FlexPod®

Flexpod Datacenter for SAP Applications with Cisco ACI and Vnomic Policy-Driven Automation

CONTENTS

Introduction

This document introduces FlexPod Datacenter for SAP applications with Cisco® Application Centric Infrastructure (Cisco ACI™) using Vnomic policy-based infrastructure. Vnomic is a member of the SAP Start-up Focus Program. This solution integrates Cisco ACI and Vnomic automation in a scale-out FlexPod architecture to automate the deployment and governance of SAP HANA and SAP Business Warehouse (BW) on SAP HANA landscapes.

This solution takes an application-centric approach in which Vnomic automation represents the application's software components and hosting infrastructure as a single logical abstraction. Vnomic translates application requirements into a set of declarative infrastructure policies that are automatically fulfilled and enforced by Cisco ACI. The result is a highly scalable application-centric platform that can efficiently instantiate, update, and decommission SAP HANA and SAP BW on SAP HANA landscapes quickly, securely, and risk free.

The notable features of this solution include:

- Prevalidated FlexPod architecture with prebuilt application-centric automation
- Fully automated software-as-a-service (SaaS) deployment of new instances of SAP HANA and SAP BW on SAP HANA landscapes in minutes
- SAP HANA Tailored Datacenter Integration (TDI)–compliant clusters
- Support for a variety of SAP BW and SAP HANA topologies with high availability and backup
- Capability to host servers on both physical and virtual infrastructure to optimize performance, increase infrastructure utilization, and reduce costs
- Multitenancy for safe concurrent hosting of mixed workloads across multiple organizations
- Flexible scale-out architecture that allows computing, network, and storage capacity to be added as needed
- Standardized configurations for computing and storage capacity, reducing zero-day costs and lead time

Application Deployment and Governance Challenges

With the growing sophistication and deployment complexity of next-generation application landscapes, such as SAP BW on SAP HANA, traditional approaches to application deployment and governance no longer meet business expectations. In a competitive world, customers must deploy services quickly and then rapidly and safely add or remove resources from already running applications to meet the demands of users, based on service-level requirements, time scheduling, and cost factors, while adhering to security and governance requirements.

Complex, large-scale, and dynamic application landscapes require a new approach to automation, which can't be achieved with today's tools that depend on risky manual processes and expert intervention. The next generation of automation must capture application requirements and lifecycle semantics in a set of abstractions to provide the fundamental operations needed to manage applications throughout their lifecycles. This application-centric perspective raises the application to the highest level of importance and views infrastructure as subservient to the application. In an application-centric world, the value of the infrastructure is determined by how efficiently and flexibly its capabilities can be used to fulfill the requirements of the application.

New Challenges with Programmable Infrastructure

Programmable infrastructure[1] has transformed infrastructure management from a logistically and human-intensive physical activity into a software configuration and programming activity. Converged infrastructure, such as FlexPod from Cisco and NetApp, has reduced the traditionally unbounded diversity of data centers to a well-defined set of resources represented as a small number of logical abstractions.

The new challenge is to map the application requirements to the semantics and APIs of each programmable infrastructure resource. For example, for a virtual server to access storage, the programmable computing resources have to be configured to interact with the programmable network resources to reach the programmable storage resources on the storage ports with the appropriate bandwidth and failover semantics. This goal requires semantic alignment across all elements of the programmable infrastructure.

Whether using graphical consoles, scripting and automation tools, or direct APIs, operations personnel are responsible for learning the semantic details of all programmable infrastructure, defining the specific configuration or programming steps, and understanding the implications and verifying the accuracy of the changes made.

Figure 1 shows the deployment cycle.

Figure 1. Deployment Cycle

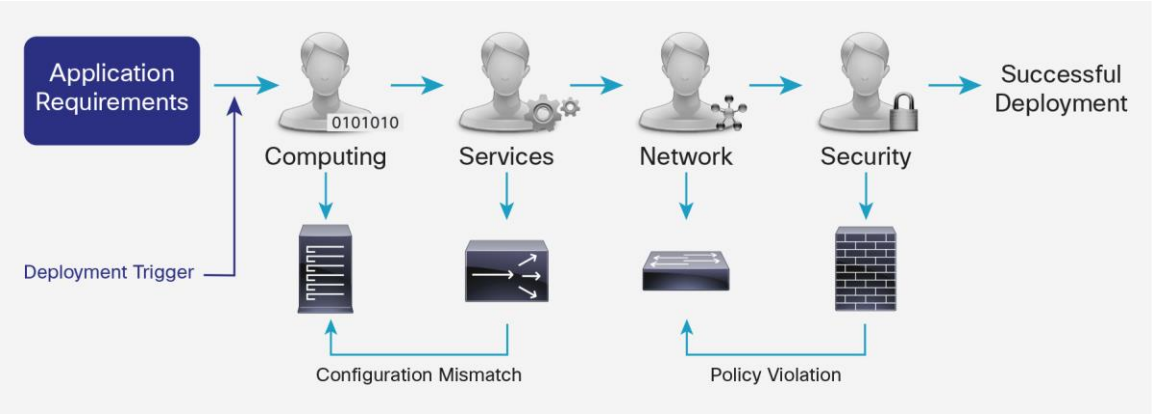

Solution Overview

SAP BW and SAP HANA landscapes consist of multiple application components across a multitier topology, with detailed requirements defining the way that these tiers communicate. The data center network must be configured to enable this complex communication while still meeting business needs for security, compliance, and service levels. Changes to a tier, such as support for increased demand, often require extensive configuration changes to the network devices.

[1] This document uses the term “programmable infrastructure” to include software-defined resources such as software-defined computing (SDC), software-defined networking (SDN), and software-defined storage (SDS).

The Vnomic Modeling Framework is used to create application models that express SAP BW and SAP HANA application semantics in terms of their structure, behavior, interdependencies, lifecycle, and infrastructure requirements.

The Vnomic Desired State Controller (VDSC) then translates the application models into a set of declarative policies that are automatically fulfilled and enforced by the hosting Cisco ACI. Cisco ACI provides flexible application policy and infrastructure control that can dramatically accelerate application deployment and operations through centralized configuration, testing, and monitoring of the entire data center infrastructure, including network connectivity, security, and other Layer 4 through 7 services (Figure 2).

Figure 2. Application Service Model for SAP BW and SAP HANA Landscapes with FlexPod

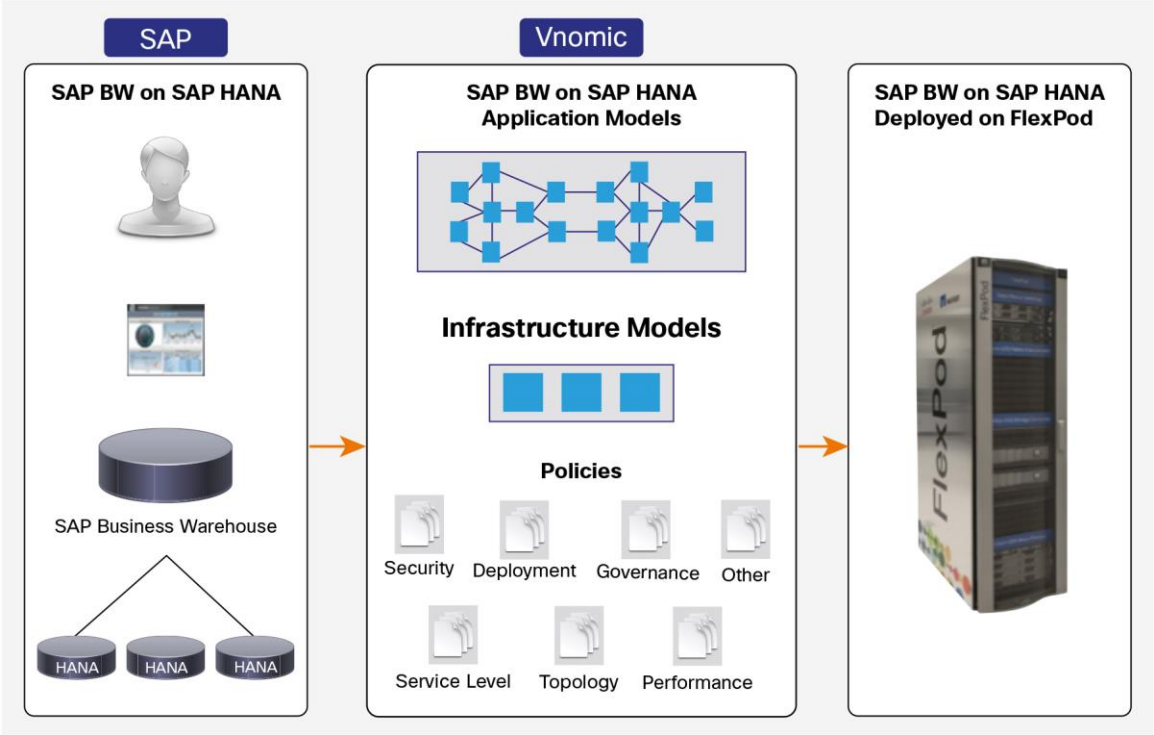

Solution Components

SAP Business Warehouse Powered by SAP HANA

The reporting, analysis, and interpretation of business data is of central importance to a company in helping ensure its competitive edge, optimizing its processes, and enabling it to react quickly and in line with the market. To achieve these goals, SAP BW provides a business intelligence platform, data warehousing function, and suite of business intelligence features. It captures, stores, and consolidates vital customer information with the enterprise data warehouse platform. It also allows customers, in real time, to tightly integrate warehousing capabilities for a single version of the truth, decision-ready business intelligence, and accelerated operations.

SAP HANA combines database, data processing, and application platform capabilities in memory, with exceptional performance. This new architecture enables converged online transaction processing (OLTP) and online analytical processing (OLAP) within a single in-memory column-based data store with atomicity, consistency, isolation, and durability (ACID) compliance, while eliminating data redundancy and latency. The SAP HANA platform provides libraries for predictive analytics, planning, text analytics, spatial processing, and business analytics. By providing such advanced capabilities, all on the same architecture, it further simplifies application development and processing across big data sources and structures. These capabilities make SAP HANA the most suitable platform for building and deploying next-generation, real-time applications and analytics.

SAP BW and SAP HANA landscapes consist of multiple application components across a multitier topology. SAP BW is deployed as a horizontally scalable application-server tier with a coordinator node and shares configuration storage with specific computing, network, and storage requirements. SAP HANA is a highly scalable in-memory database deployed as a multinode cluster with specific requirements for high-performance processing, network connectivity and performance, large memory capacity, and cluster shared state with specific storage resources.

Cisco Application Centric Infrastructure

With Cisco ACI, Cisco aims to significantly accelerate application deployment times from weeks to minutes by reducing complexity through application-based policies. By treating applications as the fundamental building blocks of IT, Cisco ACI integrates physical and virtual devices into a dynamic pool of IT resources that can be automatically provisioned and reprovisioned. With this approach, Cisco ACI enables an agile, open, and secure solution that offers businesses exceptional agility.

One of the main innovations in Cisco ACI is the introduction of a highly abstracted interface to express the connectivity of application components along with high-level policies governing that connectivity. Cisco ACI provides the critical framework that enables the complex communication requirements of sophisticated applications, such as SAP BW on SAP HANA, to be implemented in real-world converged infrastructure.

Cisco ACI provides:

- A declarative model to provision and control the network through policy, eliminating the need for special state-dependent logic and sequencing during orchestration
- A set of application-centric abstractions, including application components, contexts, endpoints, and contracts. Since these abstractions exist in both the application and network domains, the required semantics are fully communicated and thus upheld by the fabric.

Cisco Application Policy Infrastructure Controller

The Cisco Application Policy Infrastructure Controller (APIC) is the main architectural component of the Cisco ACI solution. It is the unified point of automation and management for the Cisco ACI fabric, policy enforcement, and health monitoring. The APIC is a centralized clustered controller that optimizes performance, supports any application anywhere, and unifies operation of physical and virtual environments. The controller manages and operates a scalable multitenant Cisco ACI fabric (Figure 3).

Figure 3. Cisco APIC Fabric Control

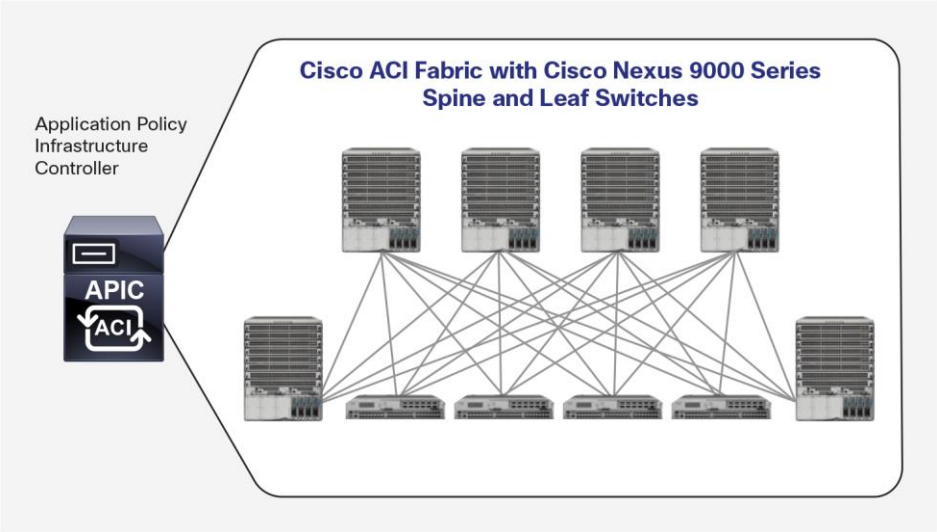

Application Network Profiles

APIC application network profiles uniquely capture application connectivity and governance requirements in terms of application abstractions, including application components, endpoints, and contracts, effectively bridging the traditional semantic gap between application requirements and network implementation. APIC uses a logical representation of application requirements to automatically provision and continuously adapt the fabric for optimized application operation (Figure 4).

Figure 4. Application Network Profile

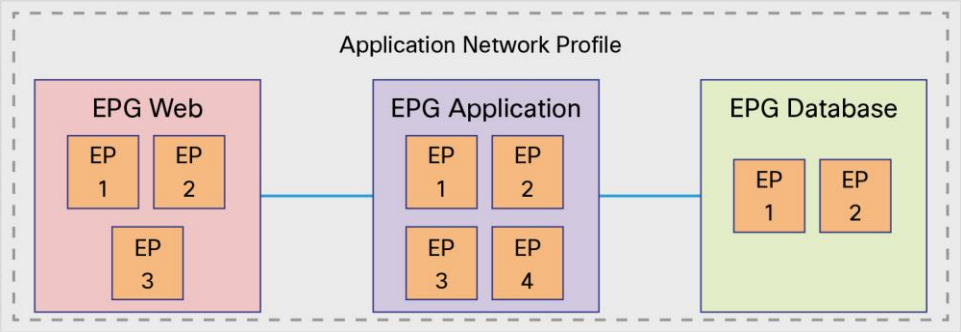

An application network profile within the fabric is a collection of the endpoint groups, or EPGs (a logical grouping of similar endpoints representing an application tier or set of services that require a similar policy), their connections, and the policies that define those connections. The application network profile is the logical representation of all the components of the application and their interdependencies in the application fabric (Figure 5).

Figure 5. Application Network Profile EPGs

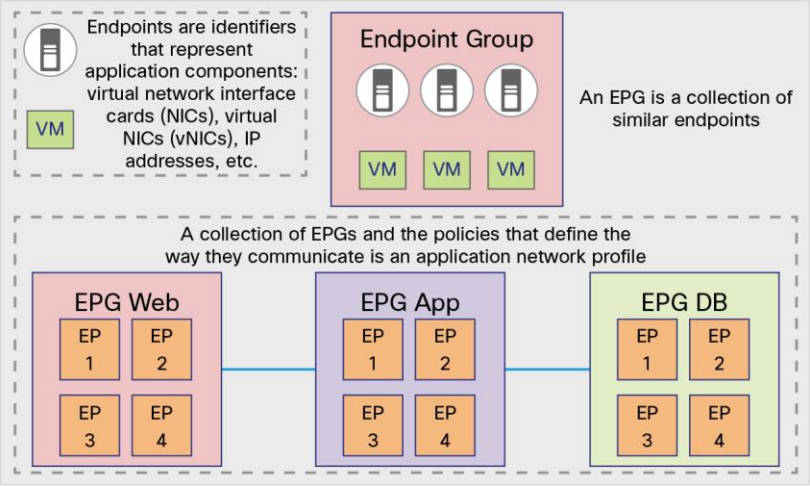

Application network profiles are designed to be semantically aligned with the way that applications are designed and deployed. The configuration and enforcement of policies and connectivity are then handled by the system through the APIC rather than by an administrator.

NetApp FAS and Clustered Data ONTAP

NetApp storage solutions are user friendly and easy to manage and automate. They can be deployed quickly, and they offer increased availability while consuming fewer IT resources, with dramatically lower lifetime total cost of ownership (TCO). A NetApp solution consists of hardware in the form of controllers and disk storage and the NetApp clustered Data ONTAP storage operating system, which is the most deployed storage operating system in the industry.

FAS storage systems with clustered Data ONTAP provide an advanced software-defined storage (SDS) platform well suited for heterogeneous workloads. They provide continuous nondisruptive operation, storage efficiency, transparent storage scalability, and snapshot and backup capabilities uniquely designed for SAP HANA.

Software-Defined Storage

FAS provides advanced SDS with storage services in storage virtual machines (SVMs). SVMs are logical storage servers created on demand on the basis of application requirements. SVMs can use any of the physical resources available across the entire cluster, and they can serve data to clients using any configured protocol. In this way, storage services can be delivered on demand to different business units or applications according to their specific requirements, while allowing storage resources to be managed as a single unit rather than as isolated silos.

Nondisruptive Operations

By using SVMs to abstract the logical storage services from the physical hardware, data volumes and host and client connections, including both block and file protocols, can be moved transparently between nodes. This capability enables workload balancing, data-lifecycle management, and nondisruptive maintenance activities. Hardware and software upgrades can be performed with client access available completely online by evacuating data volumes and host connections from a given high-availability pair, performing the necessary maintenance or upgrades, and then migrating those hosts back to the system.

In this way, a complete hardware refresh can be performed with no disruption to the clients, applications, or hosts that access the system.

FAS storage controllers are deployed in high-availability pairs to enable uninterrupted operations during maintenance activities and provide fault tolerance in the event of a controller failure. Each controller in a high-availability pair is connected to the same set of disk shelves to allow a surviving controller to maintain access to the disks and provide storage services to clients while the other controller is offline. This capability allows you to upgrade controller hardware and software and perform other hardware-upgrade maintenance activities without disrupting client data access.

Storage Efficiency

Clustered Data ONTAP provides substantial space savings, allowing more SAP HANA data to be stored at a lower cost. Volumes are created through virtual sizing. They appear to be provisioned to their full capacity, but are actually created much smaller and use additional space only when it is needed. Extra unused storage is shared across all volumes, and the volumes can grow and shrink on demand.

NetApp snapshots provide automatically scheduled point-in-time copies of only blocks that have been changed by write operations, with no performance penalty. Snapshot copies consume little storage space because only changes to the active file system are written. Individual files and directories can easily be recovered from any snapshot copy, and the entire volume can be restored to any snapshot state in seconds.

NetApp FlexClone volumes consume almost no space and provide instant virtual copies of data sets. The clones are writable, but only changes to the original are stored. FlexClone technology rapidly creates additional, space-efficient data copies well suited for SAP development and test environments.

Deduplication provides redundant data blocks in primary and secondary storage with flexible policies that determine when the deduplication process is run.

Compression can provide additional space savings when run alone or together with deduplication.

NetApp SnapMirror data replication software can asynchronously replicate SAP HANA snapshots to other storage controllers for backup and business continuity.

Transparent Scalability

Clustered Data ONTAP addresses the growth of storage requirements in an enterprise by using a scale-out approach. When requirements increase, storage controllers can be added as needed to the same storage cluster, providing additional CPU, memory, network, and disk resources to support the growing workload. Data ONTAP clusters can support up to 24 nodes per cluster with up to 104 petabytes (PB) of storage capacity, all managed as a single system.

High-availability pairs within the cluster can be different controller models with different drive types, from high-performance solid-state disk (SSD) drives to high-capacity SATA drives, to provide appropriate storage based on workload requirements while still maintaining a single management interface.

SAP HANA Snapshots and Backup with NetApp Snap Creator

NetApp Snap Creator with a SAP HANA plug-in is designed specifically for the business data backup requirements of SAP HANA and SAP HANA business applications. SAP HANA nodes can be backed up as point-in-time application-consistent snapshots. These snapshot copies can be scheduled for automatic backup at any frequency and can be created while SAP HANA is running, thus not delaying business processing. SAP HANA snapshot copies can be replicated using SnapMirror to local or remote data centers for disaster recovery and business continuity.

Vnomic Automation

Policy-based delivery and governance is a recent technical breakthrough, tying the power of application modeling to the flexibility of programmable infrastructure. The Vnomic Policy Defined Application Delivery and Governance Platform is an application-centric automation solution capable of controlling the most sophisticated applications and infrastructure.

This section describes the main attributes of the Vnomic automation technology.

Application-Centric Approach

Vnomic takes an application-centric approach in which the application intent and requirements guide the entire automation process. This application-centric perspective requires the automation platform to understand the application semantics. Application requirements must be completely and concisely described so that they can be automatically fulfilled by infrastructure capabilities (Figure 6).

Figure 6. Application-Centric Approach

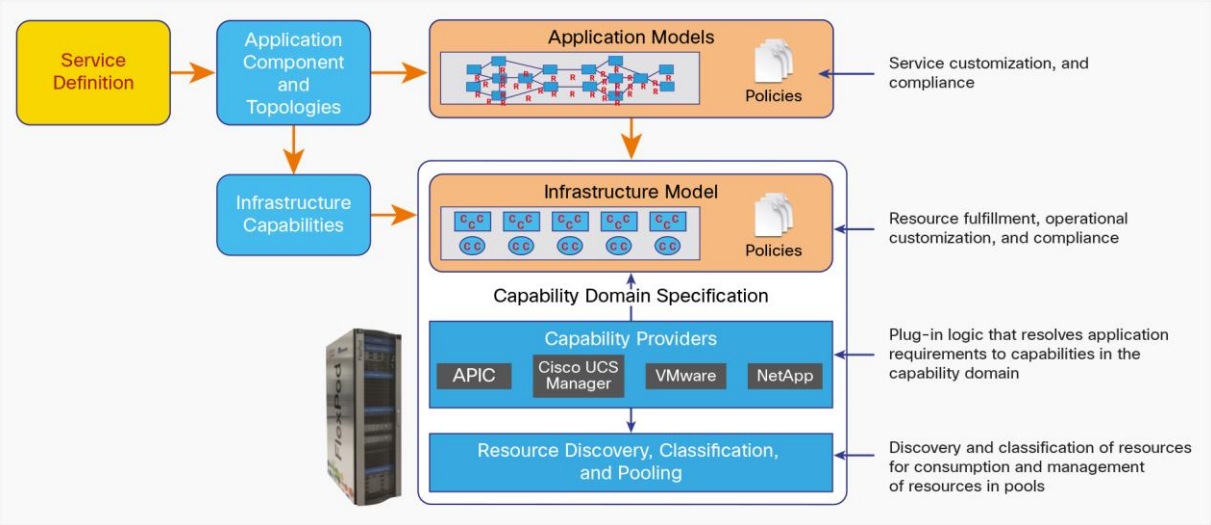

Application Modeling

Vnomic application models express application semantics in terms of their structure, behavior, interdependencies, infrastructure requirements (computing, storage, and network resource needs), scalability and availability, service-level objectives, and security and compliance. In this document, this set of application-centric semantics representations are referred to as application models. These application models represent the entire application topology and the services offered by the application as a single logical abstraction with well-defined lifecycle operations and transformations. These capabilities enable automation of the most common operations such as initial provisioning, update, topology change, migration, decommissioning, and deprovisioning processes.

Infrastructure Modeling

Given an application model, the automation process must determine how to fulfill the requirements of the application with the available infrastructure. The VDSC does this by expressing the semantics of FlexPod's software-defined computing, networking, and storage as capabilities: that is, as the resources, features, and behaviors it can provide to application components. In this document, this set of infrastructure semantics is referred to as an infrastructure model. The sole purpose of infrastructure models is to describe the way that a managed infrastructure can fulfill application requirements to enable Vnomic to automatically select the resources needed to meet application requirements. This capability eliminates the need for users to manually perform or supervise the selection of the appropriate types and amounts of infrastructure resources.

Matching Application Requirements to Infrastructure Capabilities

Using the infrastructure requirements contained in application models and the infrastructure capabilities contained in infrastructure models, the process of matching requirements to capabilities can be generalized and fully automated (Figure 7).

Figure 7. Matching Requirements to Infrastructure

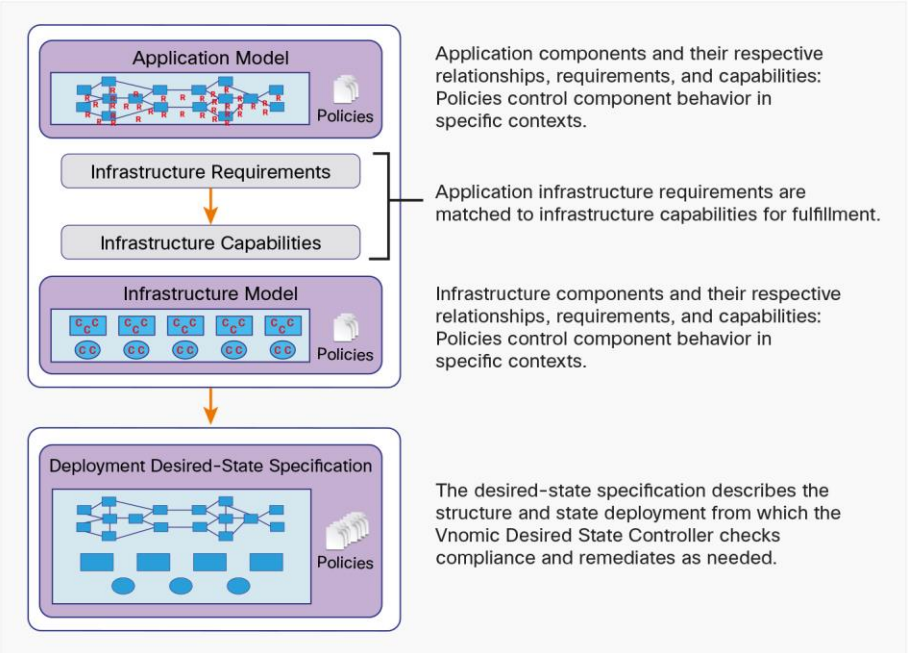

Policies Enabling Constraint-Based Control and Customization

Vnomic uses policies as declarative control and component customization specifications to avoid having to create large numbers of very specialized application models. Policies are sets of constraints that can be applied to one or more modeled entities. Sets of policies can be associated with applications and infrastructure models to:

- Specify specific configuration values
- Select specific modes of operation
- Enforce specific environment constraints
- Enforce specific IT policies
- Enforce customer business or industry policies

Policy-Defined Declarative Desired State

Application and infrastructure modeling is extremely useful for abstracting and expressing deployment semantics. However, the use of this information to physically automate the necessary application lifecycle operations, such as initial provisioning, update, topology change, migration, decommissioning, and deprovisioning operations, still presents significant challenges. Vnomic uses a declarative desired-state (DDS) model to control the state of managed software and infrastructure. The DDS uses a desired-state specification (DSS) that is automatically computed by transforming the information in the application models and infrastructure models into a set of provisioning, deployment, and configuration policies. The policies specify **what** is desired, not **how** to achieve the desired state.

The resulting policies are partitioned into two classes:

- Delegated policies (DPs): The VDSC delegates these policies to underlying infrastructure element managers and software automation tools that are responsible for fulfilling the way that the policies work. For example, many infrastructure-related policies are passed directly to the APIC and to Cisco UCS® Manager, dramatically simplifying infrastructure automation.
- Managed policies (MPs): The VSDC fulfills these policies itself, by orchestrating the necessary operations on software components or infrastructure. Managed policies are required when a declarative automation component is not available to fulfill the policies.

The VDSC applies these policies in its control loop to move the managed software and infrastructure to the desired state.

An important benefit of the DDS approach is that it can work throughout the application lifecycle, not just for the initial deployment (Figure 8).

Figure 8. Vnomic Declarative Desired State

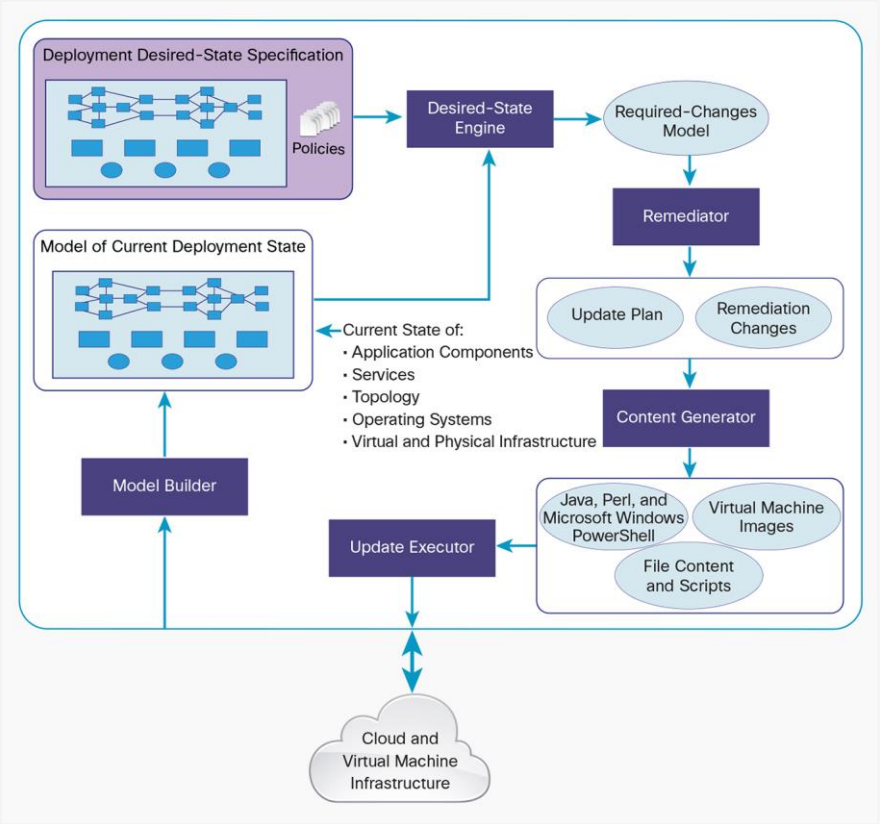

Prevalidated Application Automation for FlexPod

The application and infrastructure models included in this solution are validated in specialized labs for end-to-end system integration and certification with prescribed combinations of software and FlexPod architecture. This testing helps ensure that the automation provides the desired functions, performance, and security. This validation results in application models and infrastructure designs that enterprises and service providers can use to build robust and supportable software-as-a-service (SaaS) implementations.

Most important, enterprises and service providers don't have to take responsibility for building the automation or configuring the infrastructure in specialized ways, and they are assured that the automation will provide the required performance and security. Enterprises and service providers can instead focus on selecting the right service offerings, such as application topologies and infrastructure capacity, to support their SaaS objectives.

Cisco UCS Manager

Cisco UCS Manager provides software-defined computing resources through the unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS). Cisco UCS Manager provides unified management domain with centralized management capabilities and controls multiple chassis and thousands of virtual machines.

Service Profiles

An important feature of Cisco UCS Manager is its use of service profiles to provision and manage Cisco UCS blade servers and rack servers and their I/O properties within a single management domain. The computing infrastructure policies needed to deploy applications are encapsulated in the service profiles. The policies coordinate and automate element management at every layer of the hardware stack, including RAID levels, BIOS settings, firmware revisions and settings, server identities, adapter settings, VLAN and VSAN network settings, network quality of service (QoS), and data center connectivity. Vnomic automation creates service profiles for all Cisco UCS servers it manages.

A service profile consists of a software definition of a server and the associated LAN and SAN connectivity that the server requires. When a service profile is associated with a server, Cisco UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the service profile. Service profiles improve IT productivity and business agility. With service profiles, infrastructure can be provisioned in minutes instead of days, shifting the focus of IT from maintenance to strategic initiatives.

VMware vSphere

VMware vSphere provides the virtualization framework for this solution. Virtual capacity is provisioned and managed by Vnomic automation.

Operational Overview

This section provides an operational overview of the steps performed by the VDSC for automated application deployment.

Basic Theory of Operation

The objective of the Vnomic Policy Defined Delivery and Governance Platform is to compute the policies and orchestrations needed to implement the specifications contained in application models for the target infrastructure (Figure 9).

Figure 9. Vnomic Policy Defined Delivery and Governance Platform

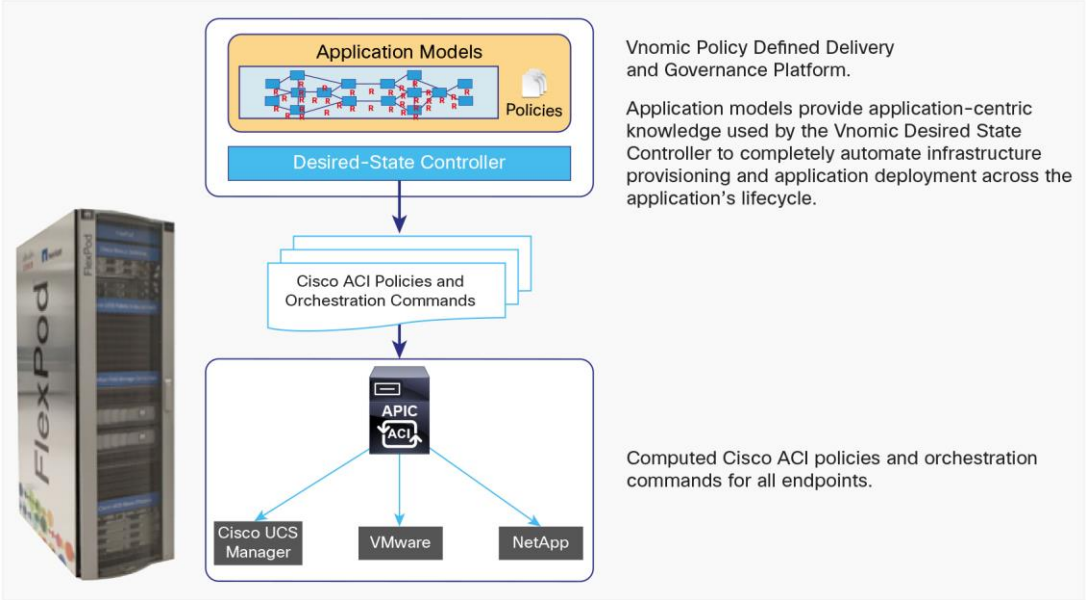

The VDSC performs three fundamental steps (Figure 10):

1. Upon receiving a service request for the initial deployment, it computes the desired-state specification. The DSS is updated only if the user changes the application topology or the application policies or requirements.
2. It computes the provisioning plan to provision all the required infrastructure resources in the correct order.
3. It computes the deployment plan to deploy any needed software and configure it correctly to meet the desired state in the correct order.

Figure 10. VDSC Process

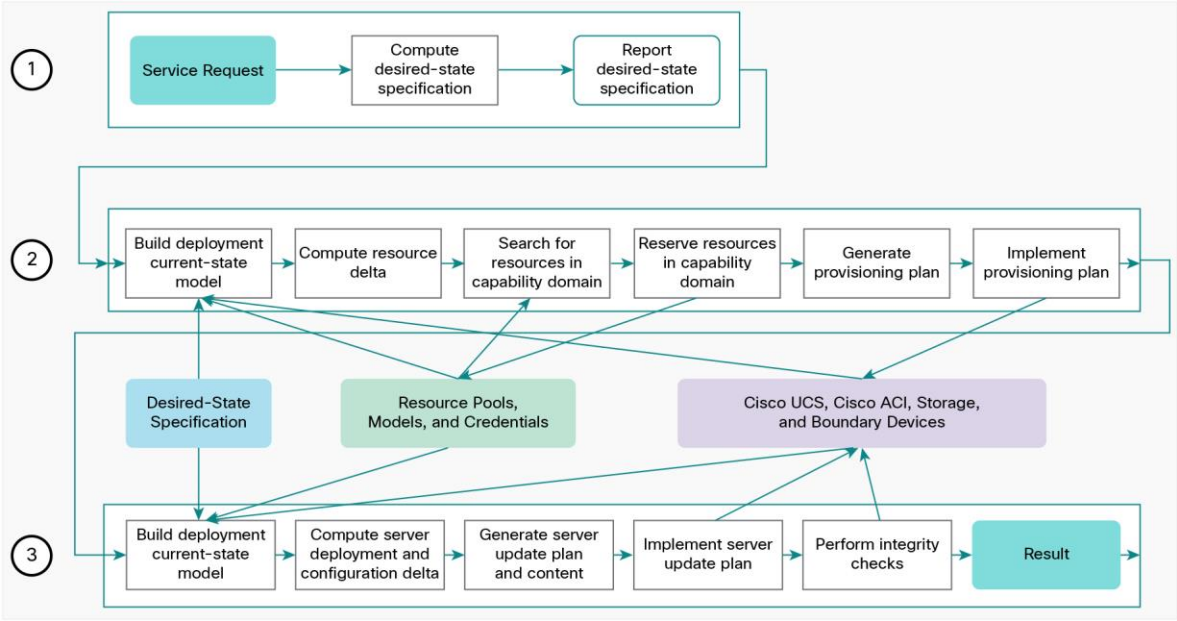

Steps 2 and 3 can be repeated periodically to check compliance with the DSS and remediate any deltas.

Provisioning is separated from deployment to help ensure that the necessary resources are available or removed before application component deployment operations are performed.

Undeployment and deprovisioning steps are not shown; these processes occur in the reverse order.

Deployment Service Requests

Application models are published in service catalogs and are accessible to users with the appropriate entitlements. Users can deploy them using a user interface or API (Figure 11).

Figure 11. Application View in Service Catalog

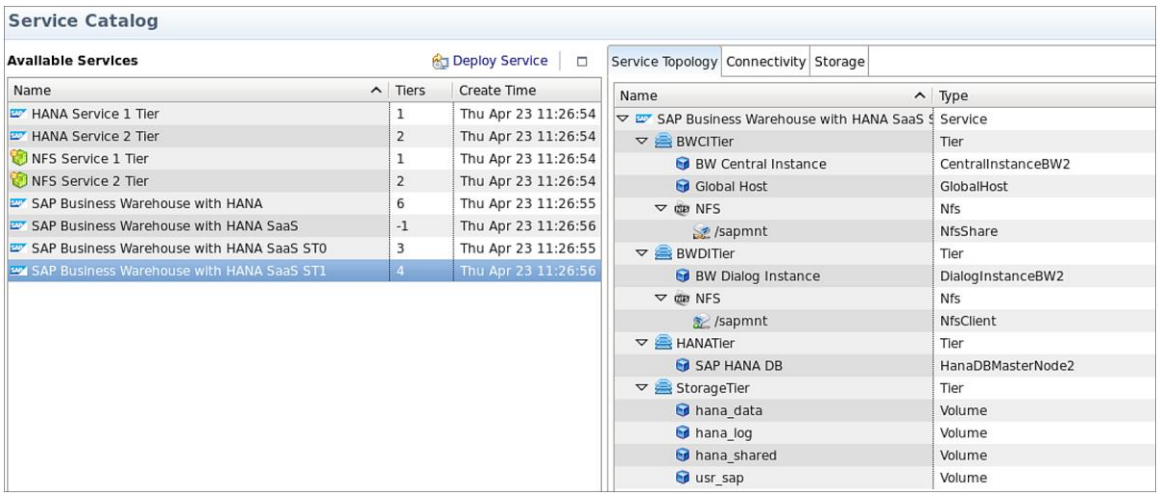

An application is deployed or updated for a tenant by issuing a service request to the VDSC using its user interface or API. Service requests typically contain the most basic information needed to instantiate a service so that users are not burdened with the need to provide low-level information. The other details (the actual topology, server types, storage, connectivity, etc.) are computed by the automation process using the application and infrastructure models. Figure 12 shows a form in the VSDC user interface that exposes parameters for the service request for an SAP BW on SAP HANA deployment.

Figure 12. Vnomic Desired State Controller User Interface

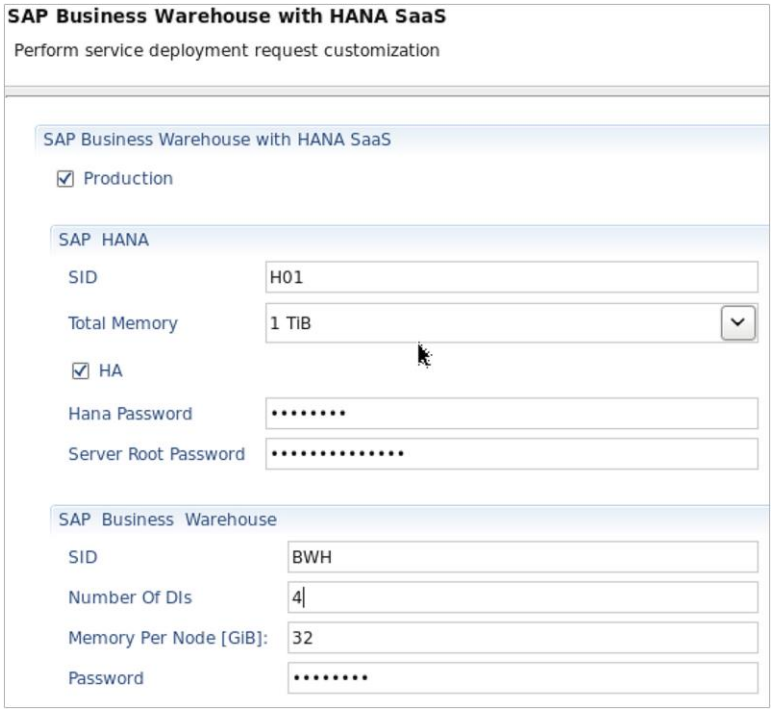

Computing the Complete Deployment Topology

Using the application model and the target infrastructure specified in the service request, the VDSC computes the complete application topology that will be deployed. This topology includes the servers and their hosted application components and interconnectivity (Figure 13).

Figure 13. SAP Landscape View of Application Topology

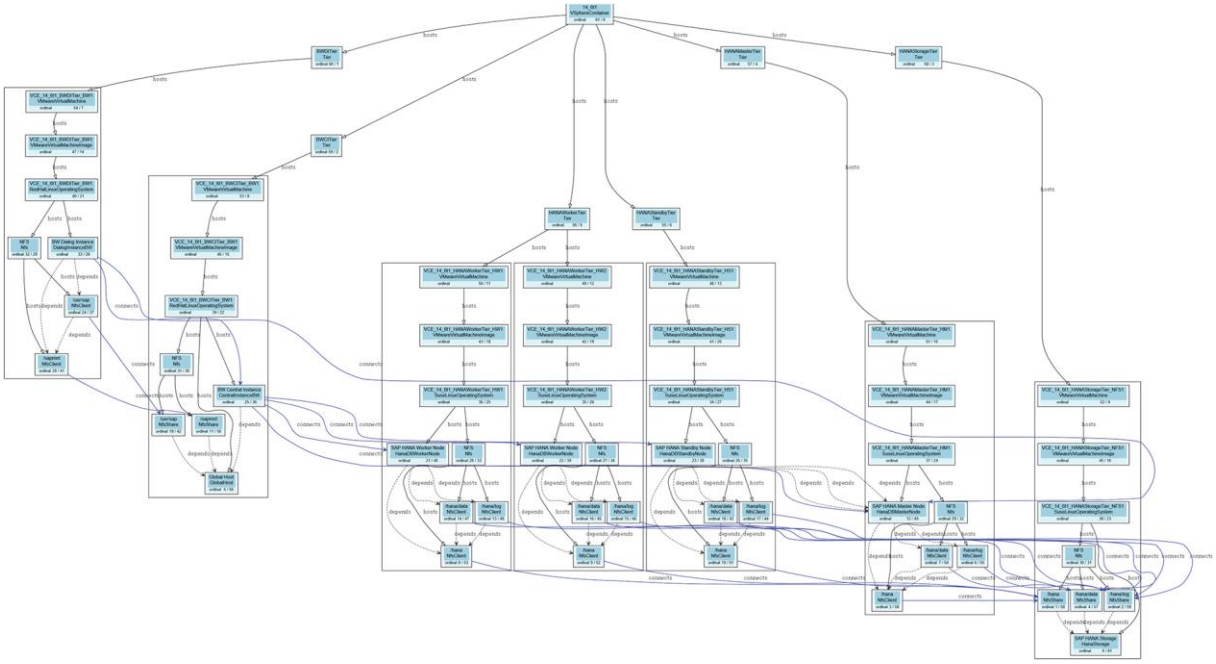

Computing Cisco ACI Policies and Converged Infrastructure Orchestration

Application models contain all the information needed to determine the computing, networking, and storage resources that must be provisioned for an application deployment. This section describes how the Cisco ACI policies and storage orchestration processes are computed.

Computing Cisco ACI Policies

The VDSC automatically determines the logical connectivity required for the application. Each set of application components that must intercommunicate is identified, and components are grouped into their own logical network. The resulting set of logical networks is referred to as the logical network model of the application (Figure 14).

Figure 14. Logical Network Model of the Application

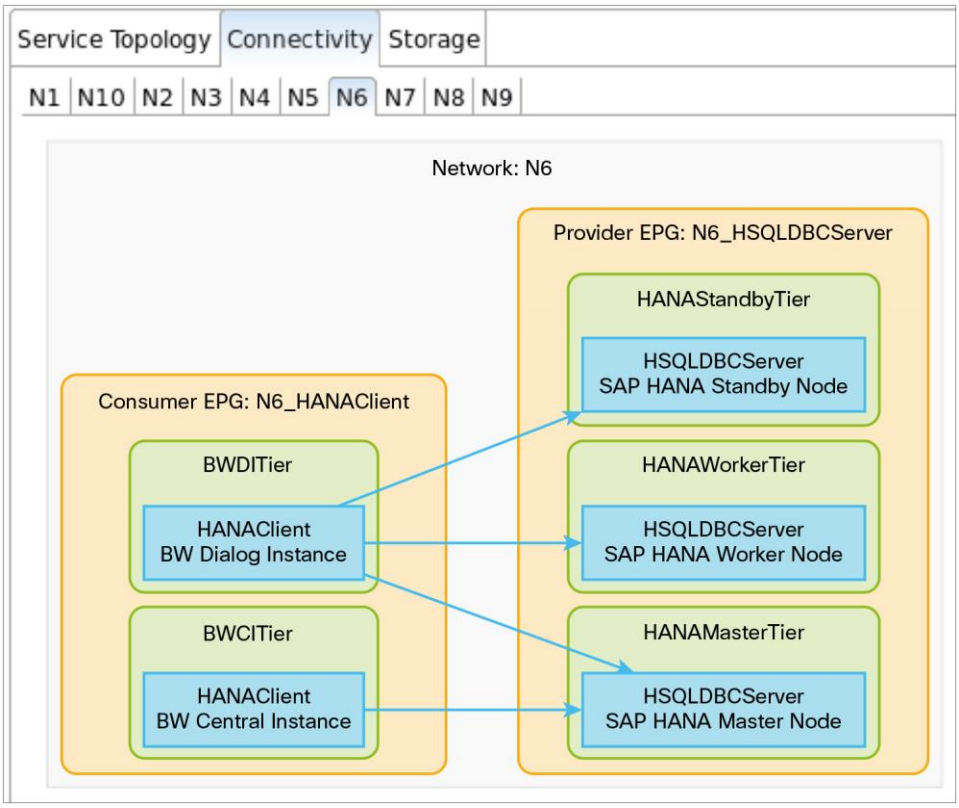

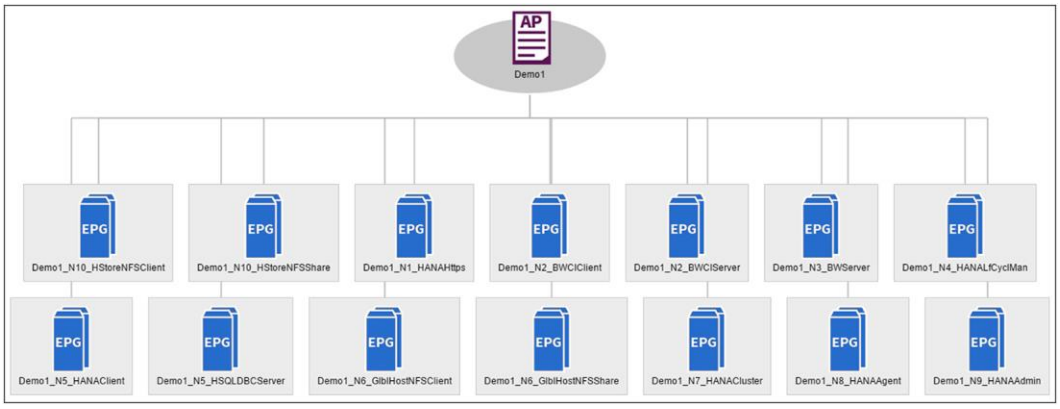

The VDSC translates the logical network model to an APIC application network profile for the target tenant. Figure 15 shows a simplified example of the steps involved. The application requirements for the connections between the endpoint groups are easily defined using Cisco ACI contracts. Applications can now consume the resources, virtual or physical, that best meet their needs, independent of location. Because all application characteristics are fully specified, compliance is assured across any changes in an agile environment. These important application-centric features are what make Cisco ACI unique.

Figure 15. Translating a Logical Network Model to a Cisco Application Network Profile

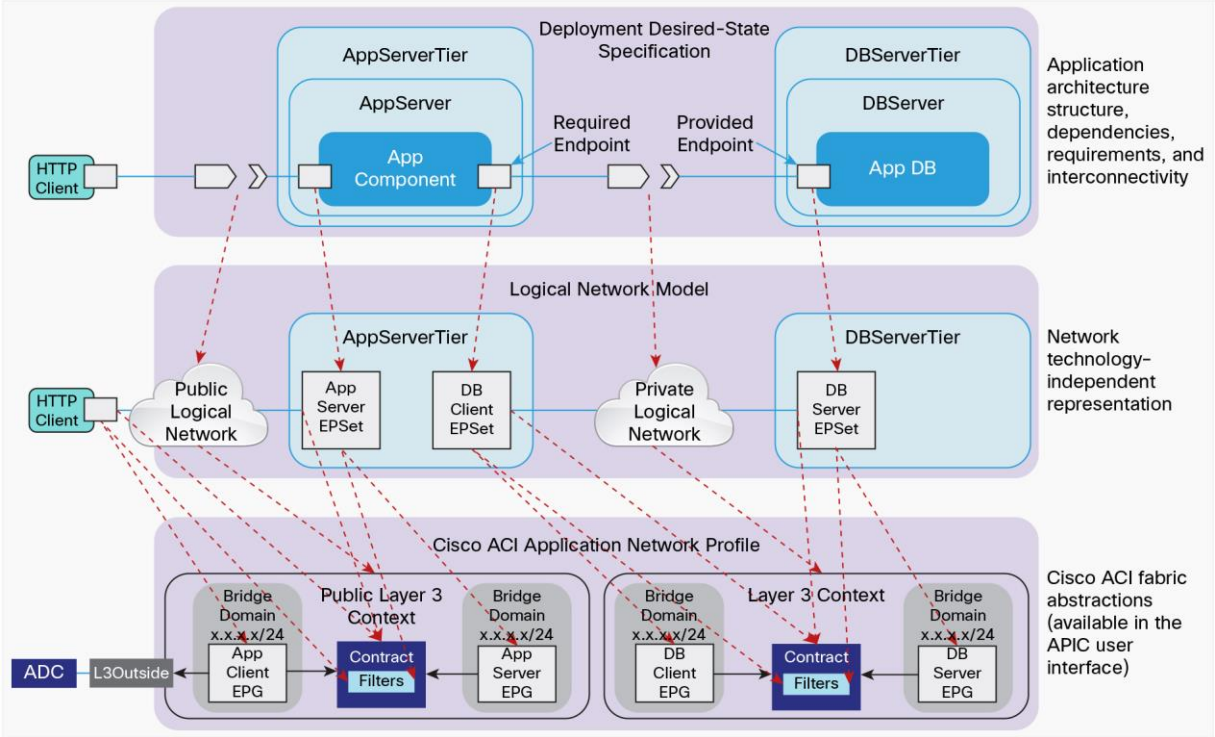

Computing the Policies and Orchestration for Servers

The application computing requirements are used to compute policies and orchestration for Cisco UCS Manager and VMware vSphere (Figure 16).

Figure 16. Computing Policies for UCS

Name	MAC Address	Desired Order	Actual Order	Fabric ID
vNIC AdminServer	00:25:B5:05:00:28	5	5	B A
vNIC Cluster	00:25:B5:05:00:38	4	4	A B
vNIC HANAClient_HSQLD	00:25:B5:05:00:48	3	1	A B
vNIC HttpServer	00:25:B5:05:00:78	2	2	B A
vNIC Vn_Mgmt	00:25:B5:00:00:05	2	3	B A
vNIC hana_dataClient	00:25:B5:05:00:18	7	2	B A
vNIC hana_logClient	00:25:B5:05:00:58	8	3	B A
vNIC hana_sharedClien	00:25:B5:05:00:68	6	1	A B
vNIC pxe	00:25:B5:01:01:09	1	1	A B
vNIC usr_sapClient	00:25:B5:05:00:08	9	2	B A

For bare-metal servers, the VDSC selects the appropriate server, instructs Cisco UCS Manager to instantiate service profiles with the necessary network and storage policies, and orchestrates the operating system installation for the server.

For virtual servers, vSphere is orchestrated to clone the necessary virtual machine templates and customize them with the necessary network and storage configuration. If more virtual server capacity is required, the VDSC will select and provision additional VMware ESXi hosts in a manner similar to that for bare-metal servers.

Computing the Orchestration for Storage

As for the storage, the VDSC automatically determines the logical storage structure required for the applications. Each set of application components that requires storage is identified and mapped to a set of logical storage abstractions, such as consumer ports, volumes, volume groups, and zones. This set of logical storage abstractions is referred to as the logical storage model of the application (Figure 17).

Figure 17. Logical Storage Model

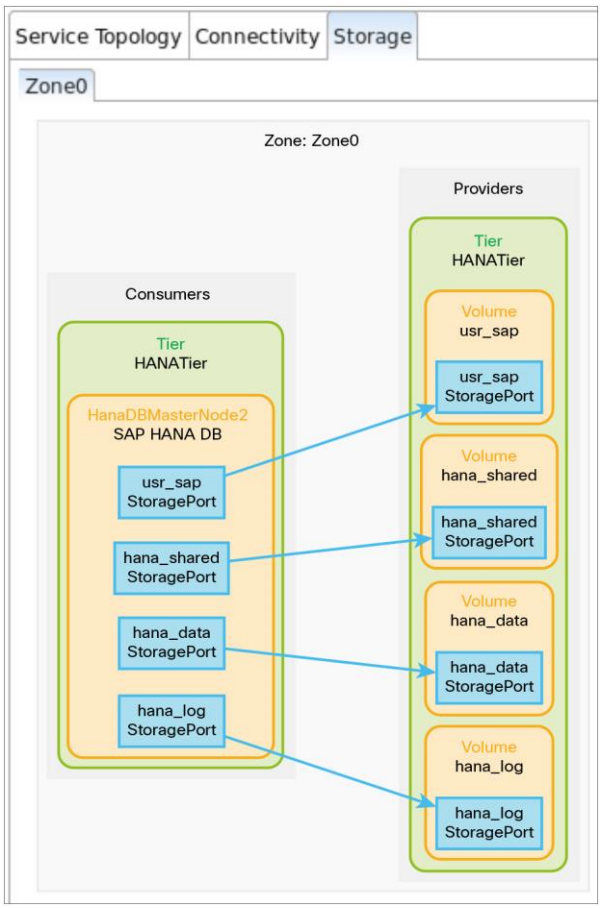

The storage model is used to compute the NetApp FAS orchestration steps that invoke the appropriate clustered Data ONTAP APIs (Figure 18).

Figure 18. Translating the Logical Storage Model to NetApp Clustered Data ONTAP

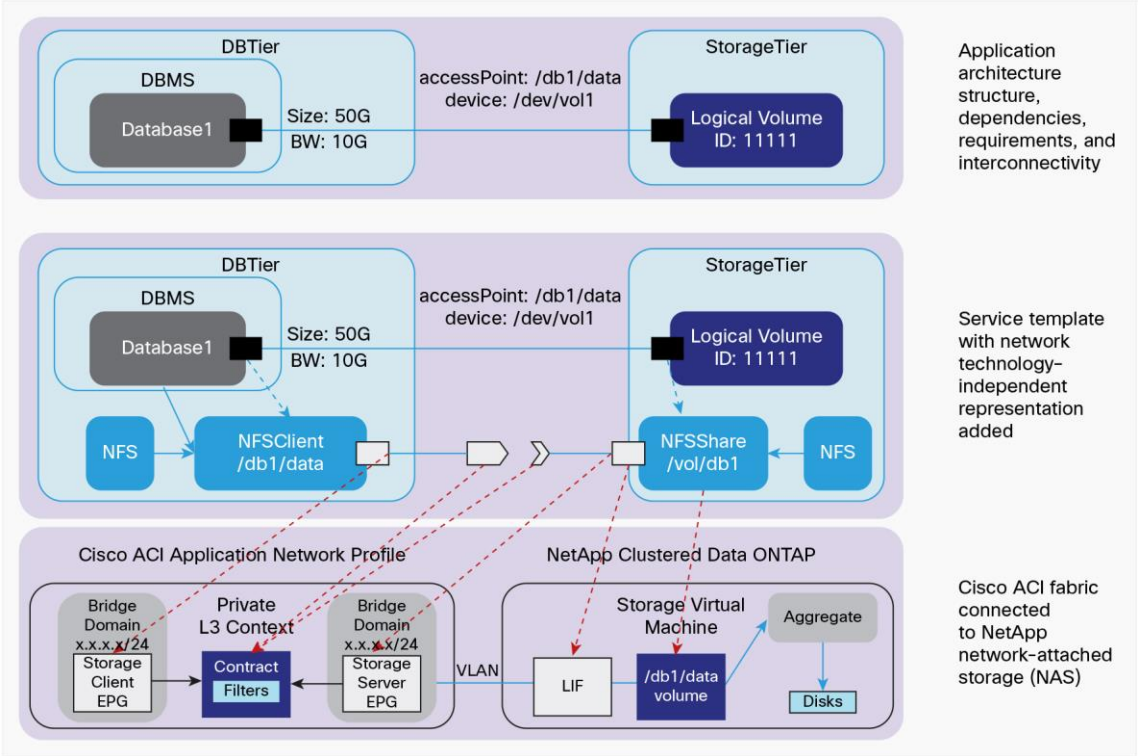

Automated Update

Computing the Provisioning Plan

Using the application and infrastructure models and the current-state model, the VDSC computes a complete provisioning plan that describes the order and state transitions required to orchestrate the provisioning and configuration of the computing, network, and storage infrastructure for the entire deployment.

Implementing the Provisioning Update Plan

Having precomputed the Cisco ACI policies and the computing and storage orchestrations, the VDSC runs its provisioning plan so that the application of Cisco ACI policies and the infrastructure API invocations occur in the appropriate order. Upon successful execution of the provisioning plan, all servers in the deployment will be ready for use for software deployment.

Computing the Software Update Plan

Using the application and infrastructure models and the current-state model, the VDSC computes a complete software update plan that describes the order and state transitions required to orchestrate the deployment and configuration of all the software components in the deployment.

Implementing the Software Update Plan

Having precomputed the required software updates and the software update plan, the VDSC runs the software update plan so that the software APIs and script invocations occur in the appropriate order. Upon successful execution of the software update plan, all servers in the deployment will be ready for use (Figure 19).

Figure 19. VDSC Software Update Plan

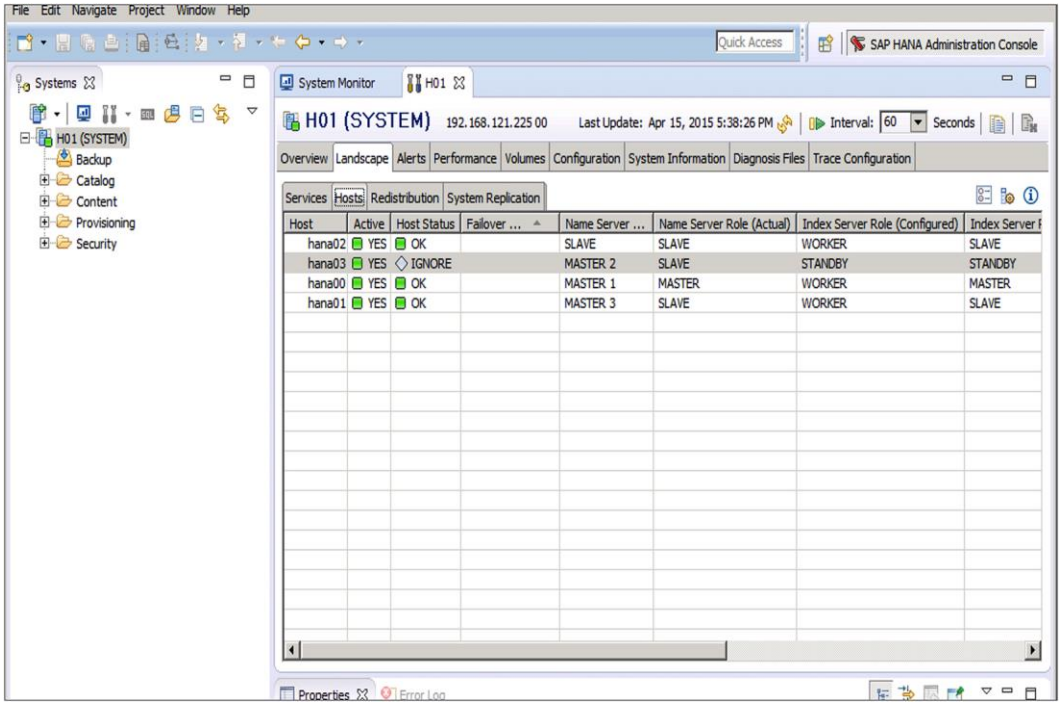

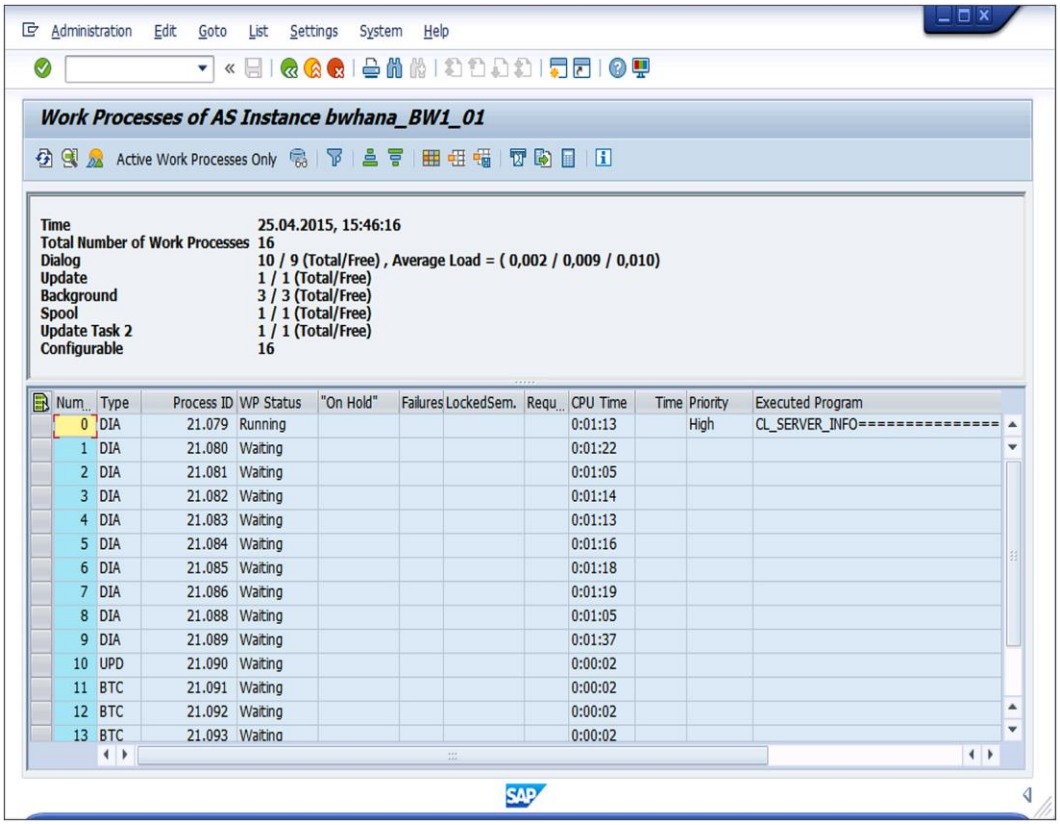

To help ensure compliance with specific key performance indicators (KPIs) or service-level agreements (SLAs), a validation test can be run after the deployment update is complete. For example, for deployments containing SAP HANA nodes, the SAP HANA TDI tests are run to help ensure that the KPIs are in the appropriate ranges.

What Is Automated?

This section describes in detail the provisioning, deployment, and configuration operations that are automated.

Infrastructure Provisioning and Configuration

The following list shows the infrastructure provisioning and configuration details that are automated. The actual automation implemented depends on the models used and the capabilities of the infrastructure.

- Selection of computing and storage resources
- Selection of fabric paths interconnecting the computing and storage nodes of the topology
- Tracking and provisioning of resources by tenant and deployment
- Tenancy with minimal connectivity and data access between components
- Server attachment to network and storage and server boot configuration
- Resource management
  - Cisco ACI network fabric
    - EPGs, Layer 3, Layer 2, contracts, tenancy, isolation, filtering, and QoS
    - Paths across the fabric
  - Computing resources (blade servers)
    - CPU architecture, speed, and memory characteristics
    - NICs, uplinks, VLANs, MAC addresses, and IP address and subnet pools
  - NetApp storage resources (software-defined storage)
    - Volumes and storage virtual machines
    - Ports, VLANs, and logical interfaces (LIFs)
    - Bare-metal servers and preexecution environment (PXE) boot
    - ESXi hosts and Small computer System Interface over IP (iSCSI) boot

Software Deployment and Configuration

The following list shows the software deployment and configuration details that are automated. The actual automation implemented depends on the models used and on the operating systems and boot techniques used.

- Addition, removal, and update of OS services
  - Startup and shutdown configuration
- Operating system configuration
  - Kernel configuration
  - File-system mounts (block and file)
  - Network interfaces and routing
- Artifact installation and staging of installation files
  - Certificates, templates, configuration files, etc.

- Configuration of intercomponent dependencies across the entire topology
  - Endpoints
  - Credentials
- Identity
  - Users and group membership
- Compliance
  - Enforcement of deployment and configuration constraints

SAP HANA

This section describes the notable aspects of the SAP HANA deployment automation.

Table 1 provides an overview of the time required for SAP HANA automated deployments.

Table 1. SAP HANA Automation Time

Software Construct	Provisioning and Deployment Time (Approximate)
SAP HANA: 1 node	• 30 minutes (virtual machine) • 60 minutes (bare metal)
SAP HANA: 4 nodes	90 minutes (bare metal)
SAP HANA: 8 nodes	120 minutes (bare metal)

Note that times are approximate and depend on software versions and the I/O load and throughput of the hosting infrastructure.

SAP HANA scale-up and scale-out topologies are automated (Table 2). High availability is an optional feature, and when it is enabled it adds a standby node to the topology. On FlexPod architecture, a complete instance of NetApp SnapCreator for SAP HANA is included in each SAP HANA deployment with a predefined snapshot schedule.

Table 2. Automating Scale-Out

SAP HANA Size	Nodes
256 GB	1 node, virtual machine
512 GB	1 node, virtual machine
768 GB	1 node, virtual machine
1 TB	1 node, virtual machine or bare metal
2 TB	2 nodes, bare metal
3 TB	3 nodes, bare metal
4 TB	4 nodes, bare metal
6 TB	6 nodes, bare metal
8 TB	8 nodes, bare metal

Each deployment meets its TDI KPIs. It is possible to provision less capacity for nonproduction use, such as for development or quality assurance deployments, but the requirements are not standardized.

Each deployment is a new installation (not a copy) of SAP HANA with the default licensing.

SAP Business Warehouse

This section describes the notable aspects of the SAP BW deployment automation.

Table 3 provides an overview of the time required for SAP BW automated deployment. The performance assumes a capable SAP HANA instance.

Table 3. SAP Business Warehouse Automation Time

Software Construct	Provisioning and Deployment Time (Approximate)
SAP BW	1.5 hours or more, depending on the size of the imported database

Note that times are approximate and depend on software versions and the I/O load and throughput of the hosting infrastructure.

Each deployment is a new installation (not a copy) of SAP BW with the default licensing.

Customer Benefits

By taking advantage of the power of the automated policy-based delivery and governance for SAP BW on SAP HANA using FlexPod Datacenter, enterprises and service providers can deploy on demand complete SaaS instances of SAP HANA and SAP BW on SAP HANA for multiple tenants across Cisco ACI in hours, with no manual errors. This solution effectively removes the traditional semantics gap between application requirements and infrastructure provisioning and configuration by representing the application's software components and hosting infrastructure as a single logical abstraction. This solution results in dramatically faster deployment times and greater operational control of dynamic applications:

- Faster deployment and change times for sophisticated applications
- Elimination of complex, error-prone processes
- Alignment of infrastructure behavior with application requirements
- Continuous assurance of security and compliance
- Simplified processes to add capacity and reduced lead times

Conclusion

Automated policy-based delivery and governance for SAP BW on SAP HANA using FlexPod Datacenter offers the first highly scalable application-centric platform that is fully prevalidated and can efficiently instantiate, update, and decommission sophisticated SaaS deployments of SAP HANA and SAP BW on SAP HANA landscapes quickly, securely, and risk free. Until now, intervention from software and infrastructure experts was required for common application lifecycle events, along with dedicated and costly infrastructure with complex manual configuration processes.

The application-centric model achieved with Cisco ACI eliminates manual errors that are responsible for more than 80 percent of application and infrastructure unplanned downtime and security breaches, helping companies avoid hundreds of thousands of dollars per hour in lost revenue. Even more important, this solution helps ensure that the company brand is not tarnished and the company is not perceived as an unreliable provider of online business services. The result is faster time to value (TTV) and reduced TCO, with ensured QoS, security, governance, and auditability.

For More Information

http://www.cisco.com/go/aci

C11-734753-01 10/15