routing contexts (VRFs) within the ACI fabric, one for DMZ/external and one for internal. This assures that there is complete isolation between the DMZ and internal security zones.

Tetration Identified Rogue Hadoop Flows within the Internal Data Center and DMZ Flows

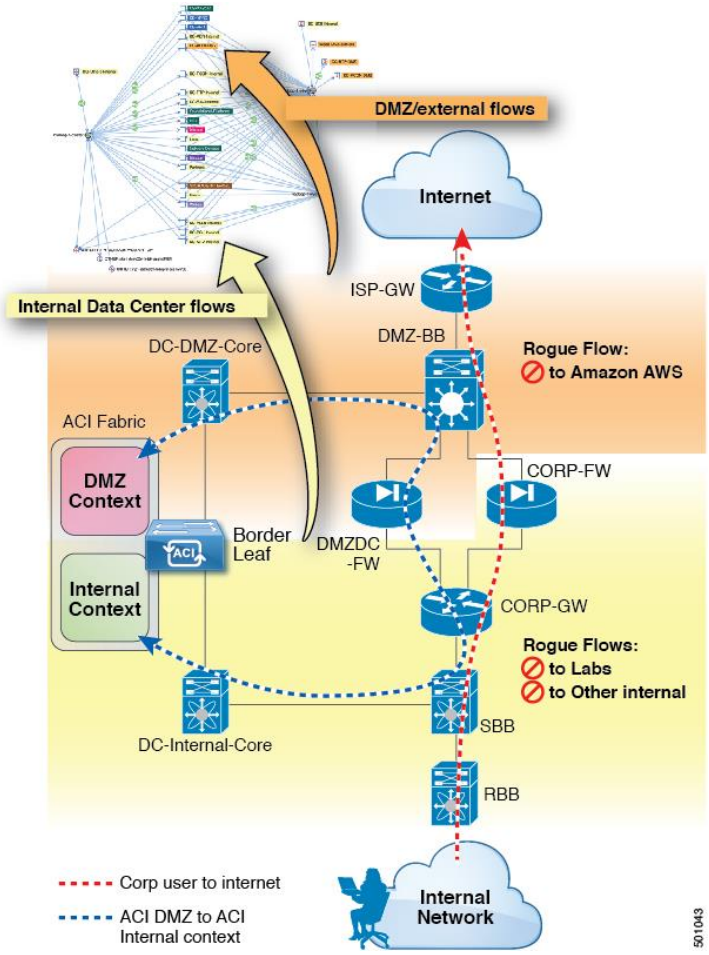

While the information Cisco application developer owners, the security team, and the networking teams provided the migration team was comprehensive, there were several surprises Tetration uncovered:

- Some Cisco enterprise internal flows were not known to any of the teams. Examples include flows to labs. These flows were not seen as problematic or security concerns.
- Some DMZ/external flows were going to Amazon AWS. This was a surprise that was a security concern.

Tetration confirmed all the TCP/IP ports that Hadoop used. This validation enabled specifying white list contract filters that would not cause problems by inadvertently blocking required ports.

The figure below illustrates the ACI allowed list contract specifications Tetration identified for the Cisco migration of its Hadoop deployment to ACI.

Tetration Flows Validate ACI clients-to-hadoop-cluster Contract Filter Specifications

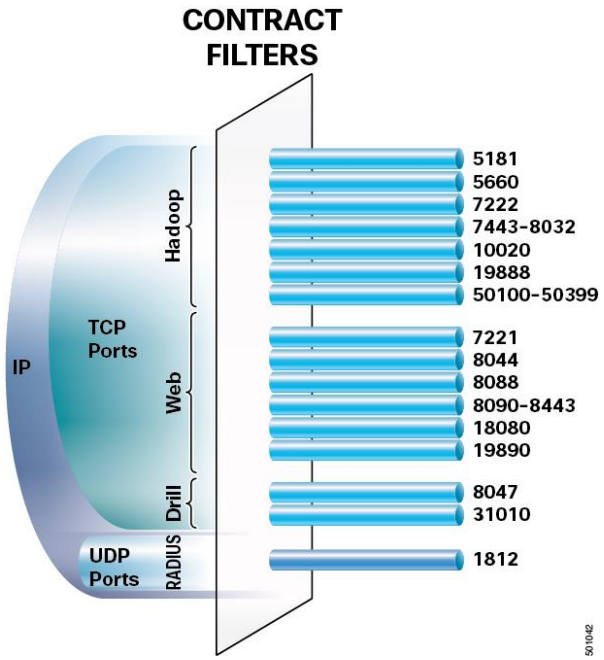

Tetration can export ACI contract specifications in various formats, including XMP, JSON, and YAML. Cisco IT chose to incorporate the contract specifications into its standard YAML library which was then posted to ACI. Cisco IT used Tetration to verify the contract specifications and assembled the YAML contract code for various contracts that specify how data flows are allowed between Hadoop EPGs.

ACI Hadoop Application Profile/EPG/Contract Policies

After understanding the application's dependency, it was very easy for Cisco IT to map the application to application profiles with their corresponding EPGs. Then, it was very simple to migrate all the Hadoop applications from the traditional network to the ACI fabric.

The application owner and the Cisco security teams chose to enforce strict limits on communications between clients and the Hadoop cluster as well as between the Platfora application and the Hadoop cluster. Communications between other Hadoop EPGs were

set to allow-all, with the expectation that these settings would be review in the future and revised accordingly.

ACI Hadoop Application Policies

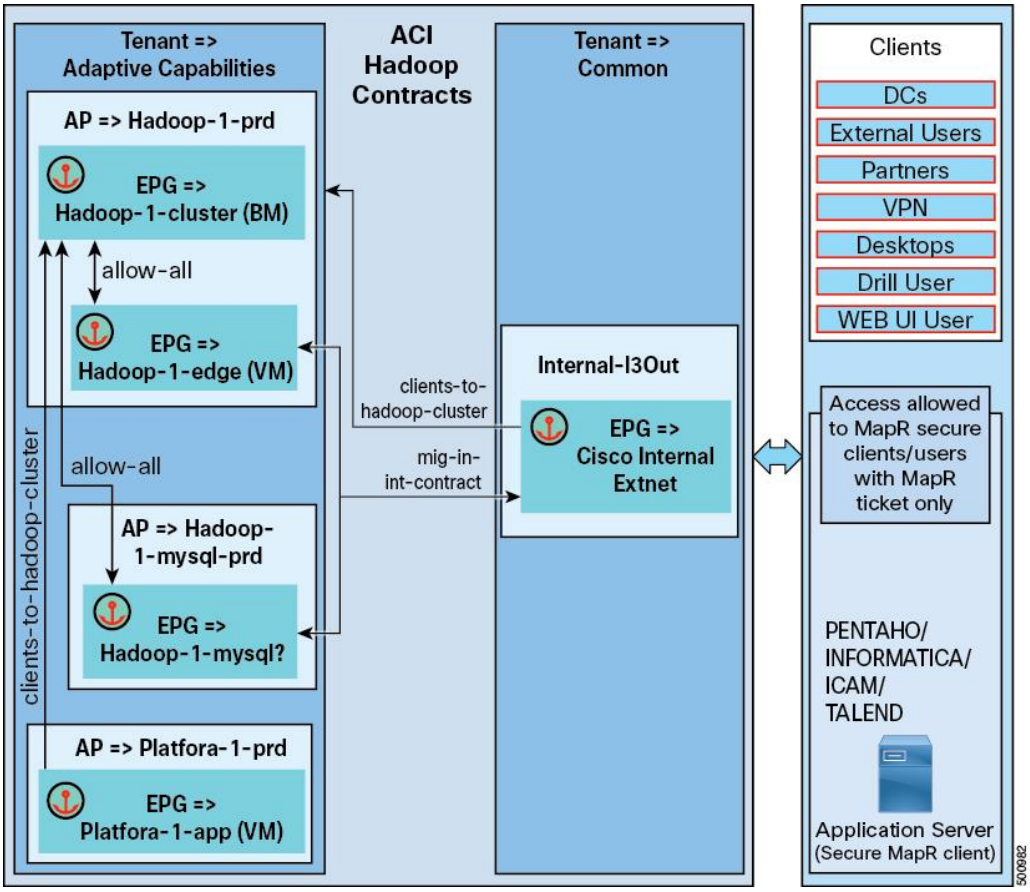

Contracts are directional; they are provided, consumed, or both. The `cisco-internal-extNet` EPG provides the `clients-to-hadoop-cluster` contract. The `hadoop-1-cluster` EPG consumes the `clients-to-hadoop-cluster` contract. The filters in this contract specify which ports are open for inbound client connectivity that connects to the `hadoop-1-cluster` EPG. The `clients-to-hadoop-cluster` contract is reused for connectivity between the `platfora-1-app` and `hadoop-1-cluster` EPGs.

The `clients-to-hadoop-cluster` contract is listed below.

ACI clients-to-hadoop-cluster Contract YAML Code

```
Contract name: clients-to-hadoop-cluster
          scope: 'Private Network' #VRF
          subjects:
             -#Hadoop
                name: 'tcp-5181'
                isUniDirectional: True
                filtersIntoEPG:
                   - 'dst-tcp-5181-filter'
             -#Hadoop
                name: 'tcp-5660'
                isUniDirectional: True
                filtersIntoEPG:
                   - 'dst-tcp-5660-filter'
             -#Hadoop
                name: 'tcp-7222'
                isUniDirectional: True
                filtersIntoEPG:
                   - 'dst-tcp-7222-filter'
             -#Hadoop
                name: 'tcp-7443'
                isUniDirectional: True
                filtersIntoEPG:
                   - 'dst-tcp-7443-filter'
             -#Hadoop
                name: 'tcp-8032'
                isUniDirectional: True
                filtersIntoEPG:
                   - 'dst-tcp-8032-filter'
             -#Hadoop
                name: 'tcp-10020'
                isUniDirectional: True
                filtersIntoEPG:
                   - 'dst-tcp-10020-filter'
```

```
-#Hadoop
    name: 'tcp-19888'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-19888-filter'
-#Hadoop
    name: 'tcp-50100-50399'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-50100-50399-filter'

-#Web
    name: 'tcp-7221'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-7221-filter'
-#Web
    name: 'tcp-8044'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-8044-filter'
-#Web
    name: 'tcp-8088'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-8088-filter'
-#Web
    name: 'tcp-8090'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-8090-filter'
-#Web
    name: 'tcp-8443'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-8443-filter'
```

```
-#Web
    name: 'tcp-18080'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-18080-filter'
-#Web
    name: 'tcp-19890'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-19890-filter'

-#Drill
    name: 'tcp-31010'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-31010-filter'
-#Drill
    name: 'tcp-8047'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-tcp-8047-filter'

-#RADIUS
    name: 'udp-1812'
    isUniDirectional: True
    filtersIntoEPG:
        - 'dst-udp-1812-filter'
```

Graceful Cutover to the ACI Hadoop Deployment

ACI Hadoop Production Hardware Topology

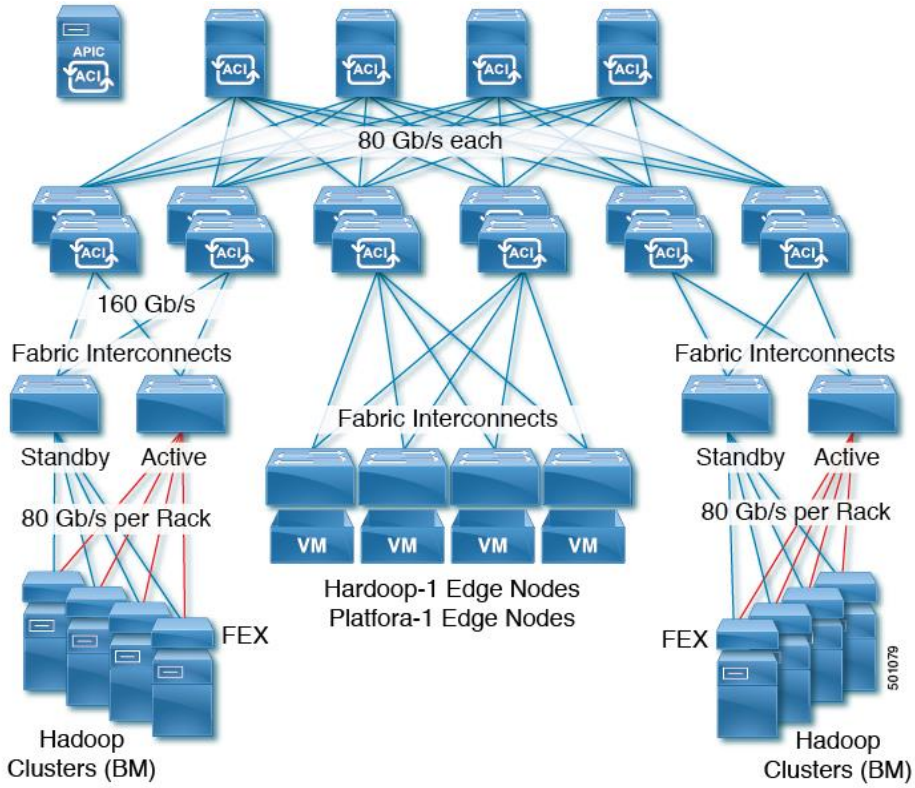

The key points Cisco IT observed during the Hadoop migration cutover were these:

1) Gracefully migrated the entire production Hadoop to ACI and performed a large scale in-place upgrade over a weekend with virtually no down time.

2) Client connectivity preserved to Hadoop was preserved.

3) Hadoop security was enhanced by adopting the allowed list contract policy model.

4) The ACI fabric provides a faster network that enabled scaling the Hadoop footprint from 116 nodes to 250 nodes across 22 RACKS.

Automated Provisioning of Secure Application Centric Cloud

The migration to ACI enables Cisco IT to greatly expand provisioning automation while letting end users do so in a secure self-service fashion.

Cisco IT Target ACI Deployment Model

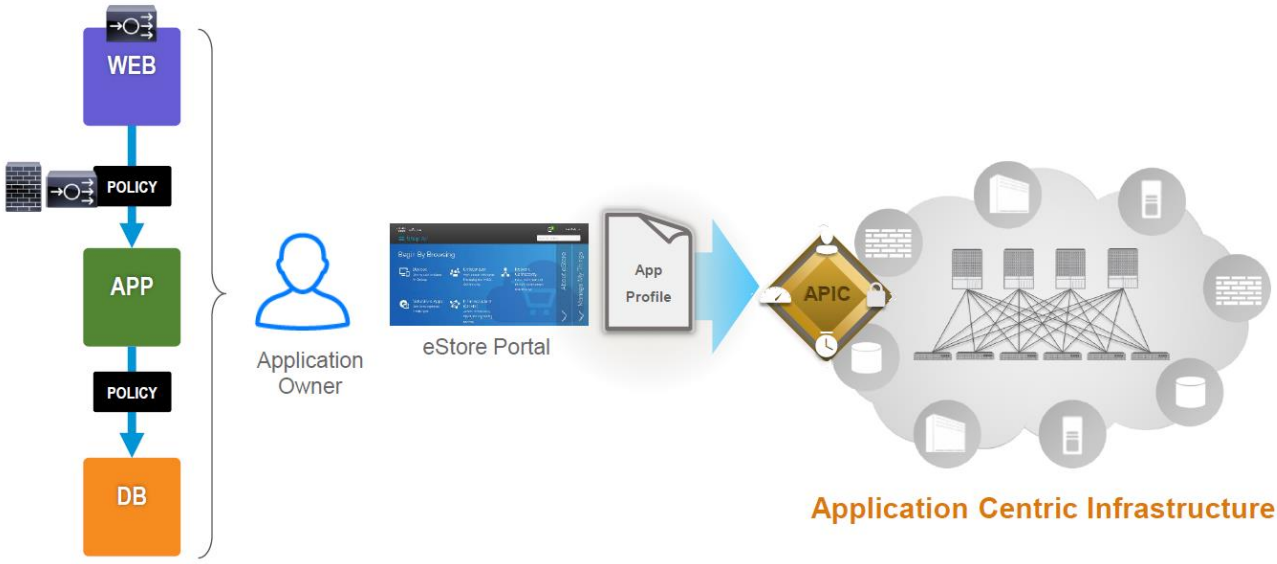

ACI enables Cisco IT operations to provide application developers direct access to standardized fabric infrastructure in a highly automated fashion through an open API while enforcing the security and governance requirements of the organization.

Cisco IT eStore Portal Enables Secure Self-Service Provisioning of ACI Resources

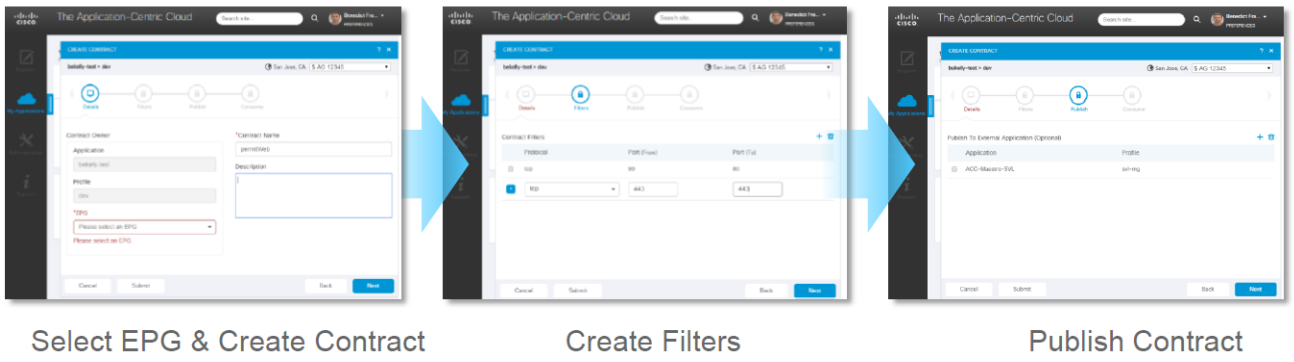

With the self-provisioning capabilities of the Cisco IT eStore Portal, application owners can provision the data center resources they need directly.

Conclusion

During the migration to ACI, Cisco IT found that the following guidelines prove useful as they migrate multiple data centers:

- Start small and gradually ramp up.
- Test and certify new code, features and functionality – focus on automation to speed up the certification process.
- Transition the workforce to think and operate differently. The productivity gains from ACI enable faster execution of business transforming objectives. So, IT staff are more closely aligned with business objectives and can think accordingly as they deliver on meeting those objectives.
- Scripting skills will help you on your journey.