

TOMORROW starts here.



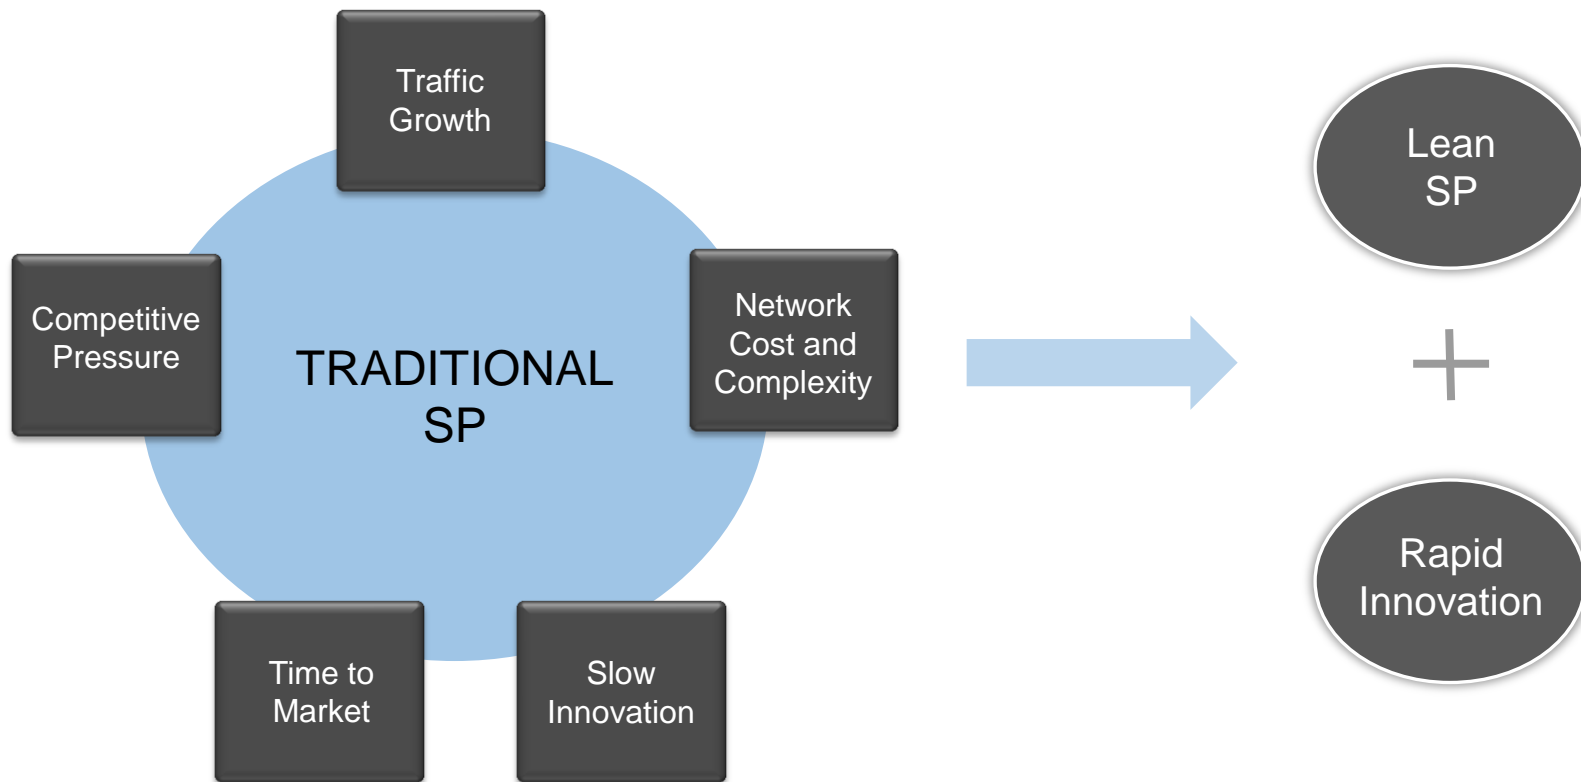
Cisco *live!*

Cloud Enablement Architecture and NFV Services Delivery

BRKSPG-3864

Rex Fernando
Lead Architect
Distinguished Engineer, CAO

Key SP Challenges and Path Forward




Transformation To Carrier as a Service

Traditional SP



Lean SP



Evolved Services Platform

Evolved Services Platform

... an open, standards-based, modular architecture and platform for services orchestration

... manages the physical & virtual network, as well as the compute & storage infrastructure to deliver carrier-class services

... which range from VPC to NFV services

ESP - End to End Architecture for Service Orchestration

Evolved Services Platform (ESP)

Service Catalog
“Business Intent”
catalogs

Orchestration Engine
“Execution”
configuration,
workflow,
automation,
provisioning

Service Catalog

Routing / VPN

Transport

Security

Virtual Private
Cloud

Mobility

Video/ Content

Managed
Services

Cross Domain Orchestration

Service Provisioning (Physical & Virtual Services)

- Provision WAN services
- Provision NFV in DC

Network Orchestration

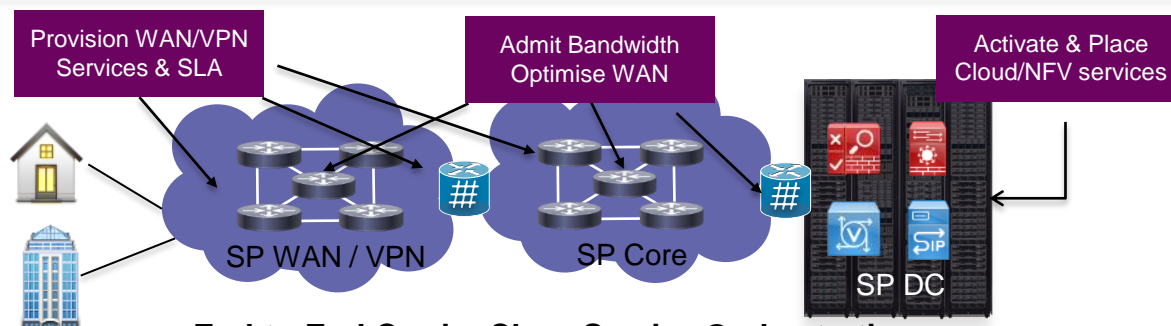
- DC SDN
- WAN Orchestration & Optimisation

Compute & Storage Control

- Elastic Services Control
- Service Lifecycle management

Carrier Class Reliability
and High Availability

Physical & Virtual
Network, Compute &
Storage



End-to-End Carrier Class Service Orchestration

ESP – Evolved Services Platform

“A flexible multi-tenanted cloud services orchestration platform for the virtualised data centre”

NETWORK AND APPLICATION CONNECTIVITY MODELS

WHAT CAN BE VIRTUALISED?

ENTERPRISE APPS

3-tier Apps, Web Servers, DB Servers, Hadoop Clusters, Distributed Storage

TRANSIT NFV

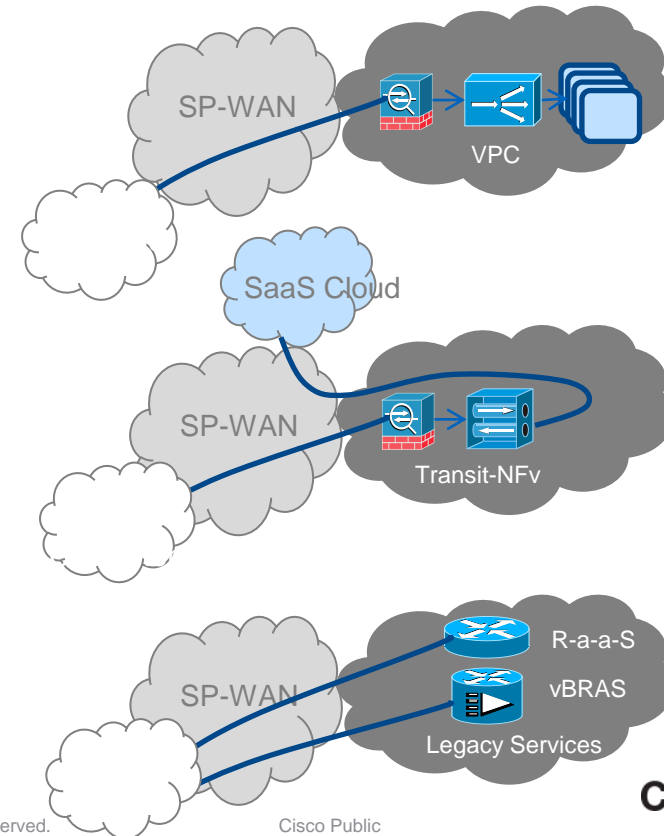
DPI, Firewall, NAT, Load Balancers, WaaS, GI-LAN Applications

TERMINATE NFV

IPSec Gateways, SSL VPNs, vEPC Applications

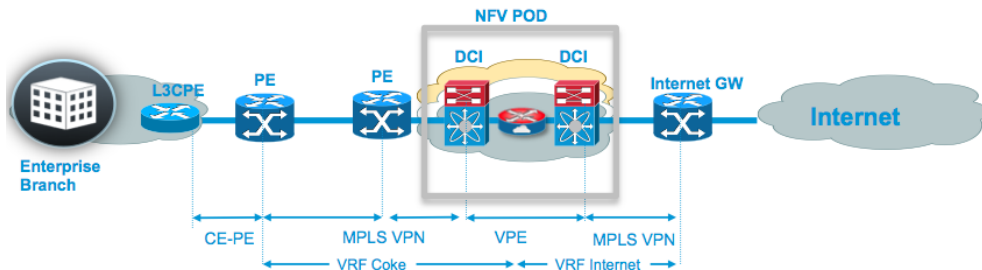
NETWORK SERVICES

DNS, Routing, BRAS, NTP




Transit NFv Examples

NFV – Internet FW – 1A



- Provide internet connectivity for VPN customers and apply NAT and Firewall policies per customer.
- VNF = CSR per customer VRF instance

NFV – Remote Access – 1B




- Provide remote branch of an enterprise with ability to access headquarters over a secure tunnel using IPsec
- Map IPsec tunnel to a enterprise VRF
- VNF = CSR per customer VRF instance

- Provide connectivity between 2 different enterprise VPNs. Apply firewall policies and translate addresses.

- VNF = CSR per customer VRF instance

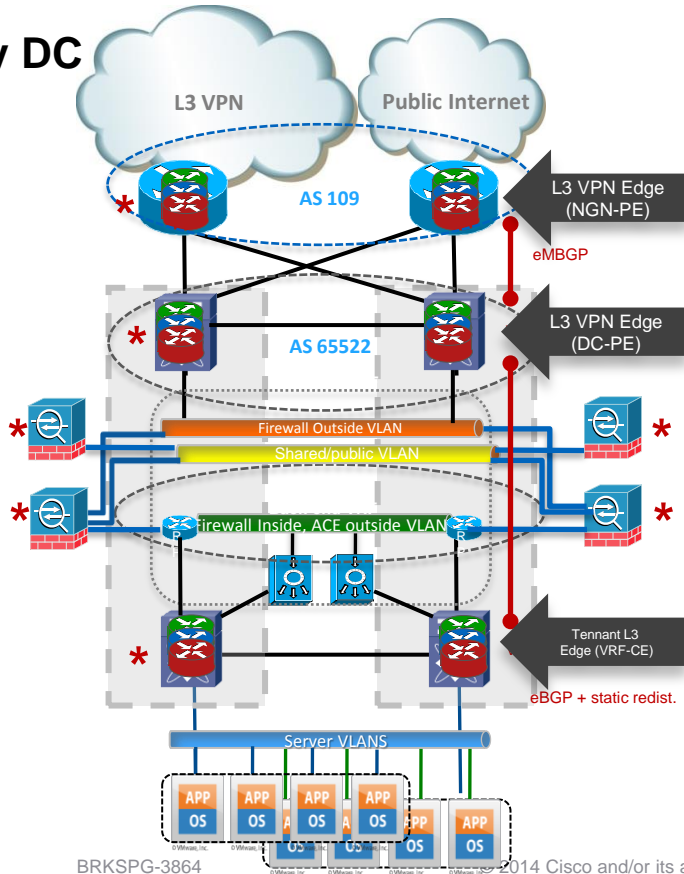
- Use case based on many tier 1 SP customer requirements

NFV – Inter-VPN Firewall – 1C

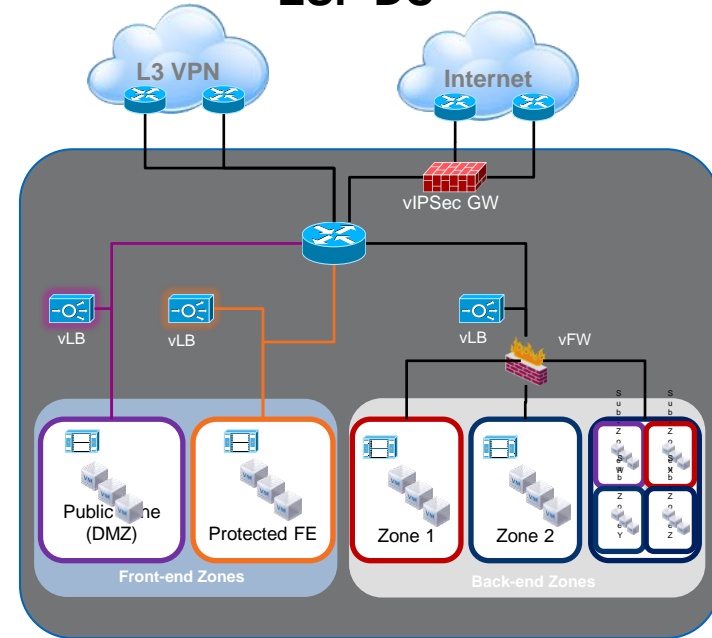


Data Centre Evolution

Legacy DC



ESP DC



Virtualised Compute and DC overlay

Agility (Create/Delete), Scale, Flexible Topologies, BYOD, Elasticity, Utility Based Pricing



ESP – Automated Cloud Services Delivery

TRANSFORMATION



Service provisioning
from days to minutes

From Cabling to Service Chaining

Simple Logistics & Common
Sparing

Dynamic & Elastic Scale


Seamless Integration with IP NGN



What are the Use Cases for SP Virtualisation

Virtualisation of SP Infrastructure

Virtualisation of foundation SP infra such as routing and mobility packet core.

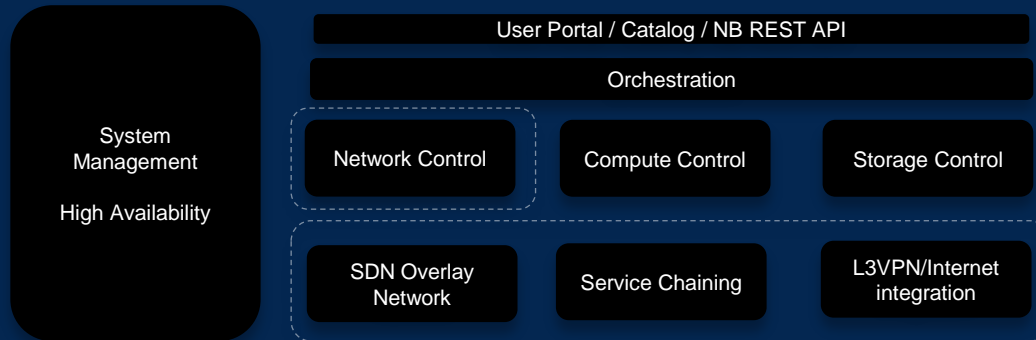


NFV for Enterprise Managed Services

Virtualisation of Network Services that can be delivered as managed services for enterprise



SP Cloud Services Platform



Cloud and Data Centre Requirements

Scale

- Data centres of varying sizes
- Large number of servers/VMs
- Multi-tenancy
- High bisectional bandwidth within DC

Services

- Network Virtualisation, instant Insertion of network services
- Service Chaining, Services networking
- Robust network availability and redundancy
- Seamless integration with WAN, DCI

Flexibility

- DC Underlay network agnostic
- Add network capacity and load incrementally
- Workload and VM mobility
- Variety of server, access connectivity options, multi-homing

Manageability

- Network orchestration and operations at scale
- Simplified network, service provisioning for tenants
- Ease of data collection and troubleshooting
- Support for OAM and proactive monitoring

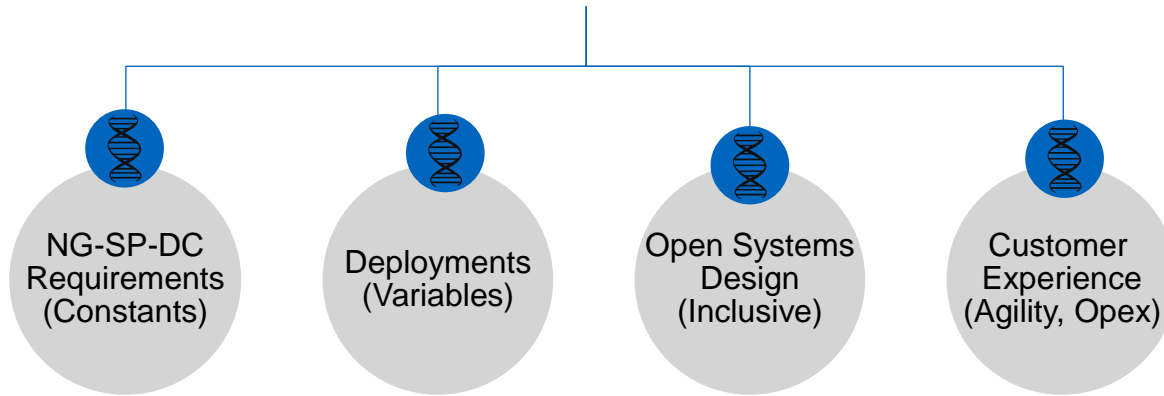
Openness

- | | | |
|-------------------------|-----------------------|--------------------|
| • Yang Models | MPLS-over-UDP, L2TPv3 | • KVM |
| • REST, RESTConf | • OVF, VMDK | • Ganglia |
| • BGP | • Linux/Ubuntu | • Puppet & Cobbler |
| • MPLS-over-GRE, VXLAN, | • Openstack | |

Architectural Goals



ESP Architecture




Architectural Goals



ESP Architecture


NG-SP-DC
Requirements
(Constants)


Deployments
(Variables)


Open Systems
Design
(Inclusive)


Customer
Experience
(Agility, Opex)

**LARGE NUMBER OF SERVERS
AND VMs**
MULTI-TENANCY
HIGH BISECTIONAL BANDWIDTH
**OPTIMAL L2 and L3
FORWARDING**

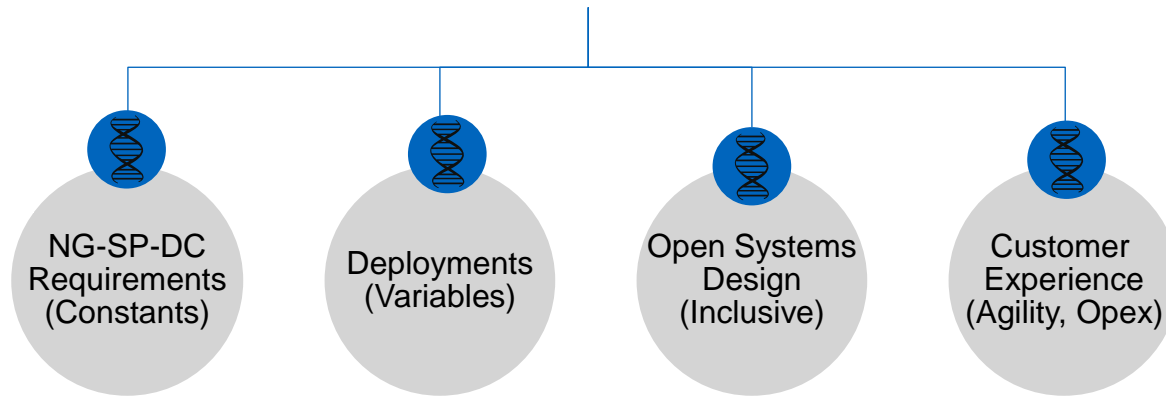
**NETWORK
VIRTUALISATION**
SLA ASSURANCE
**AVAILABILITY AND
RELIABILITY**
**SEAMLESS DCI
INTEGRATION**
NETWORK SERVICES

ELASTIC SCALING
WORKLOAD MOBILITY
**MULTIPLE CONNECTIVITY
OPTIONS**
**PHYSICAL DEVICE
INTEGRATION**

Architectural Goals



ESP Architecture



THIRD PARTY SWITCHES, SERVERS AND DCI
DIFFERENT PHYSICAL UNDERLAY CONNECTIVITY
OPENSTACK & VMWARE SUPPORT
CEPH & NETAPP SUPPORT
BARE METAL INTEGRATION

3RD PARTY NFVS
L2 VPN, L3VPN, INTERNET FOR WAN ACCESS
ENCAPSULATIONS (VXLAN, MPLS-O-GRE, L2TPV3)
SERVICE TOPOLOGIES
TENANT APPLICATIONS


IDENTITY MANAGEMENT
ADMINISTRATIVE SEPARATION
OSS DIFFERENCES
BROWN FIELD DEPLOYMENTS

Architectural Goals



ESP Architecture


NG-SP-DC
Requirements
(Constants)


Deployments
(Variables)


Open Systems
Design
(Inclusive)


Customer
Experience
(Agility, Opex)

YANG MODELS
REST, HTTP, RESTCONF
BGP, IGP
ETHERNET/IP
MPLS-OVER-GRE, VXLAN, MPLS-OVER-
UDP, L2TPV3
L3VPN & L2VPN INTEGRATION
OVF, QCOW2, VMDK

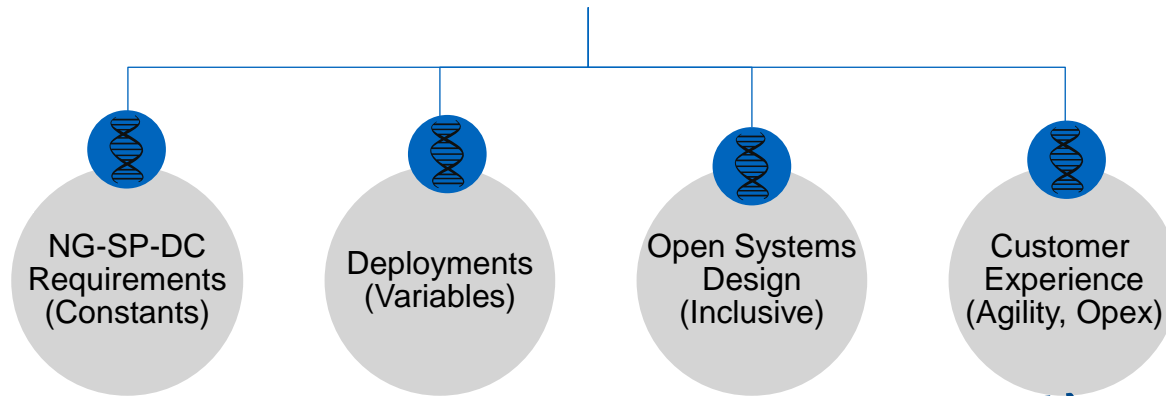
LINUX/UBUNTU
OPENSTACK/KVM
NAGIOS
MYSQL
GANGLIA
PUPPET & COBBLER

JUNIPER AND ALU DCI
PLUGGABLE DHCP AND DNS
AGENTLESS NFV
3RD PARTY SERVERS AND NICS

Architectural Goals



ESP Architecture

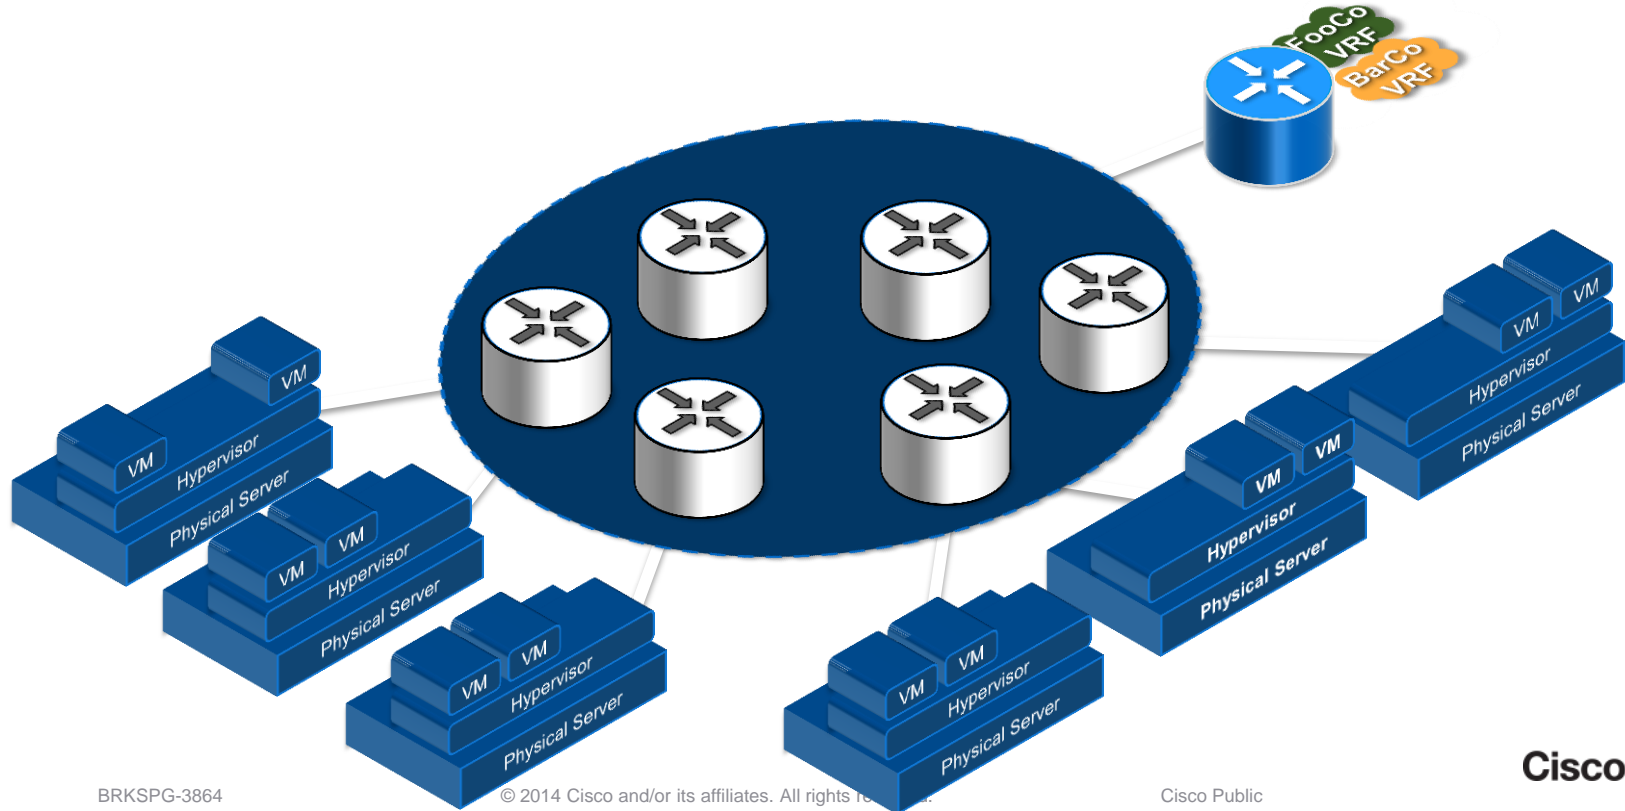


AUTOMATED INSTALLER
AUTOMATED UPDATES
ISSU
SCALE OUT
MODULAR POD
APPLICATION CENTRIC PODS

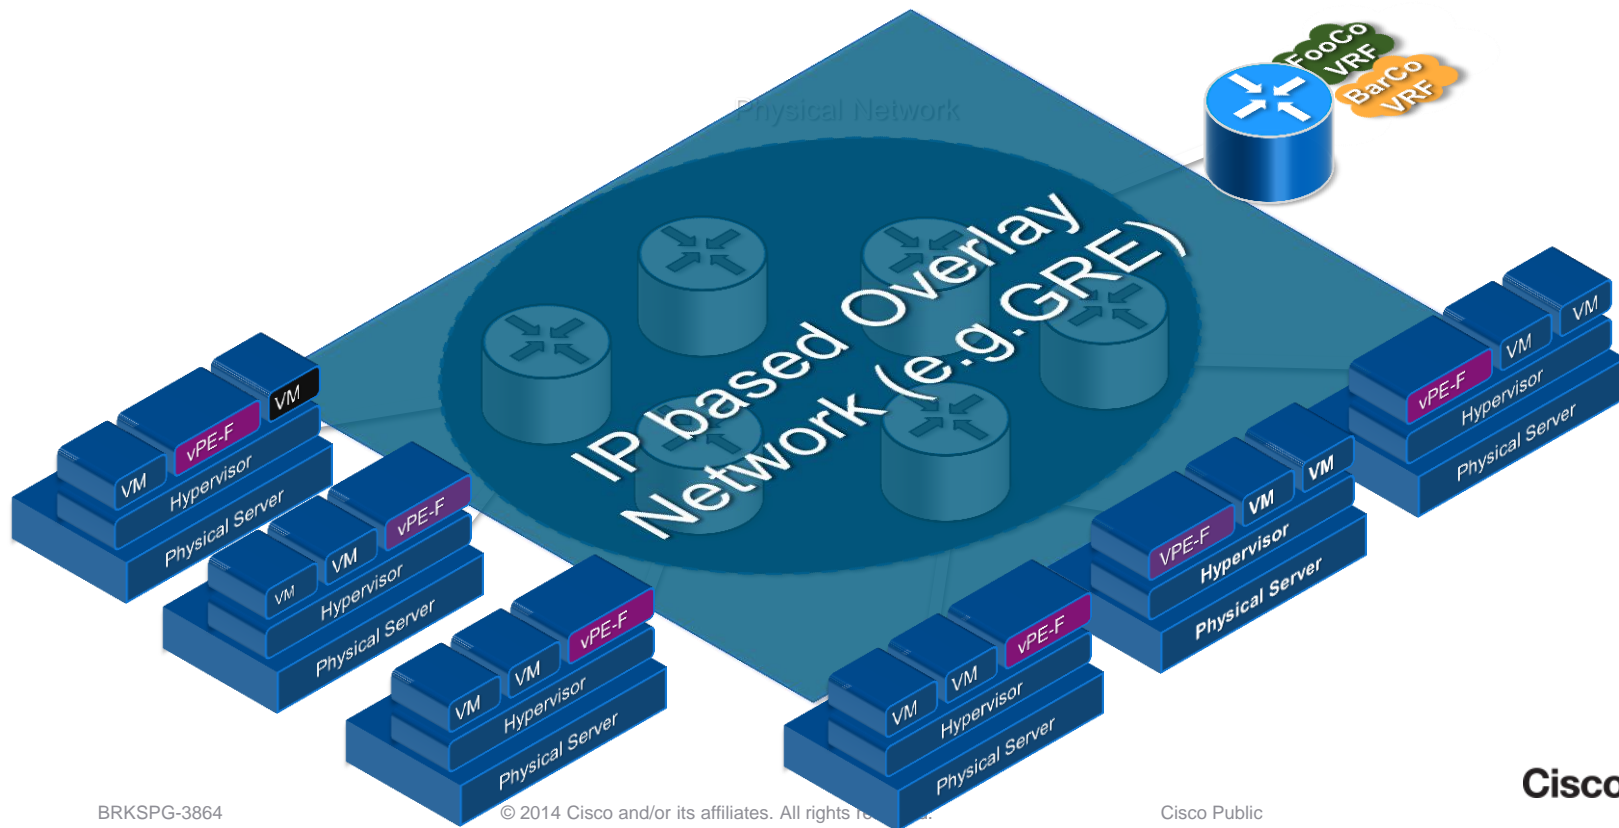
ZERO TOUCH PROVISIONING OF SERVERS AND DCI
CENTRAL TRACING OF SYSTEM EVENTS
APPLICATION ORIENTED SERVICE PROVISIONING

CONTROL HA
PHYSICAL NODE HA
DCI HA
GEO-REDUNDANCY FOR SERVICE TOPOLOGIES
AUTOMATIC RESTART OF FAILED PROCESSES
OVERLAY NETWORK OAM

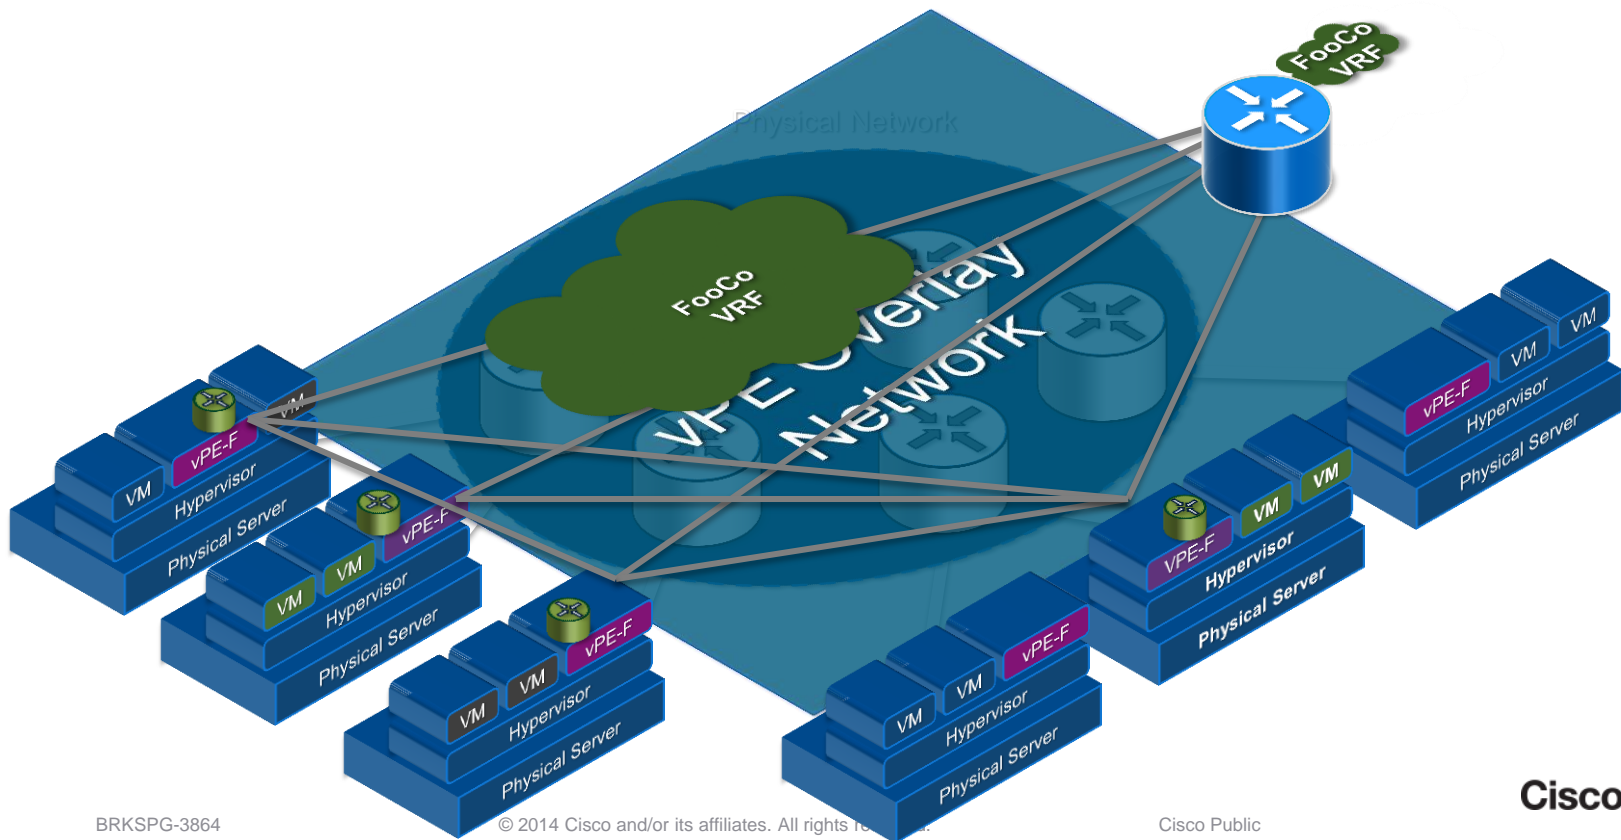
The Data Centre Infrastructure



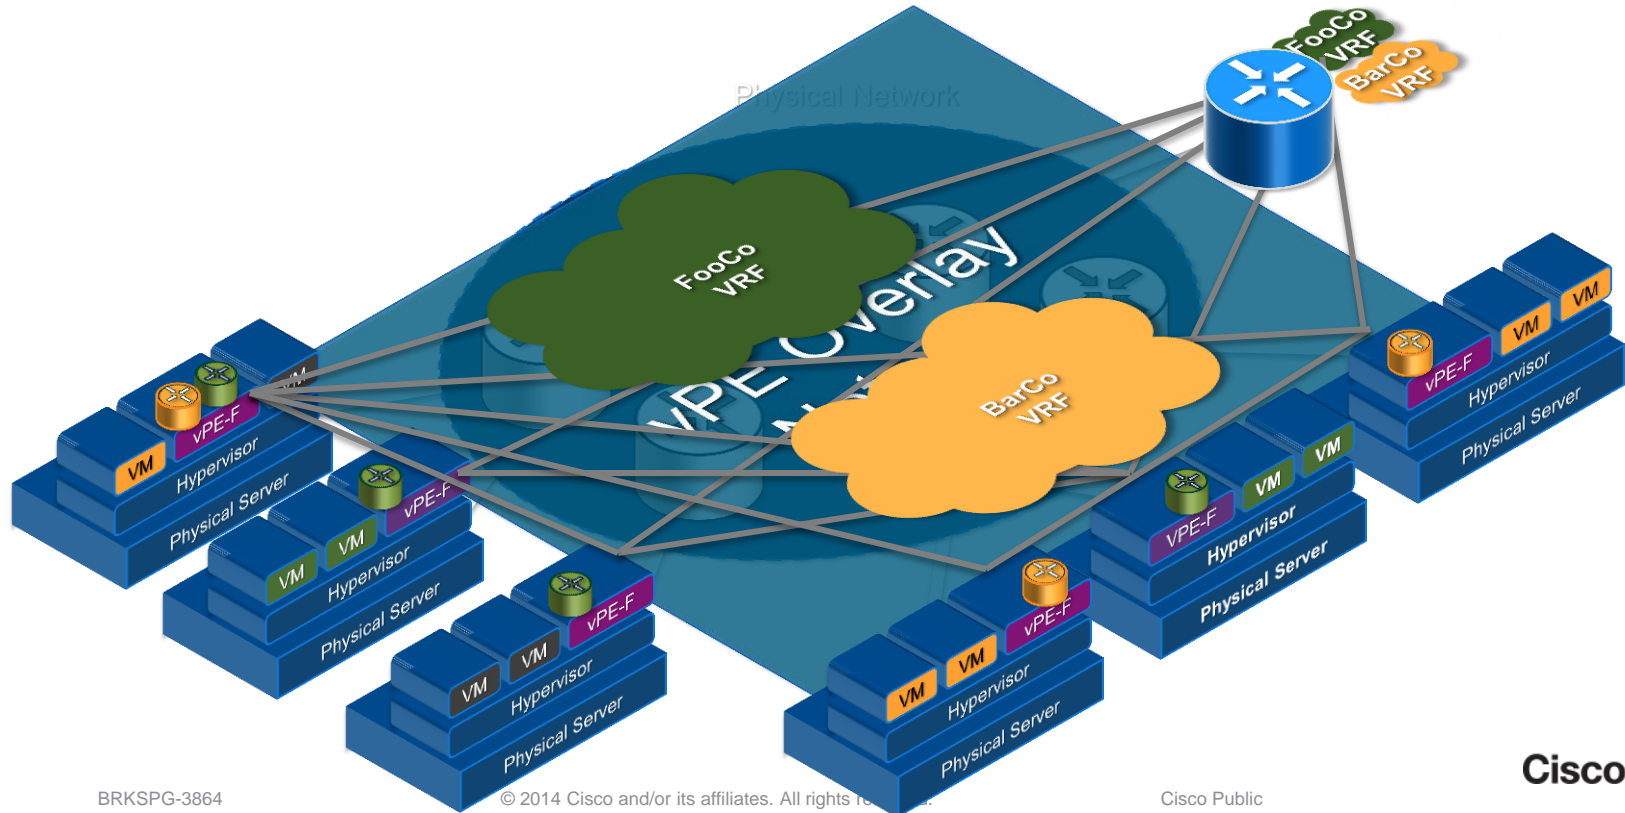
Building an Overlay



Connecting VMs to VPNs

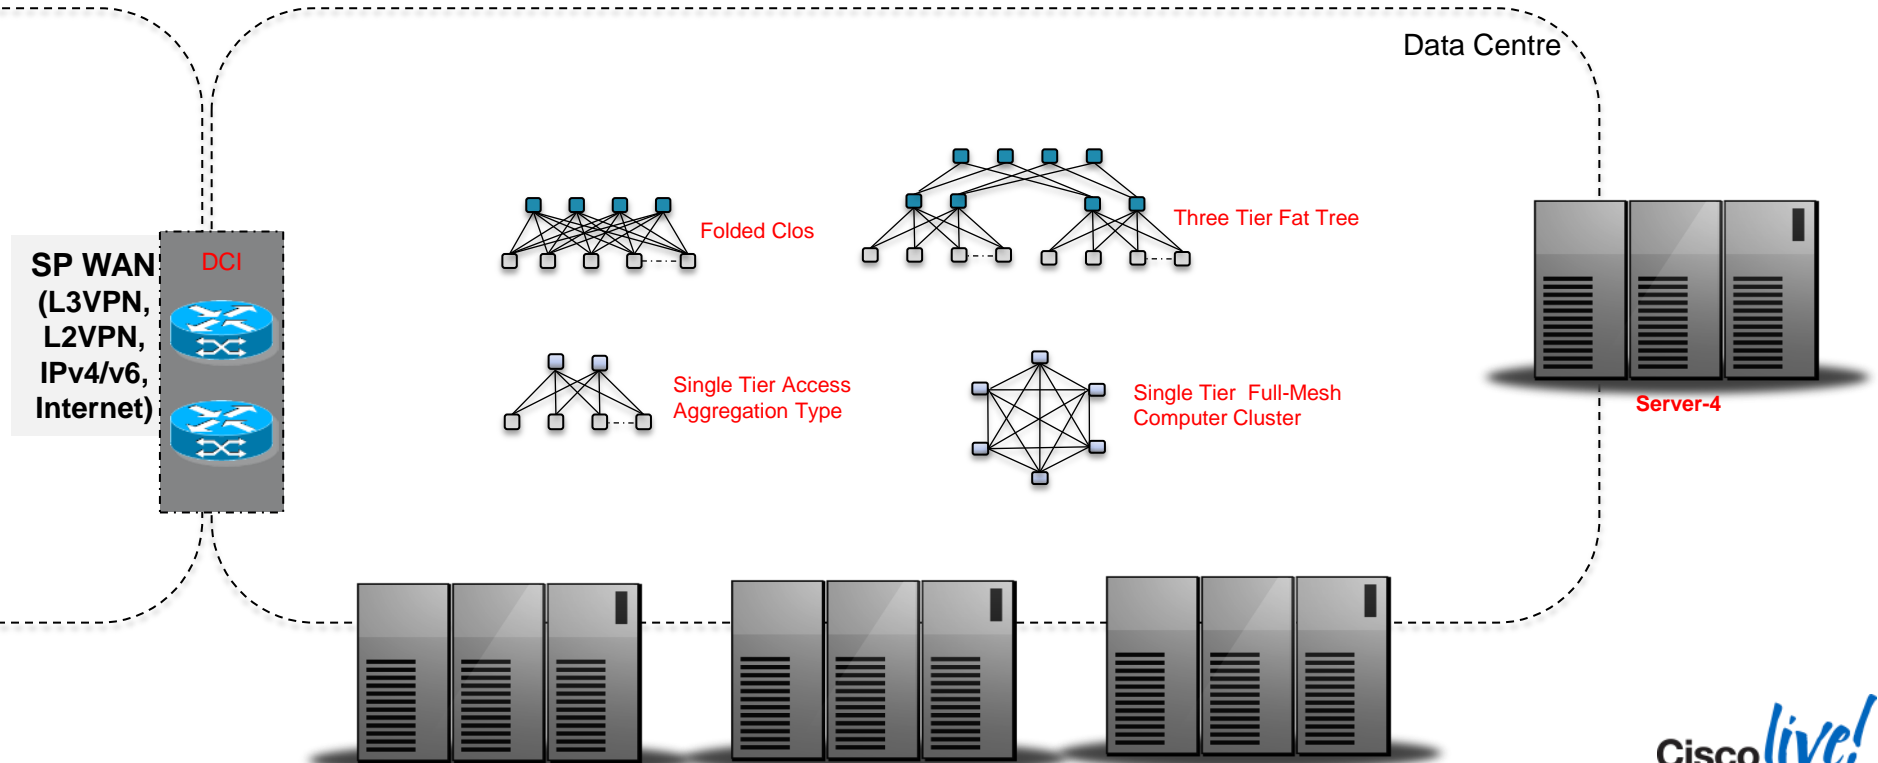


Connecting VMs to VPNs



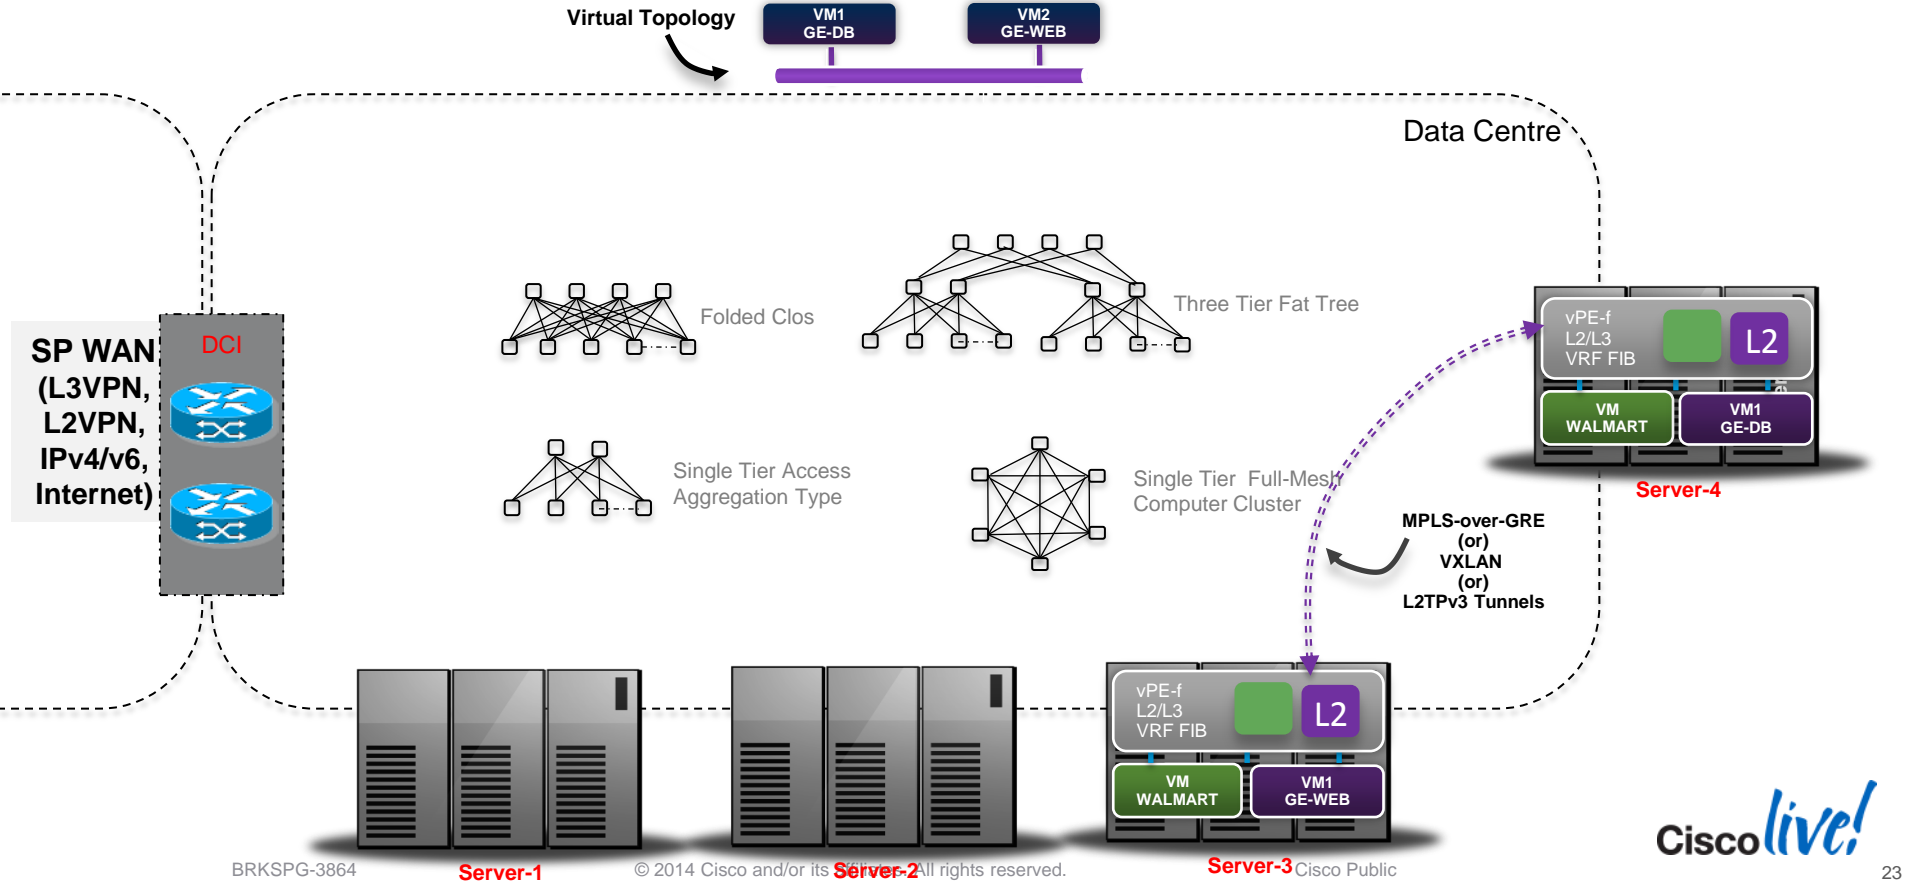
Data Centre Fabric – The Underlay Network

- Many Options for building the underlay
- Provides Fast Reliable Network Connectivity
- Should support P2P and P2MP Capabilities
- Hardware optimised for cost and efficiency



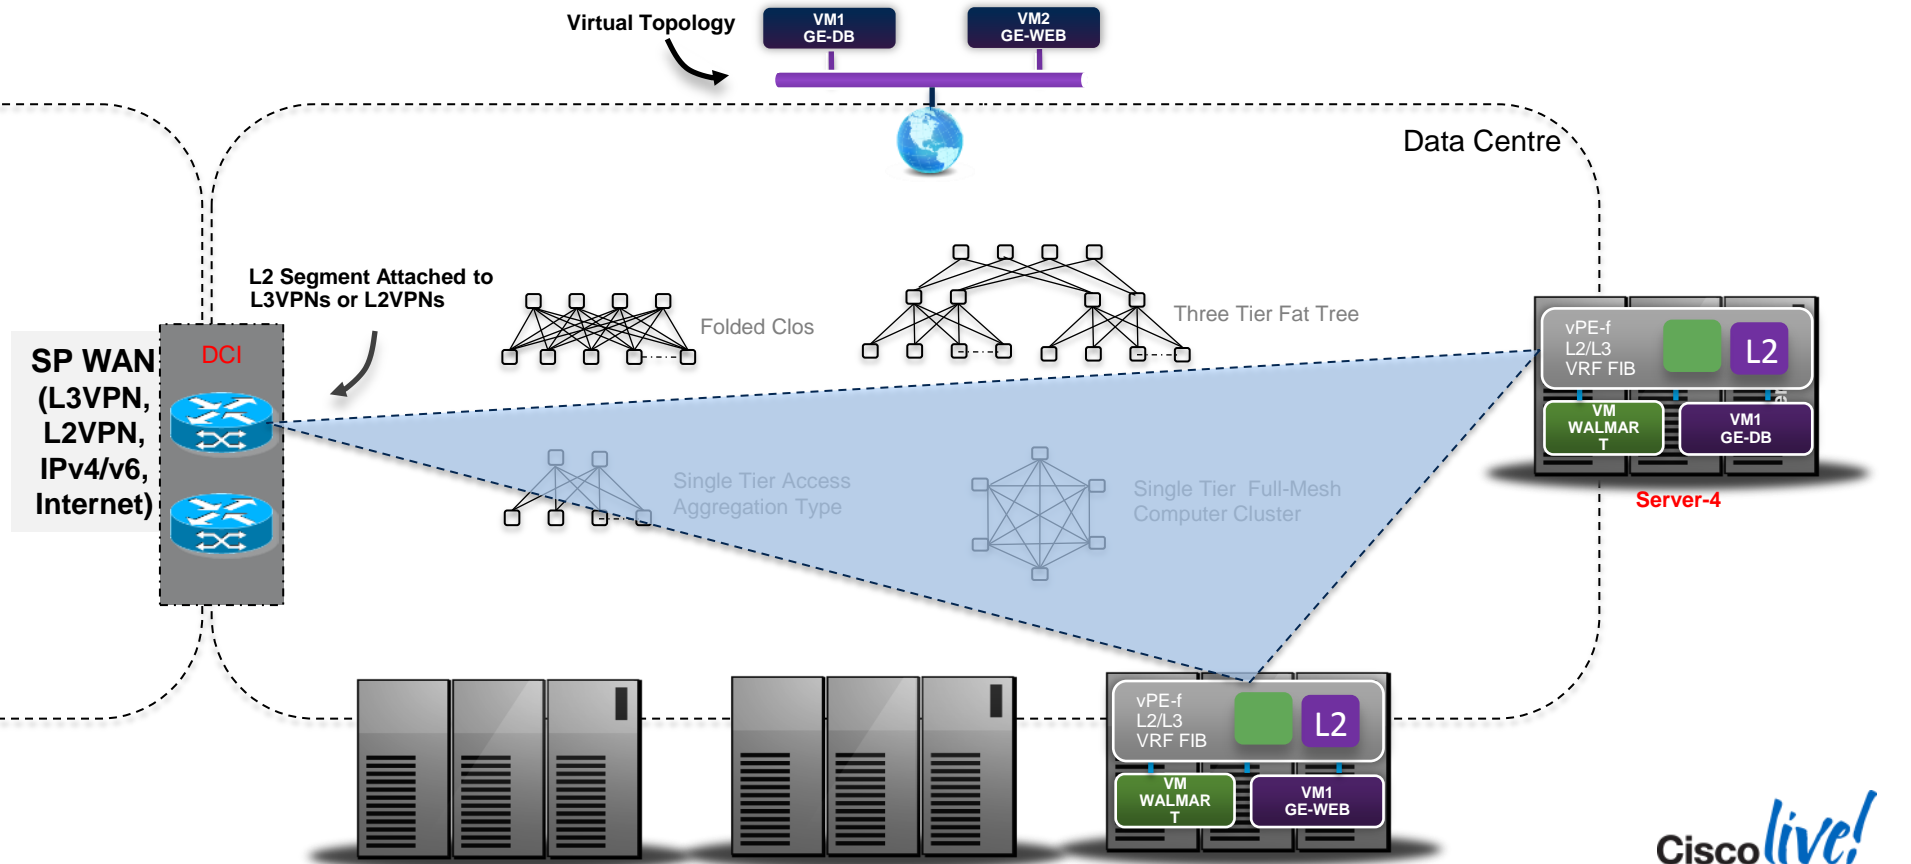
L2 Segment and Forwarding

- Each vPE-f has VRF L2 tables
- vPE-f populated with MAC entries
- VMs see each other in an L2 segment
- MT traffic encapsulated in single transport tunnel
- Only a small class of applications need strict L2 connectivity



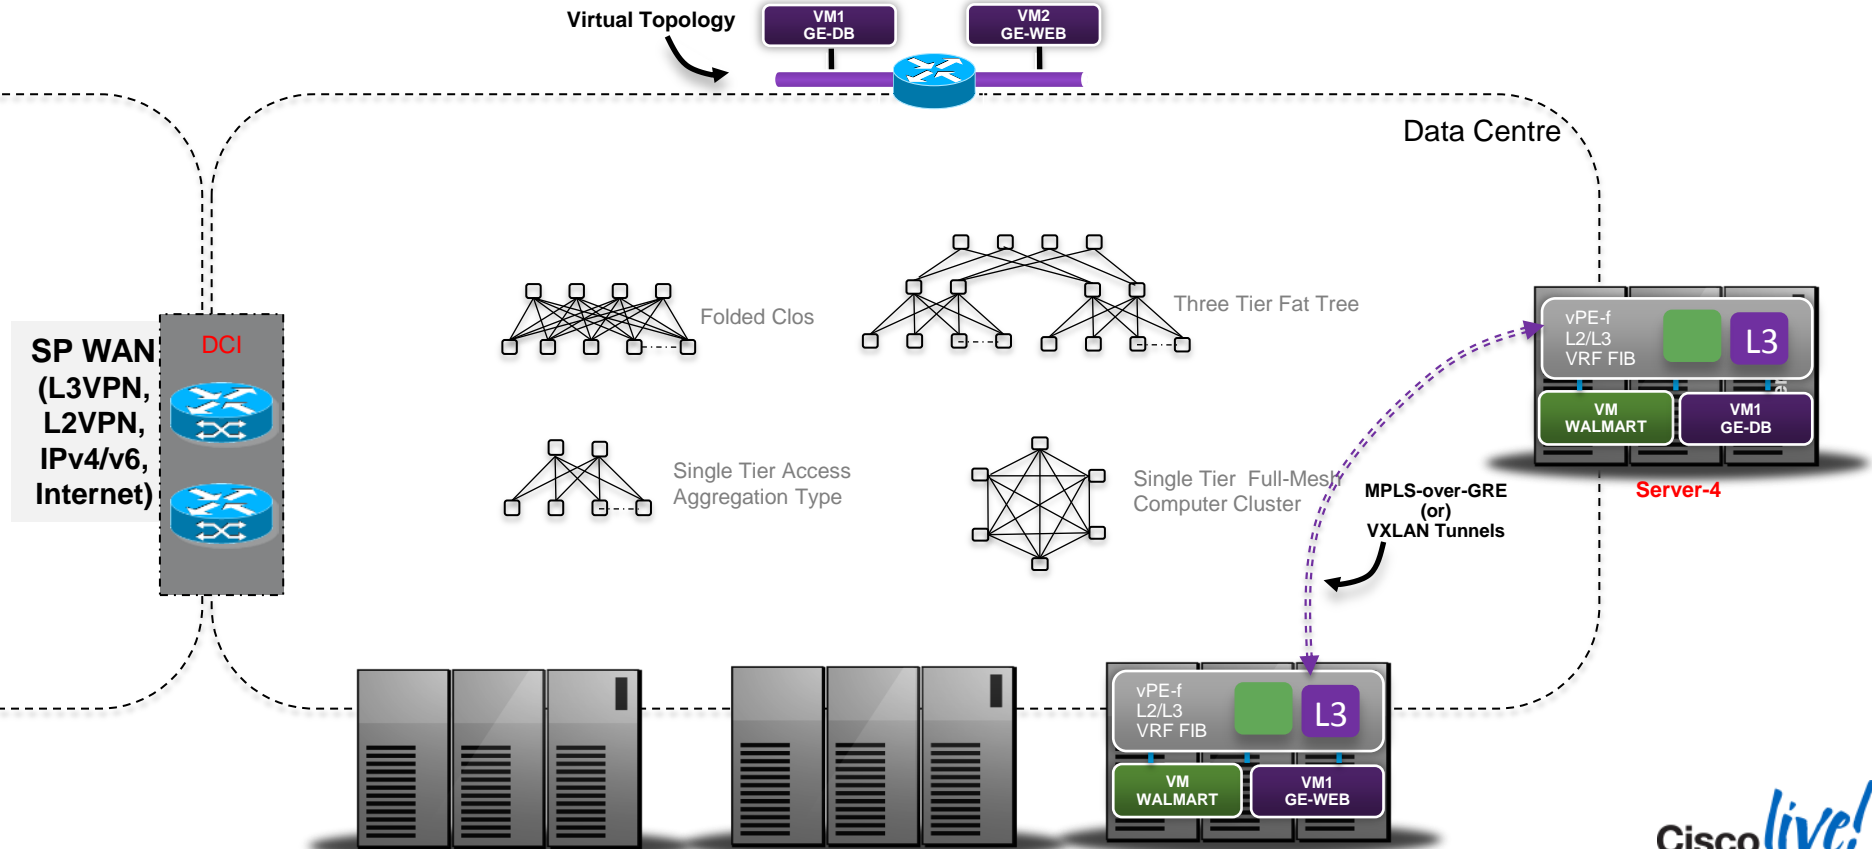
L2 Segment and DCI

- Each vPE-f has VRF L2 tables
- vPE-f populated with MAC entries
- VMs see each other in an L2 segment
- MT traffic encapsulated in single transport tunnel
- Only a small class of applications need strict L2 connectivity



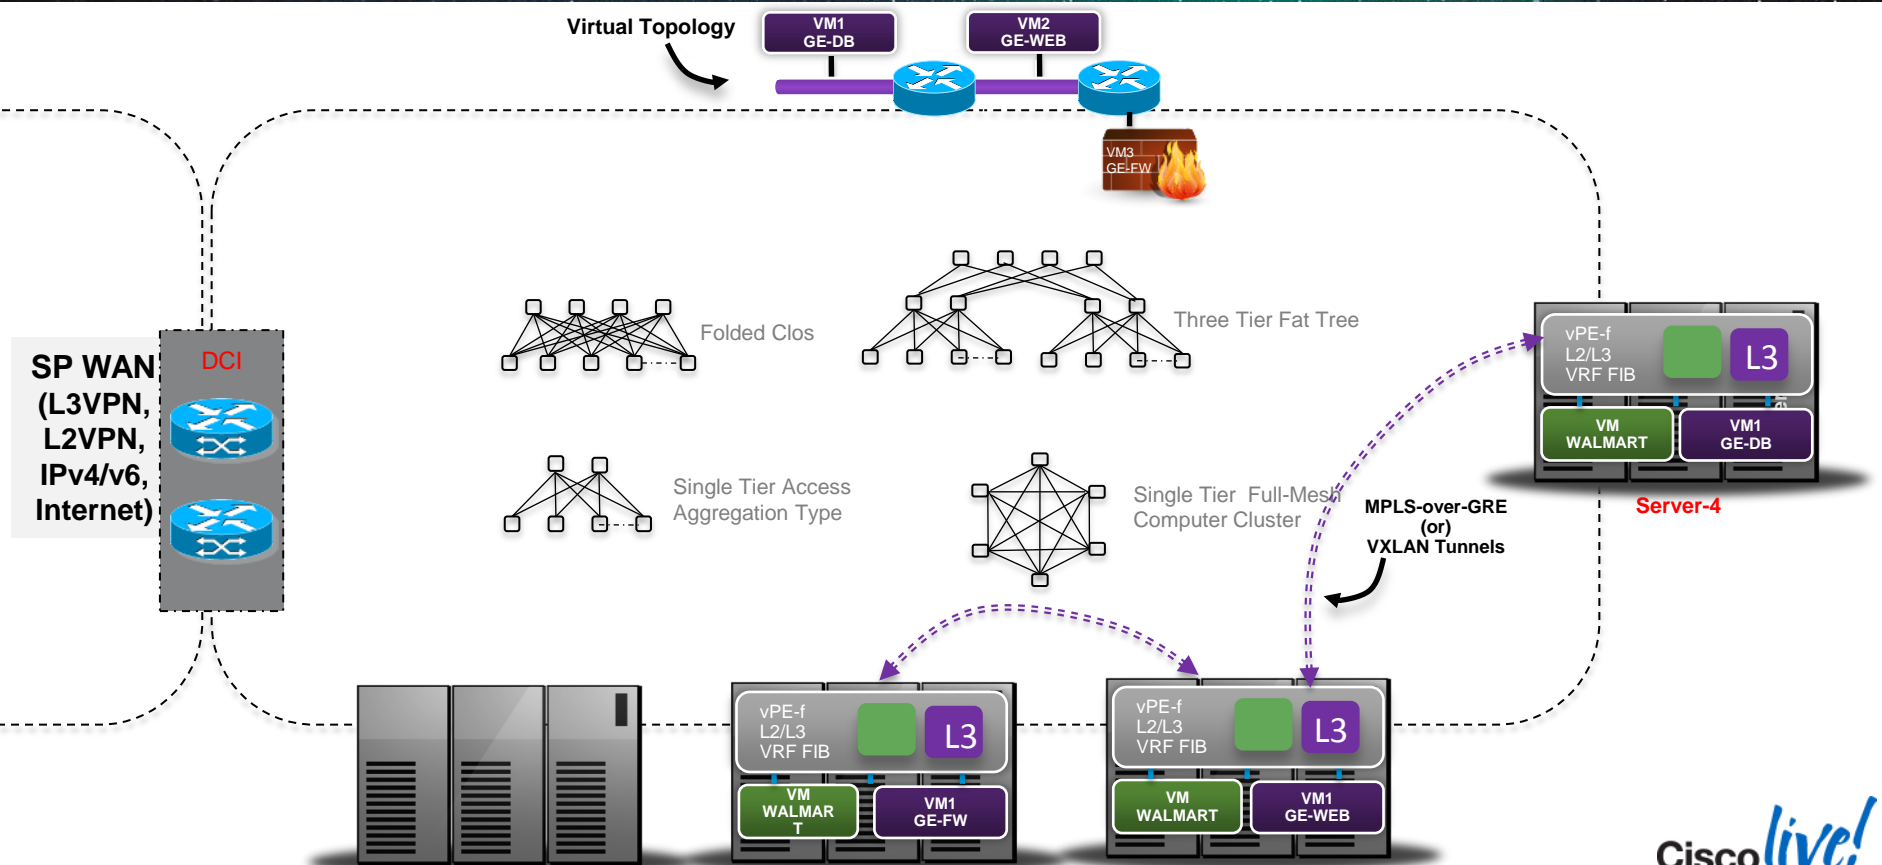
IP Forwarding

- Each vPE-f has VRF L3 tables
- vPE-f populated with L3 /32 or /128 entries
- vPE-f is first hop router/DHCP Relay
- VMs can reach each other in L3 network
- MT traffic encapsulated in single transport tunnel



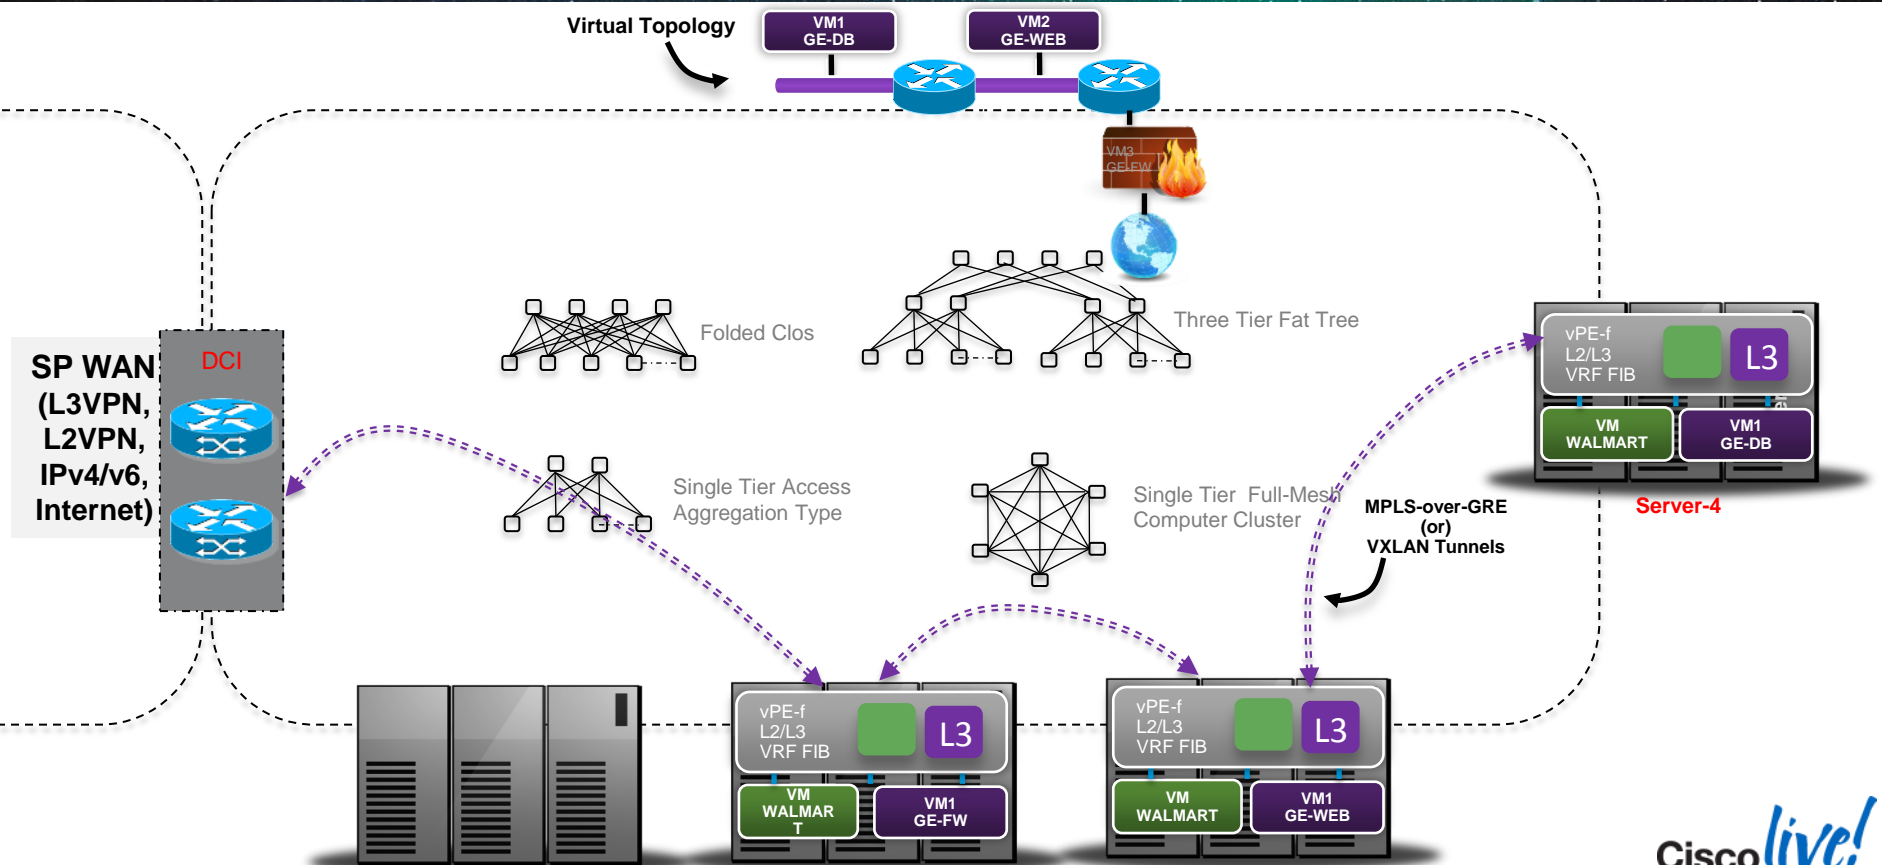
Network Function Virtualisation

- Network Services could be bump-in-the wire services or termination services
- Using L2/L3 entries in tables an arbitrary services topology can be created
- No hair-pinning of traffic as it moves from service to service
- Control Plane responsible for computation of paths and optimal routing of traffic
- Bring-your-own-Service or choose from Cisco service catalog



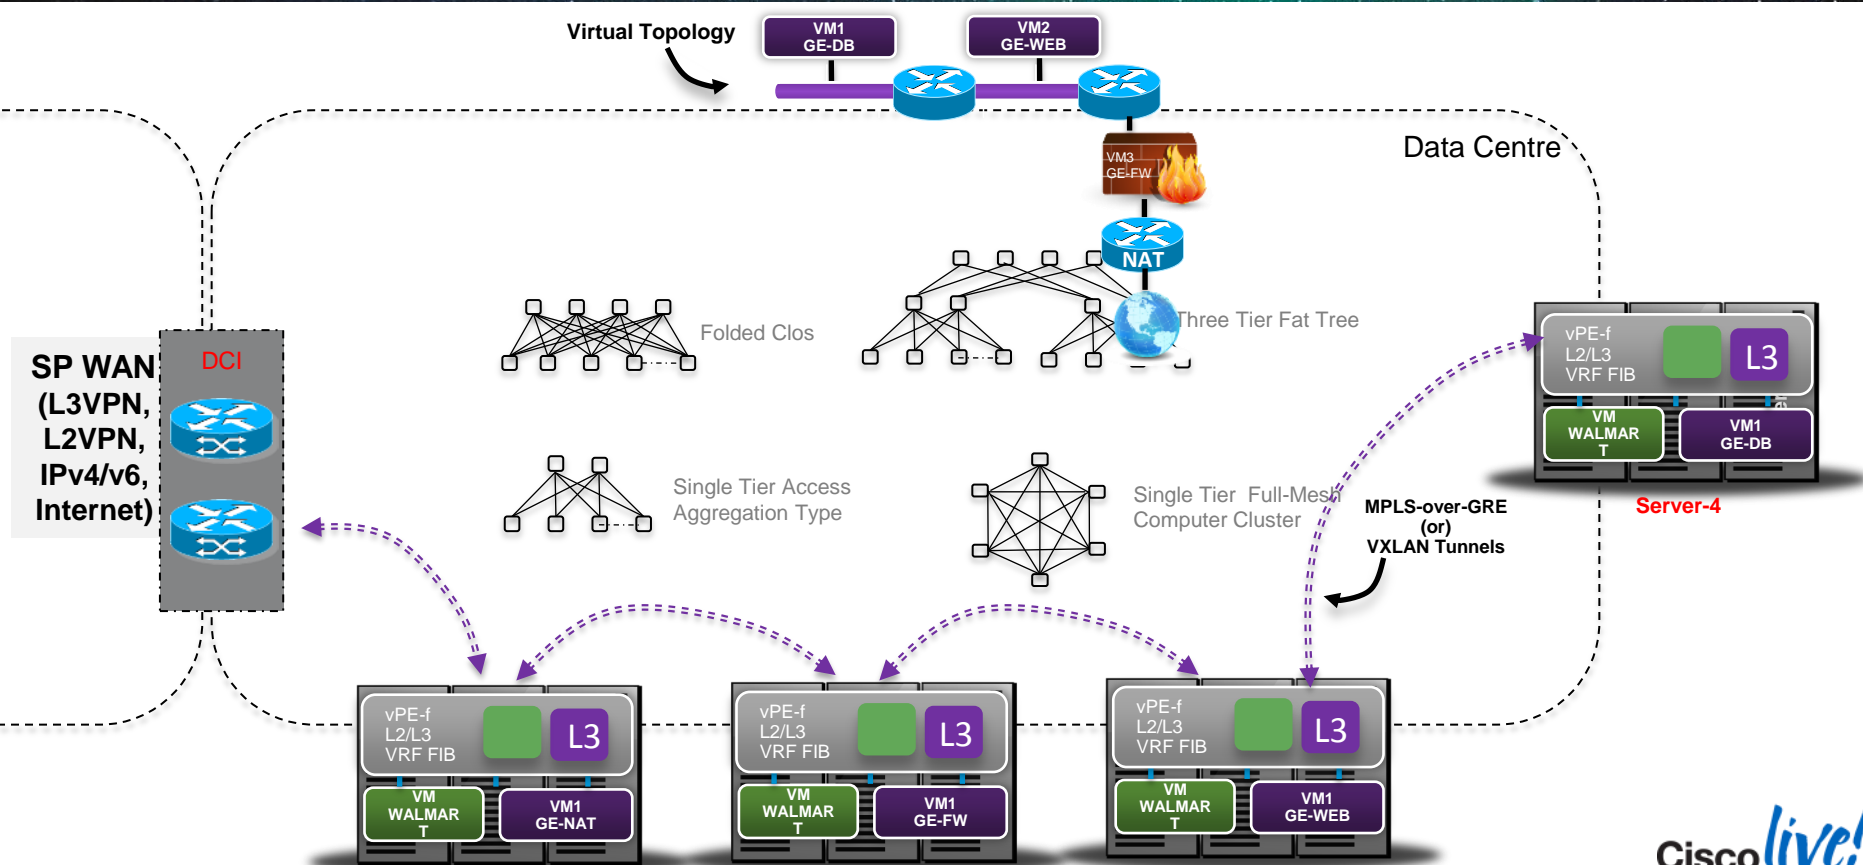
L3VPN, L2VPN & Internet Access

- DCI can be either by injecting /32 or aggregates in SP-WAN MP-BGP
- All VMs default route to DCI for unknown destinations

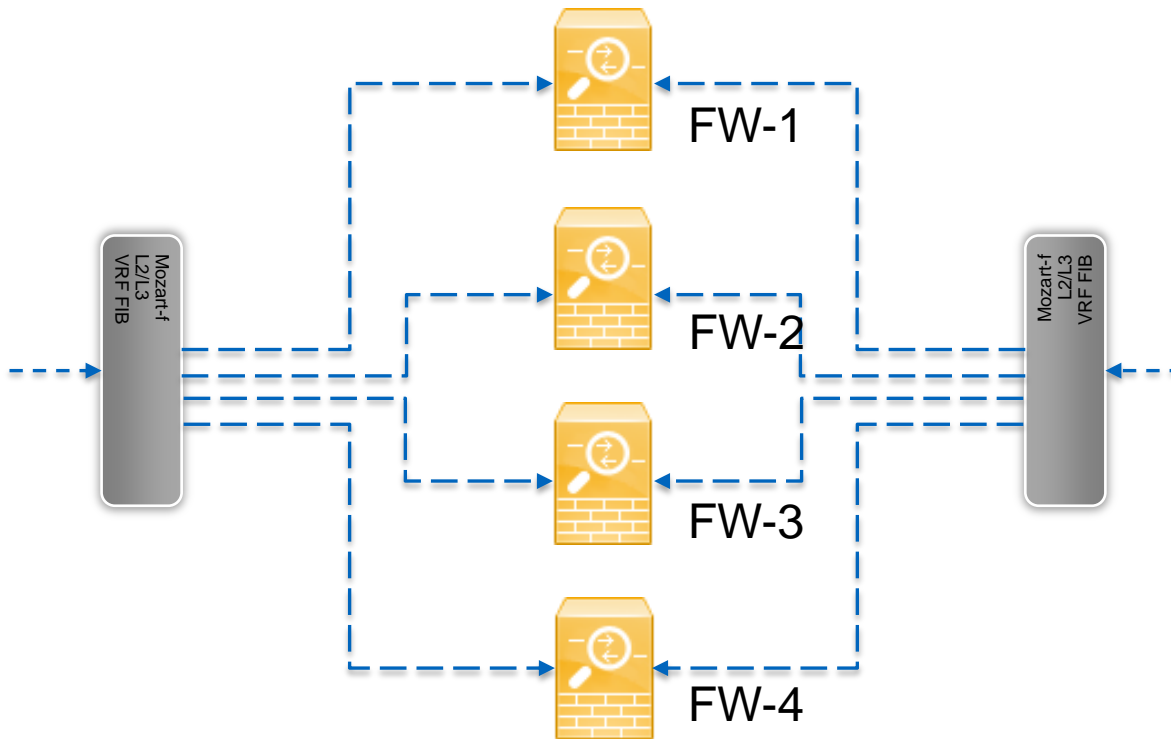


ESP Service Chains

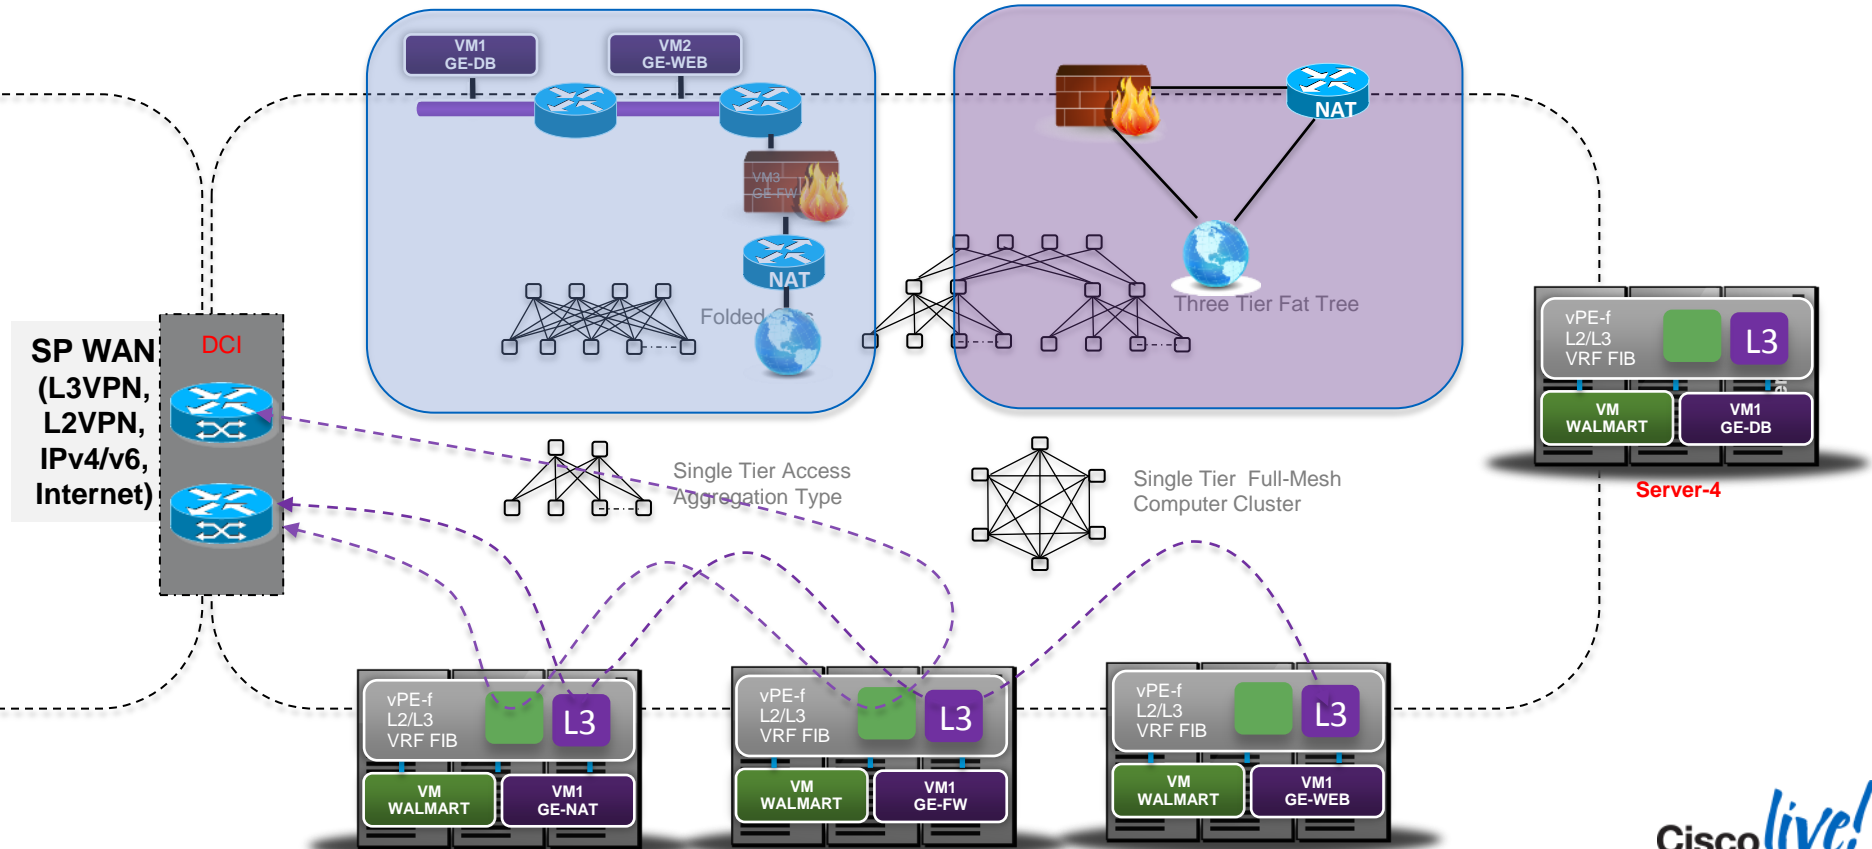
- Network Services can be daisy chained
- No restriction on the number of services in a chain
- Services can be dynamically inserted in the chain



NFv Horizontal Scale, Stateful Load Balancing, Elasticity & Flow Stickiness

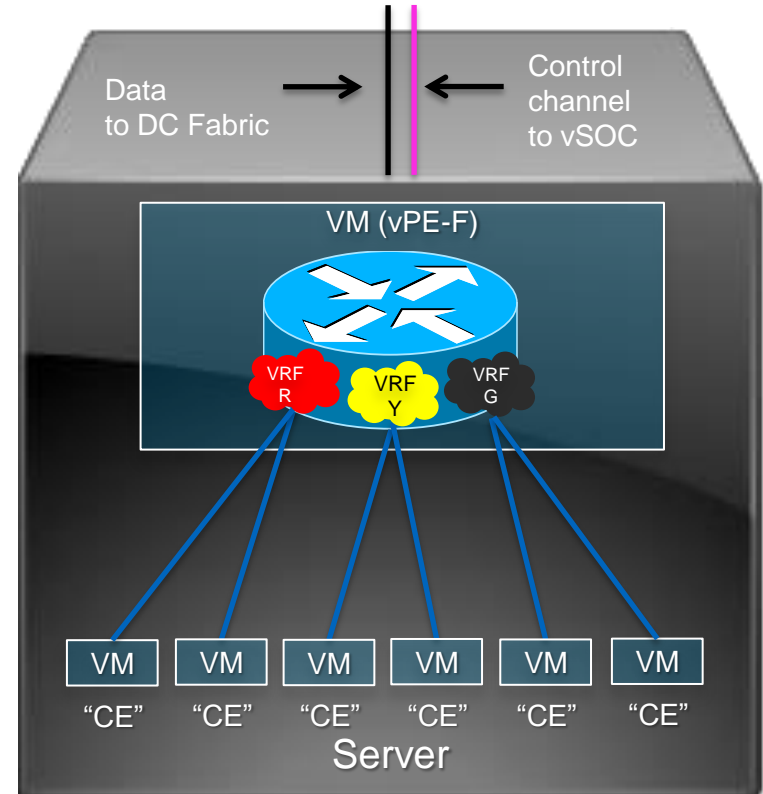


Multi-Tenancy, Varied Topologies

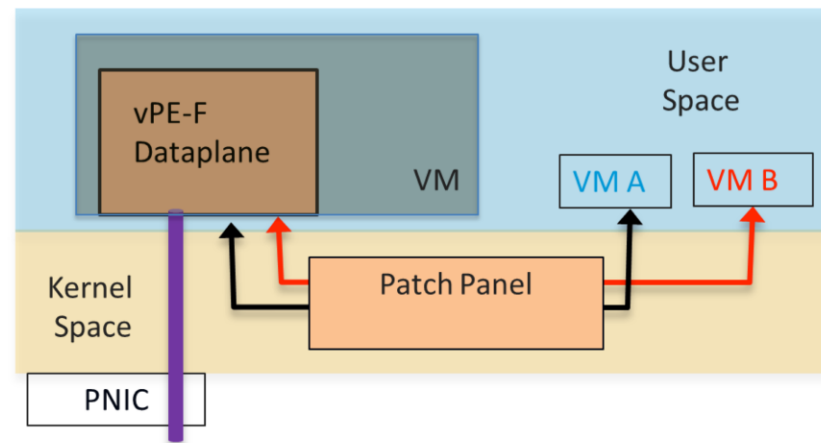
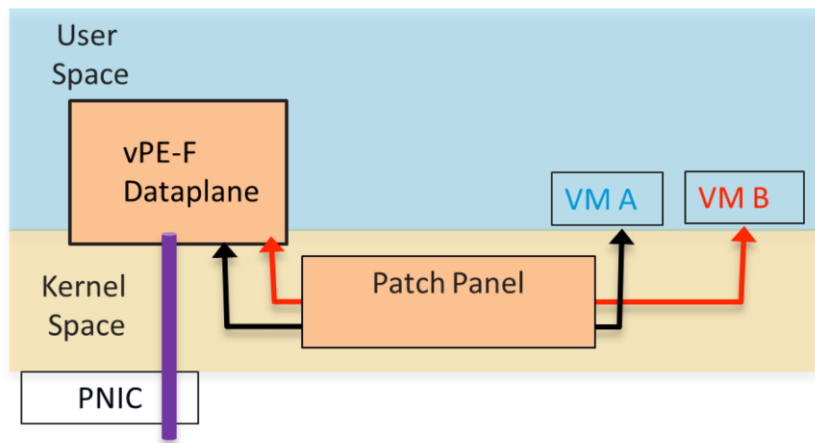


Virtual Packet Edge Forwarder (vPE-F)

- Light weight, high performance software forwarding plane
- Provides highly optimised forwarding in x86 environment
- Runs once on each server
- Contains a unique forwarding context per tenant
- Provides per-tenant L3, L2 and PBR forwarding for service chaining
- Provides IP routed and L2 P2P transport
- Provides DHCP relay, ARP function
- Programmed by vSOC Controller using YANG over RESTConf
 - All forwarding controlled centrally
 - Granular L3 and L2 forwarding entries
 - N-tuple match



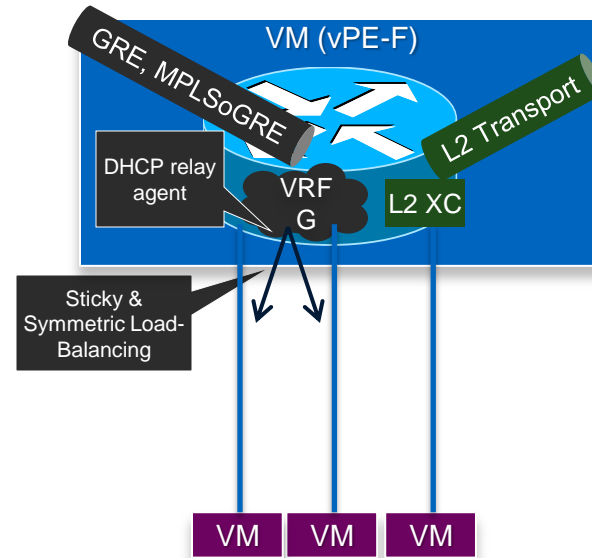
VPE-Forwarder Deployment Modes




- The patch panel provides a virtual point-to-point connection from the tenant VMs to the vPE-f dataplane
- Patch panel is a L2 switch running as host kernel module configured for Point-to-point connectivity without Mac learning
- VM deployment model: easy portable, high performance

VPE-Forwarder Capabilities

- L3 IP stack and Forwarding
 - IPv6
 - IPv4 (ARP, ICMP, etc.)
 - VRF aware FIBs for all address families
 - un-equal-cost multipath forwarding
 - ARP/ND Proxy
 - DHCPv6 Relay
- L2 Forwarding
 - VLAN crossconnect
 - L2 P2P
 - L2 Bridging
- Load-Balancing
 - Sticky load balancing onto stateful services (e.g. firewall)
- Tenant Encapsulation
 - Ethernet
 - 802.1q (single-tag) VLAN sub-interfaces
- Network Encapsulation
 - Routed: GREoIPv4, MPLS-o-GREoIPv4
 - L2 Forwarded: L2tpv3 L2 cross-connect



vSOC – Virtual Systems Operations Centre



▪ Orchestration

- Exposes a North Bound ReST API that allows provisioning of services
- Implements model driven workflows to realise the services
- Secure REST NB API with RBAC support
- Implements service templates for easy OSS integration

▪ Compute/Storage Control

- Service VM Lifecycle management
- VM Monitoring & VM Recovery
- Scale up/down of VM based on elasticity criteria
- Integrates with NAS, SAN systems (CEPH, NetAPP)

• Network Control

- Controls forwarding entries in vPE-forwarder
- Controls routing to DCI through XRvR

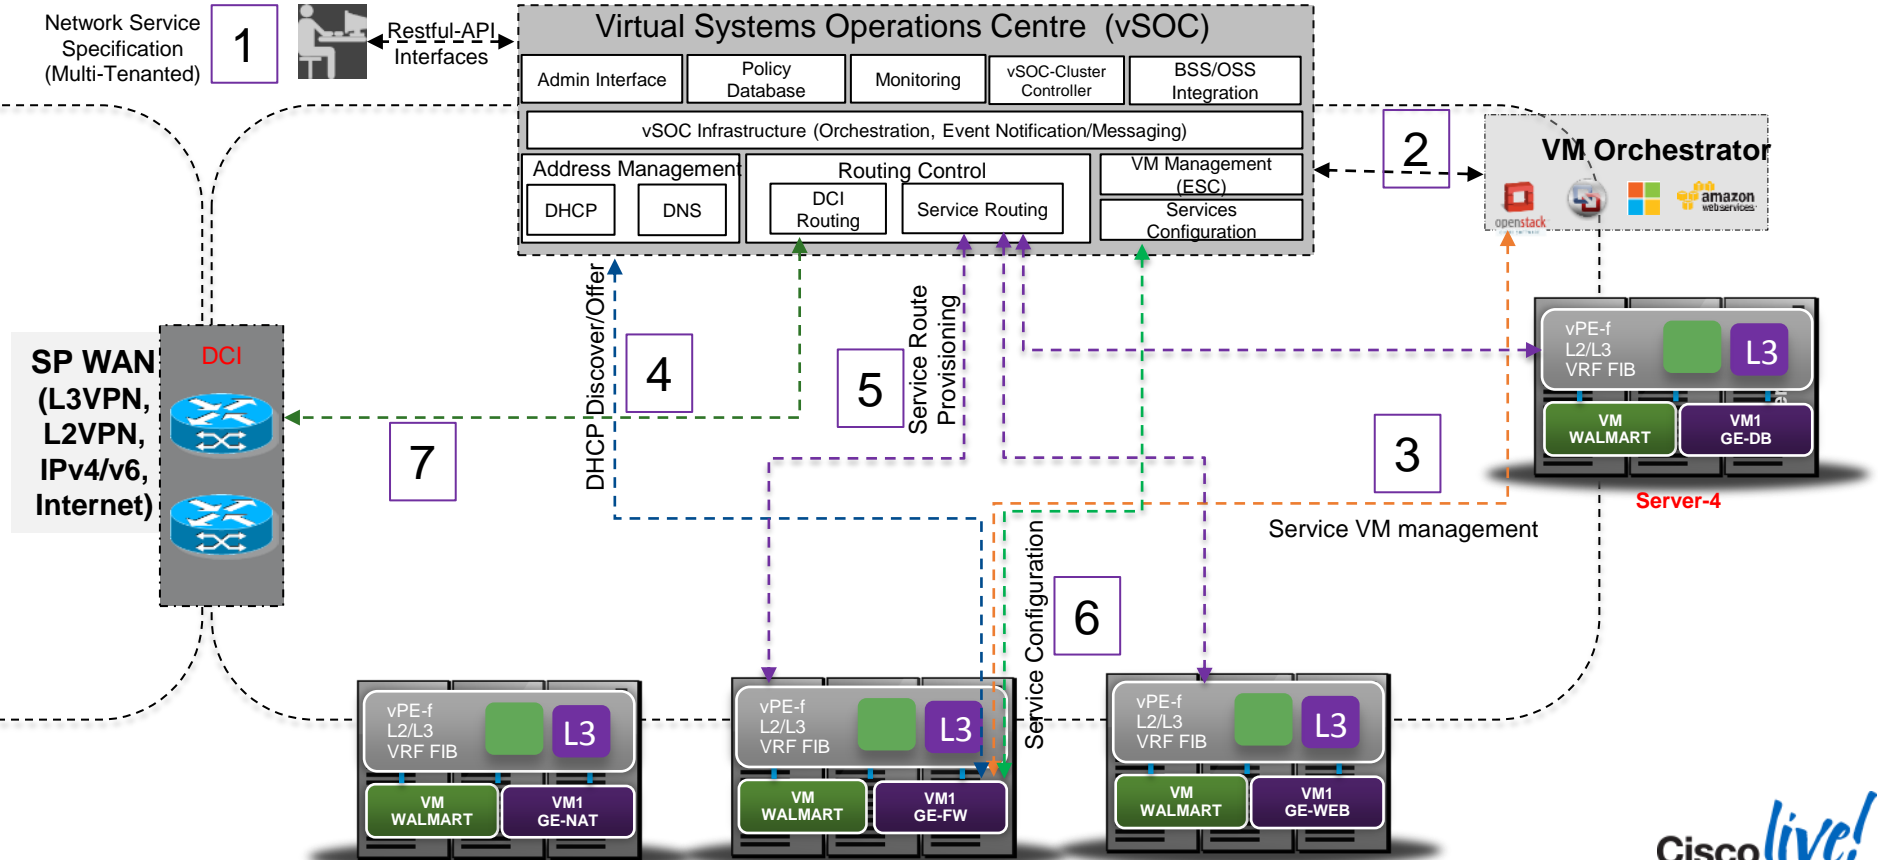
• Service Provisioning

- Configures DHCP Server
- Configures Service VMs e.g. ACL, Firewall, etc. on CSR
- Configures DCI router for L3VPN VRF & MPLSoGRE tunnel for connection to vPE-f
- Configures remote PE and CPE's.

• System Management

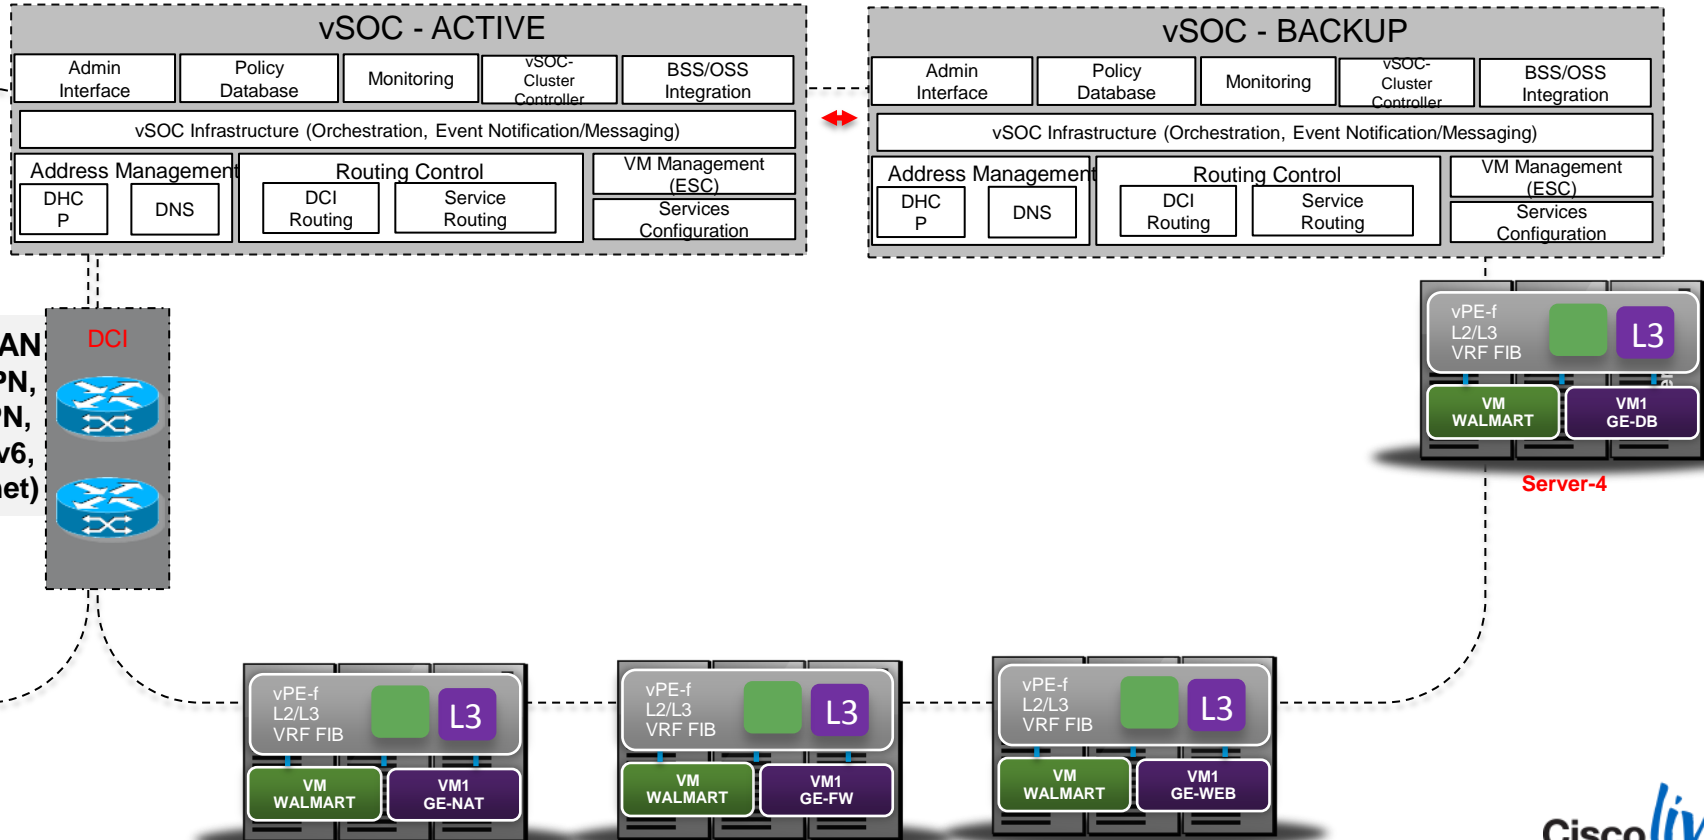
- Auto Installation of the system
- vSOC HA Control
- ISSU Control
- Packaging

vSOC Call Flow

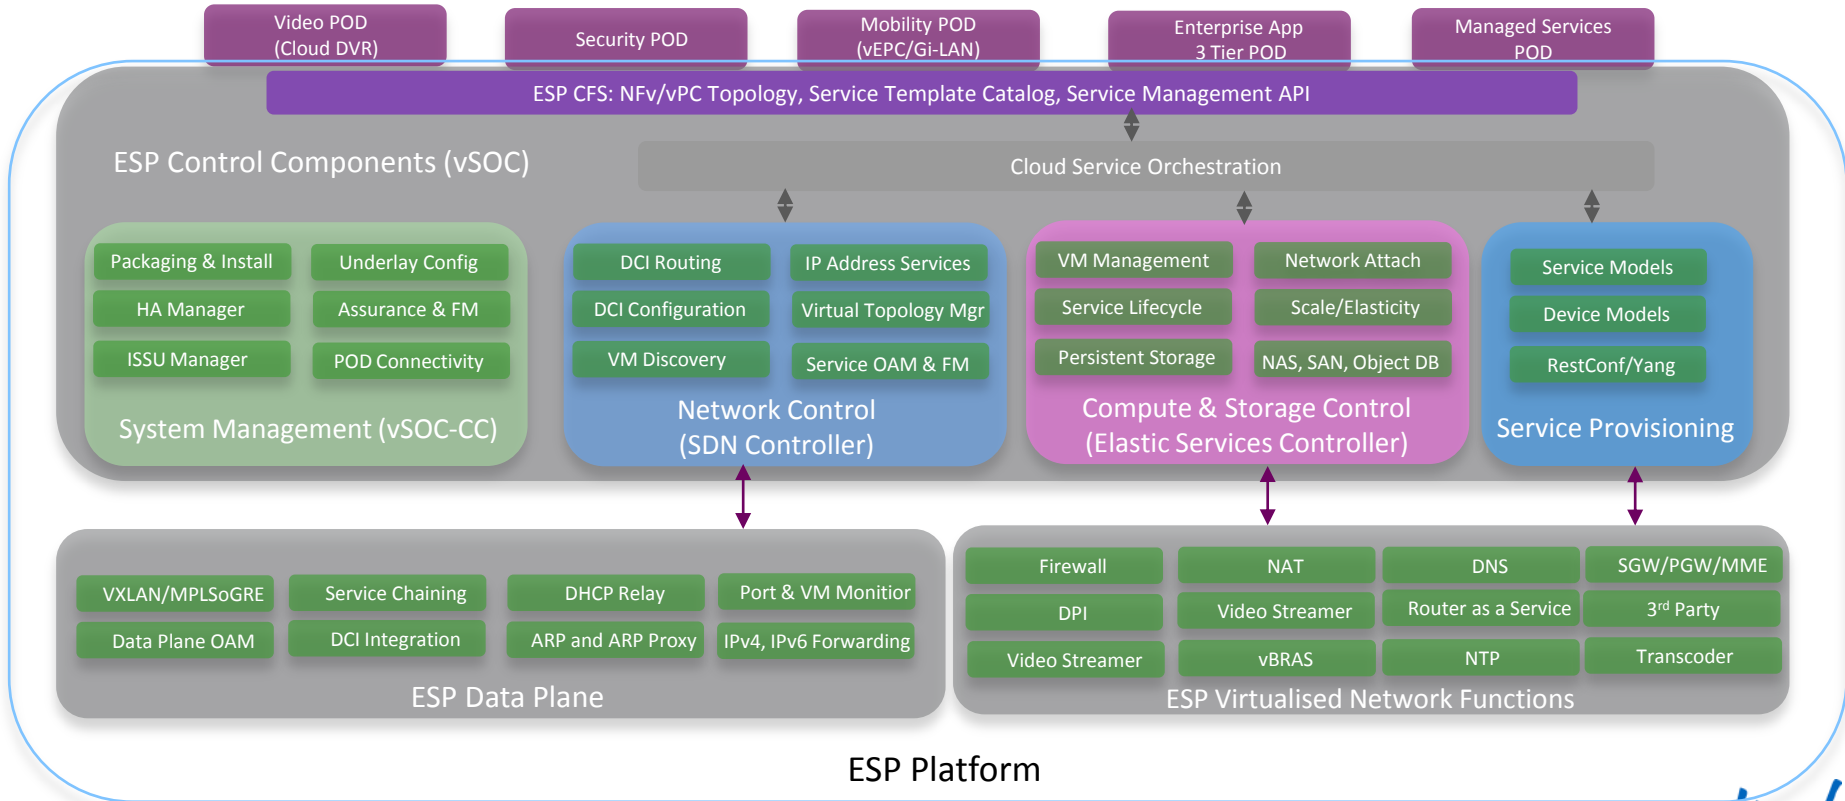


ESP High Availability

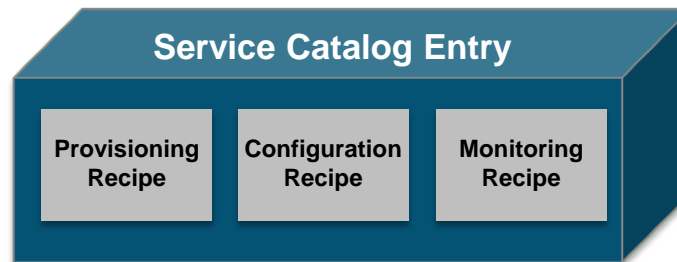

- vSOC is not required in steady state
- vSOC supports Active-Standby HA



ESP Architectural Components



vSOC Elastic Services Control



Provisioning

- **VM Disk Images** - can be several images in a topology
- **Virtual CPUs, MEMORY** – resource requirements for service
- **Network** – NIC interface type & network topology/configuration (basic or complex)
- **Hypervisor** - supported hypervisors for this service

Configuration

- **Puppet/Chef**- Service has a puppet/chef agent that allows it to have it's configuration pushed to the VM after boot-up
- **Inject** – Orchestration system can inject the configuration into the VM image file-system at provision time

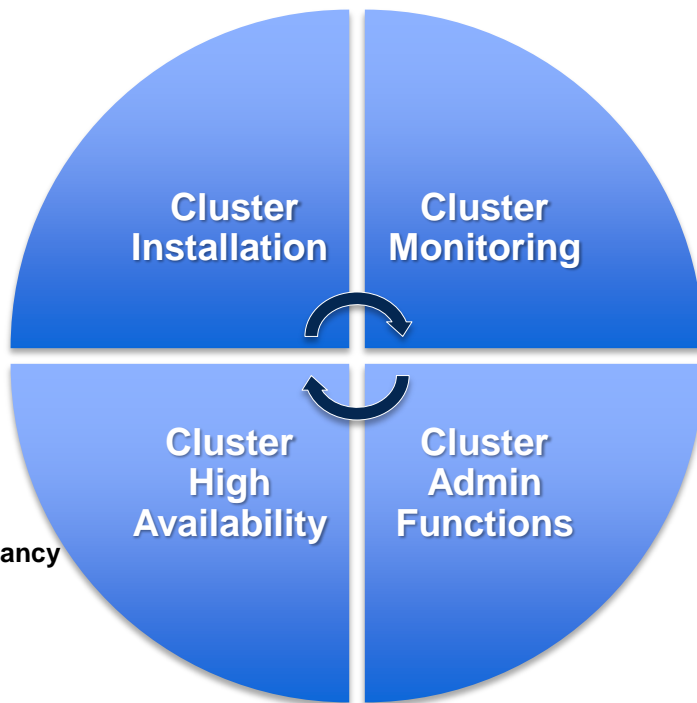
Monitoring

- **SNMP**- service has an SNMP agent & metrics that can be monitored
- **Ganglia** - service has a ganglia agent & metrics that can be monitored
- **PING** – service has no monitoring support so is deemed to be alive when VM responds to pings

vSOC - System Management Overview

- Packaging
- Zero Touch Install
- Policy Based Declarative Install
- System Underlay Inventory
- Software Versioning
- ISSU

- Fault Detection
- NIC Failure Detection & Recovery
- Server/VM Failure recovery
- Control VM Switch over
- NIC Teaming, bonding and Redundancy
- DCI Redundancy
- Storage HA
- Service Assurance



Cluster
Installation

Cluster
Monitoring

Cluster
High
Availability

Cluster
Admin
Functions

- Server Monitoring
- Control Plane VM Monitoring
- NIC Monitoring
- Storage Monitoring
- Process Monitoring
- DCI Monitoring

- System Log Management
- Time synchronisation
- Name Resolution
- License Management
- Crypto Key Management
- Backup and Restore
- Storage Management


Provisioning the ESP System

MODEL BASED – DECLARATIVE SERVICE DEFINITION


REST BASED API



Tenant Identifier, Tenant Specific VPN Identifier, L3VPN & L2VPN Extended Communities, Organisation Definition, Global Tenant Specific IP Address Pools



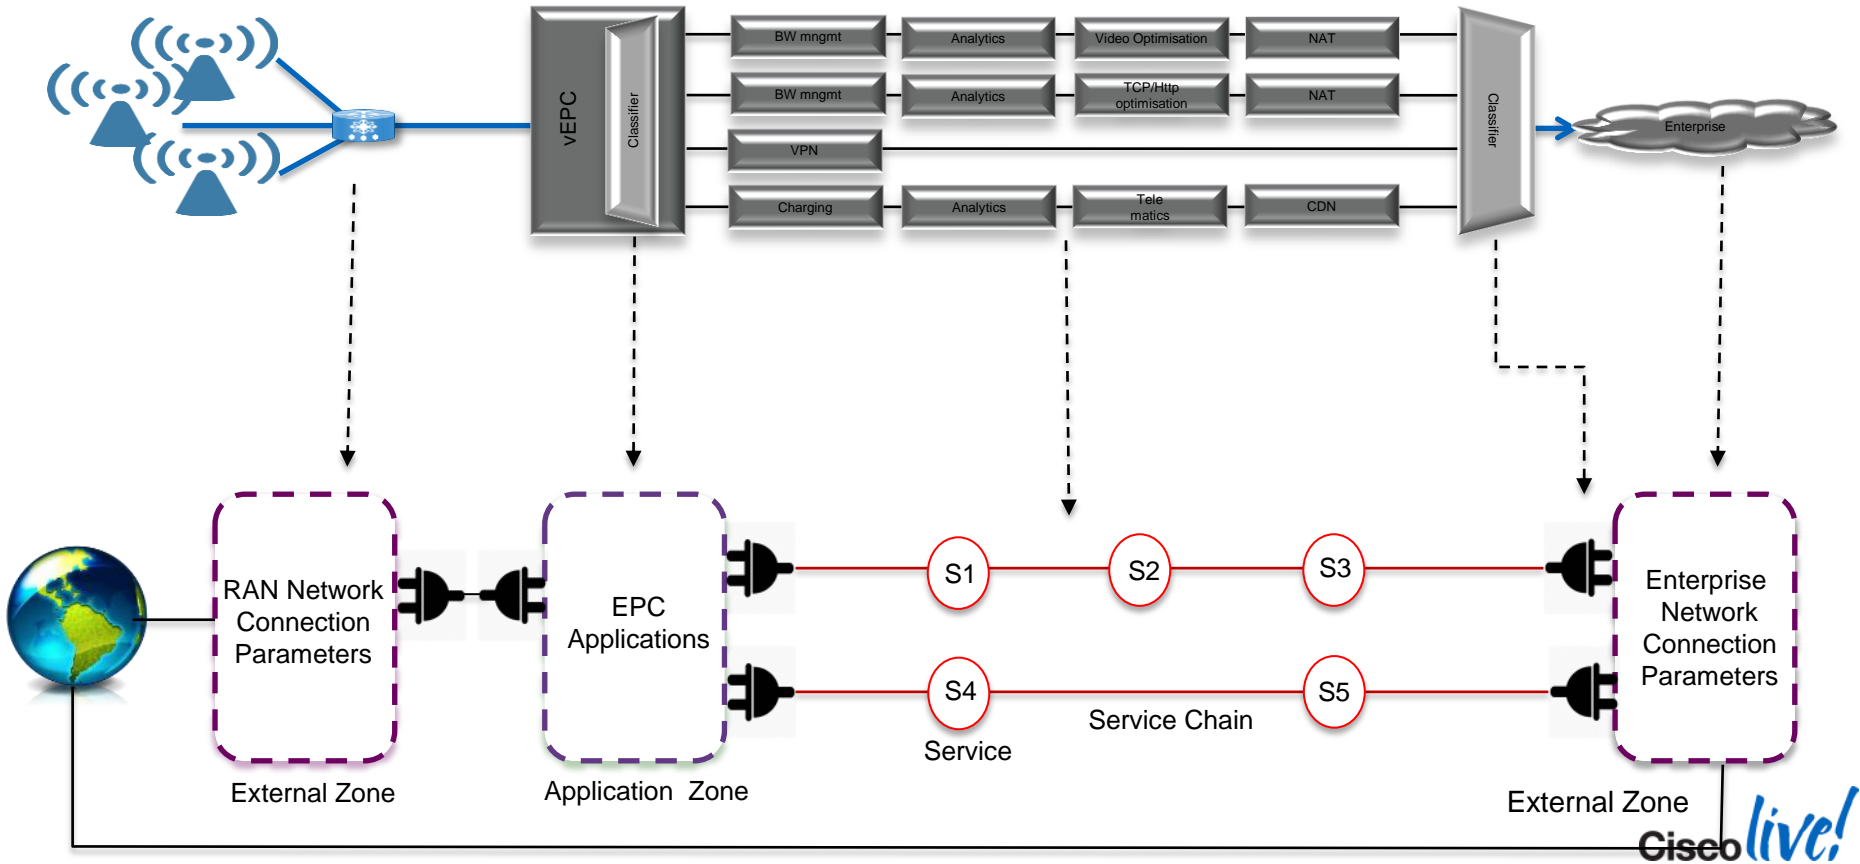
Network Zones, Zone Connectivity, External Zones, Managed Zones, Transit NFv Appliances, Terminate NFv Appliances, Service Topology Definition, Service Chain Definition, Multi-Path Requirements



Define CPU, Memory, Network Interfaces, Horizontal Scale Factor, Elasticity, Disk Storage, Persistency Requirements, Service Configuration



ESP Models Example: Mobility (vEPC + Gi-LAN)

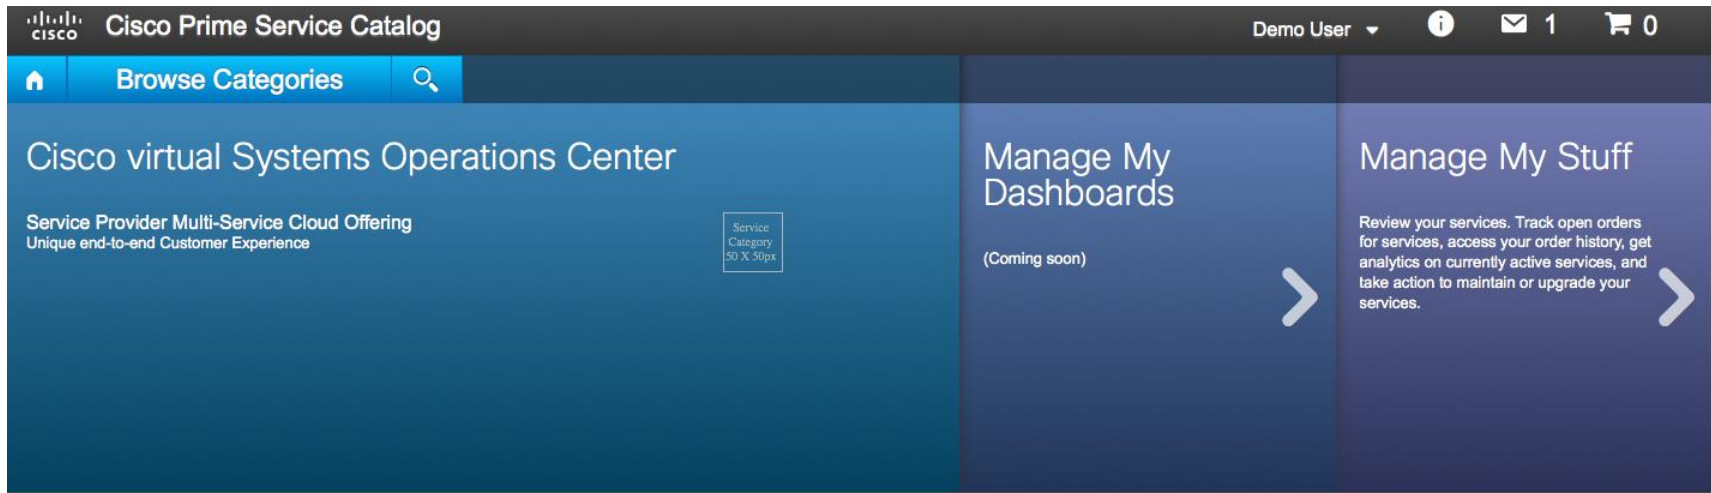


Customer Experience - GUI

- Single portal for customers to login and provision their network and application VMs
- Each customer can create multiple topologies
- Traffic for a topology could come from Internet, existing L3VPN network, L2VPN network
- Topology composed of multiple zones
- Inter zonal traffic subjected to one or more services (FW, NAT, DPI, Load Balancer)
- Ability to provide pre-packaged end application services such as Web Server, Video Server, Mail Server, Database Servers, Hadoop Cluster, etc
- Design template library and custom network topology templates for provisioning ease.
- BYOS – Ability for customers to bring their own service appliances



Customer Experience - GUI



Tenant and User Management



Tenant Management



Organization Management



User Management

Topology Management

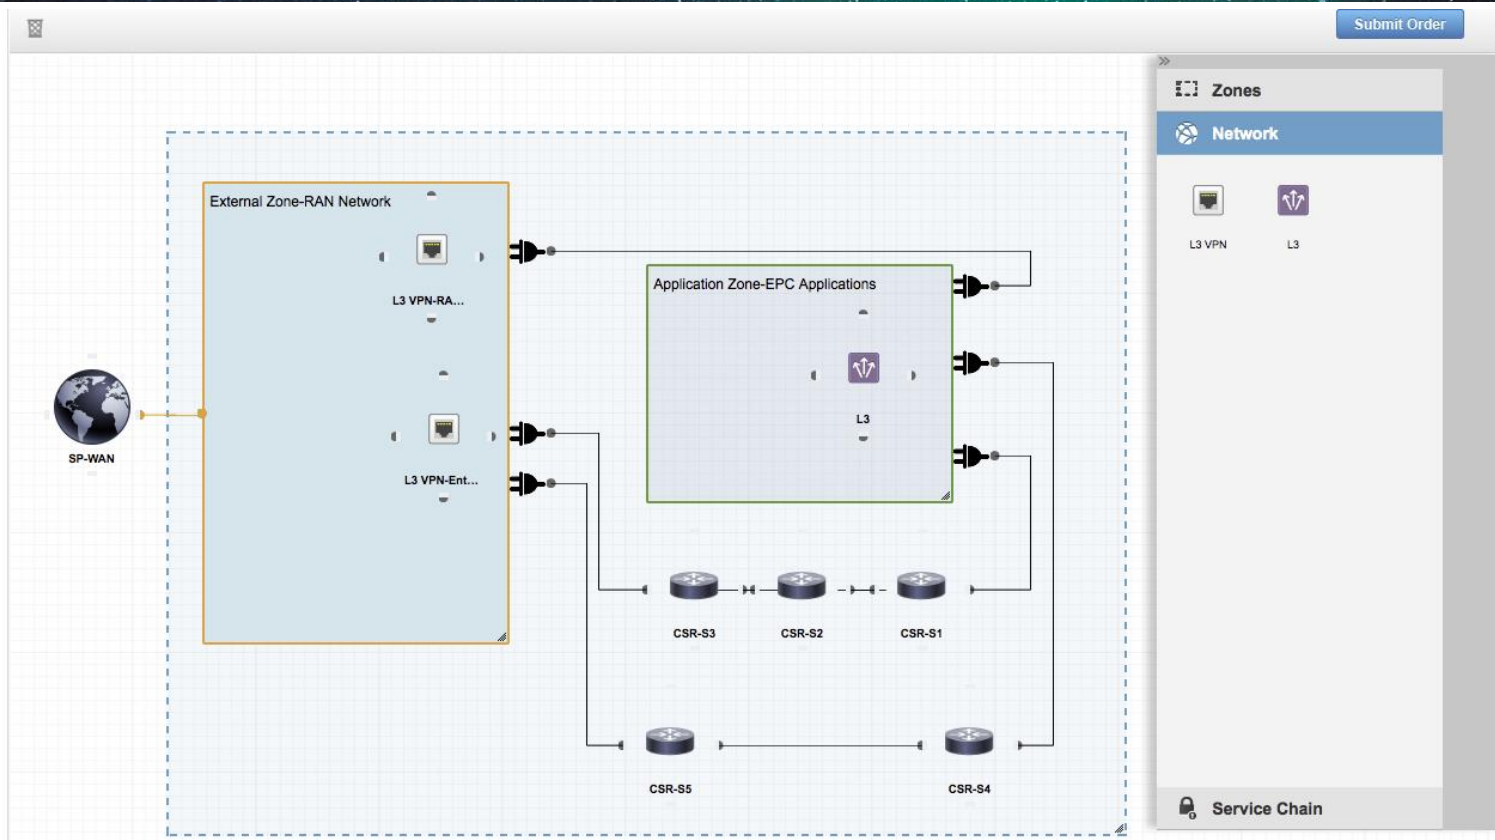


Network Management

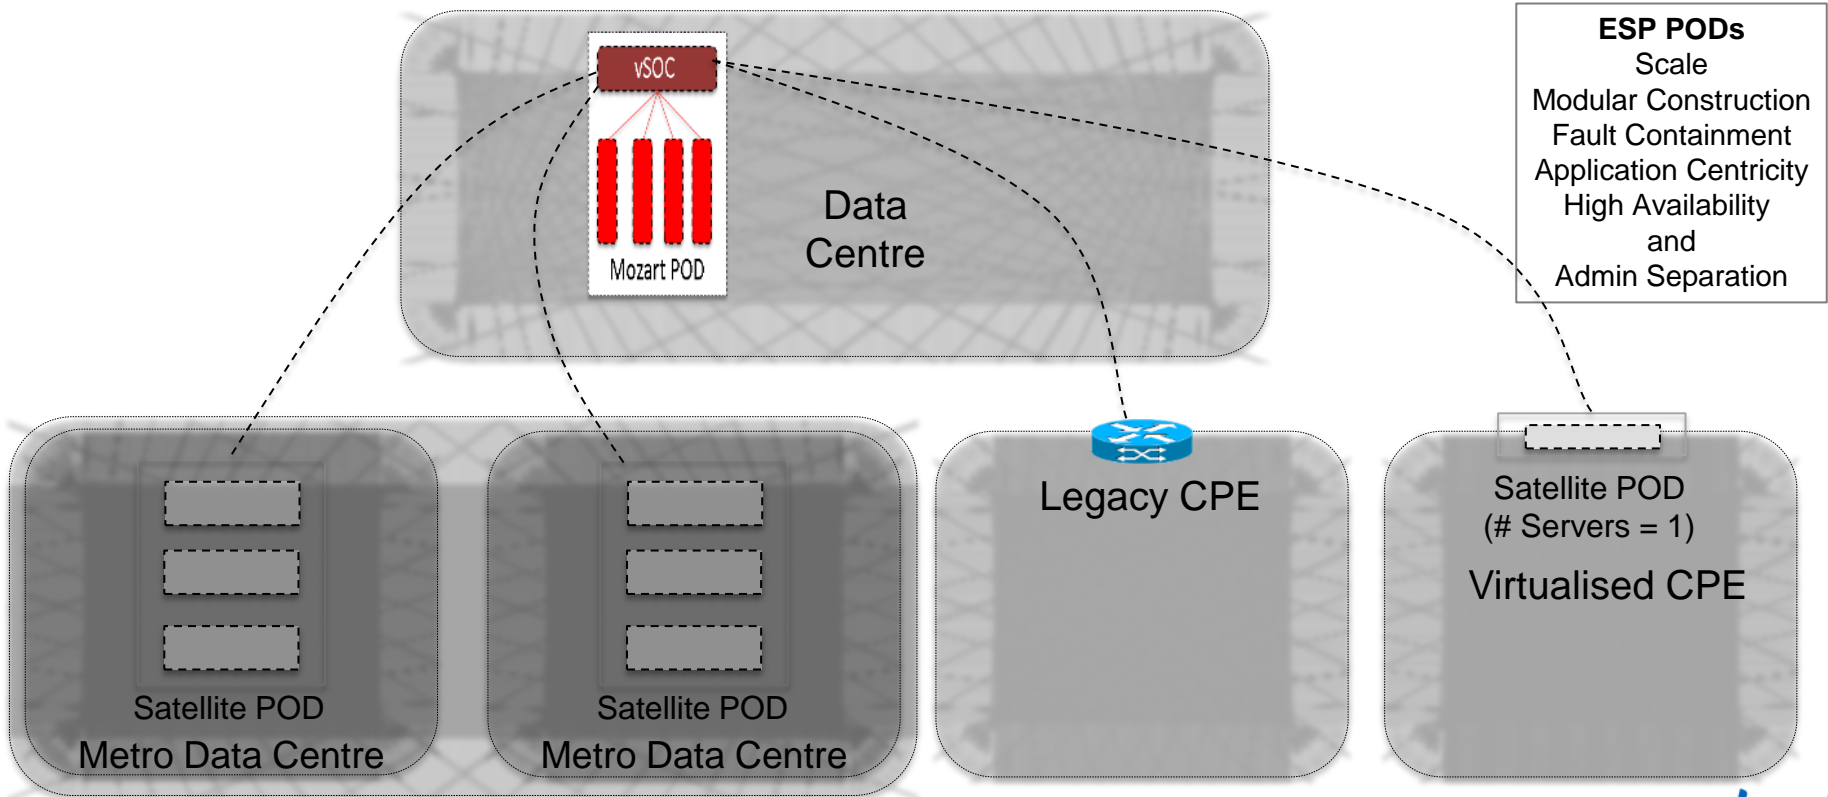


Zone Management

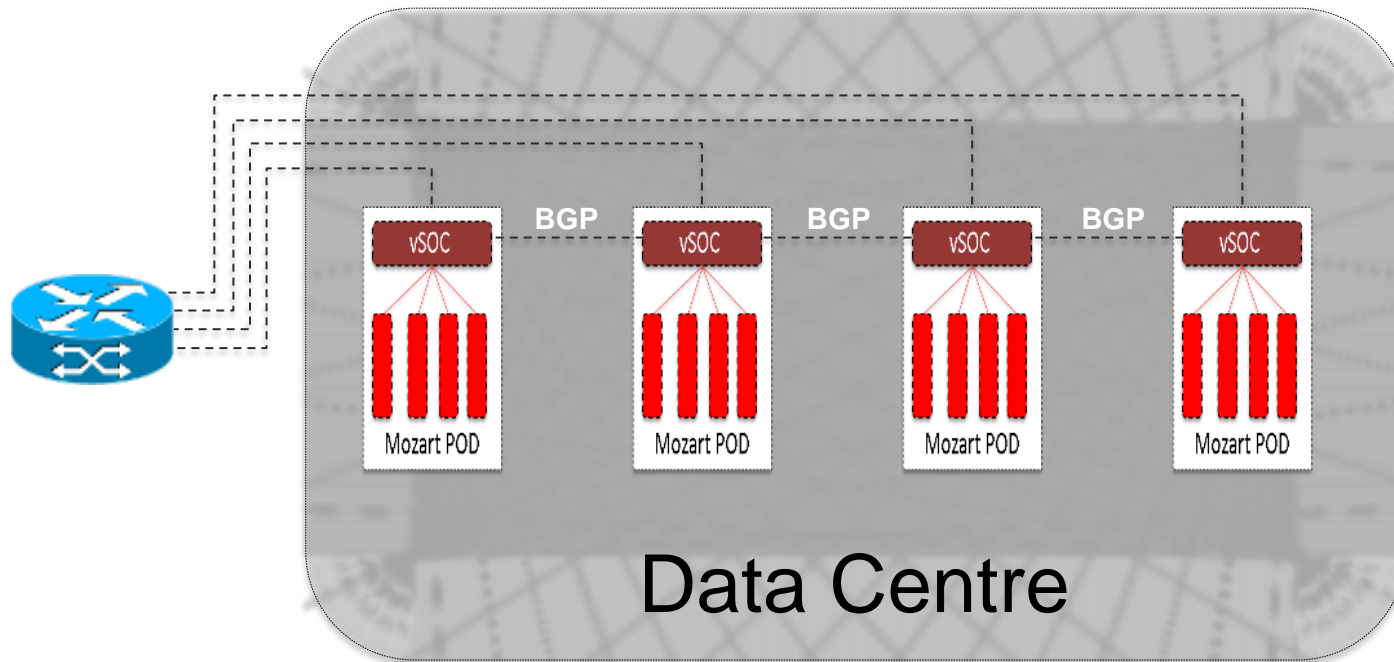
ESP GUI: Designing The Mobility Service



ESP PODs & Satellite PODs

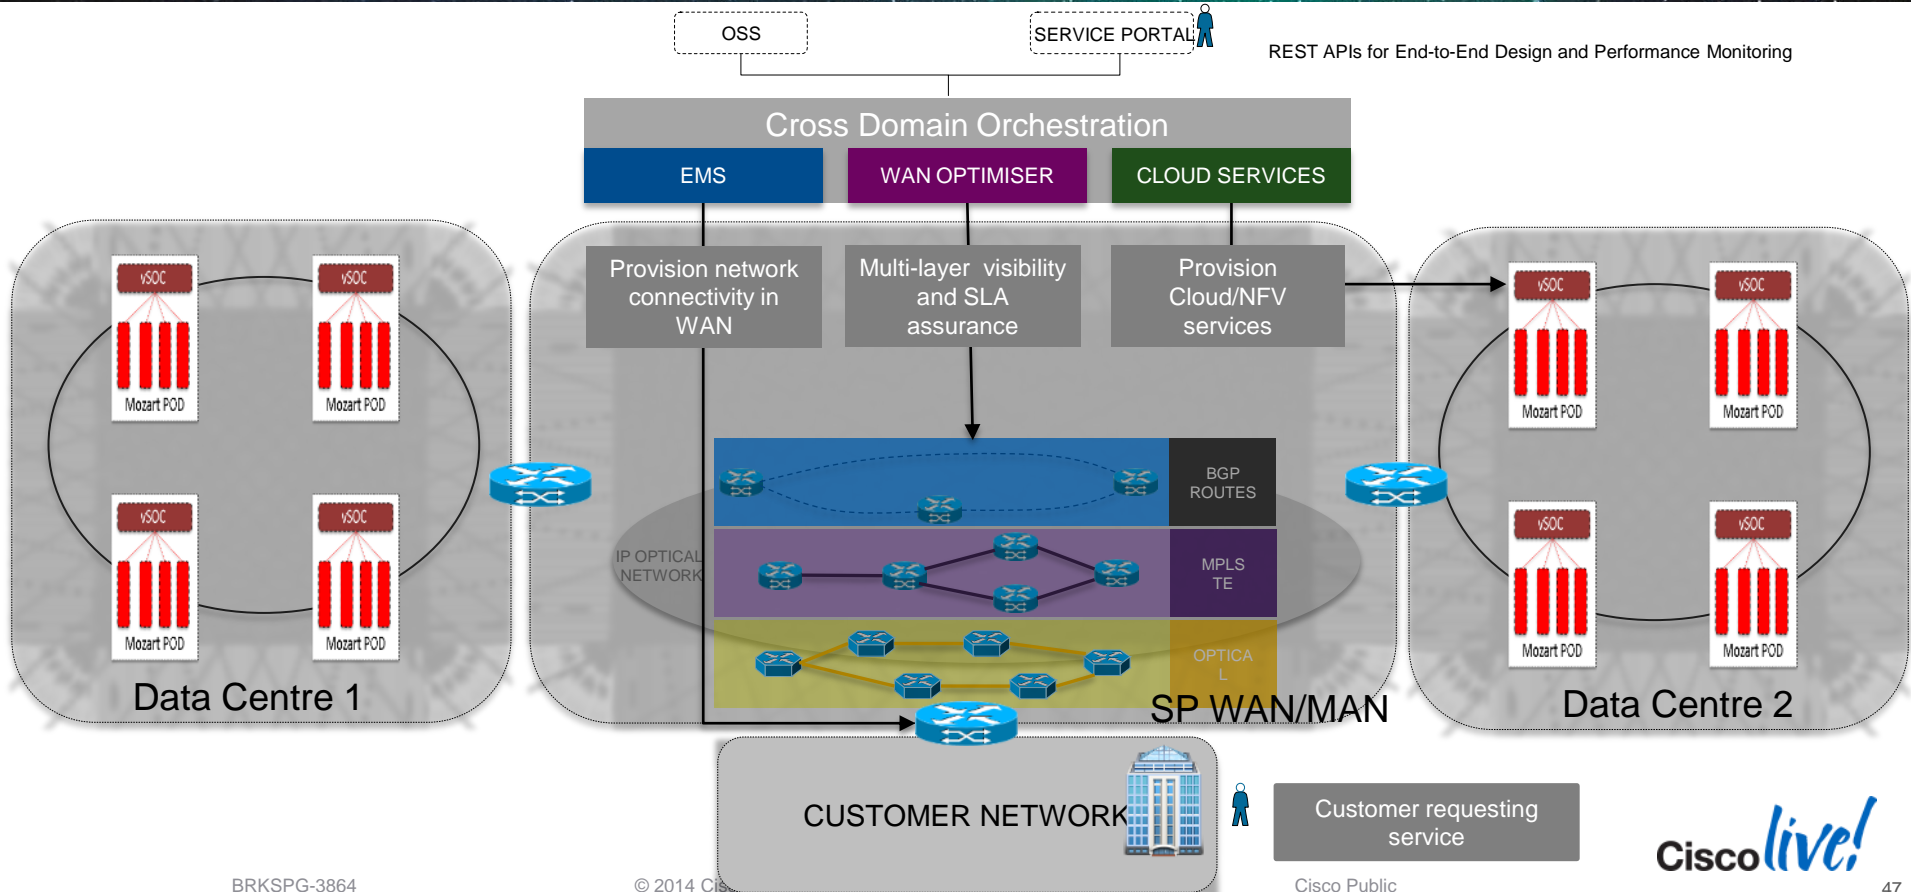


ESP PODs

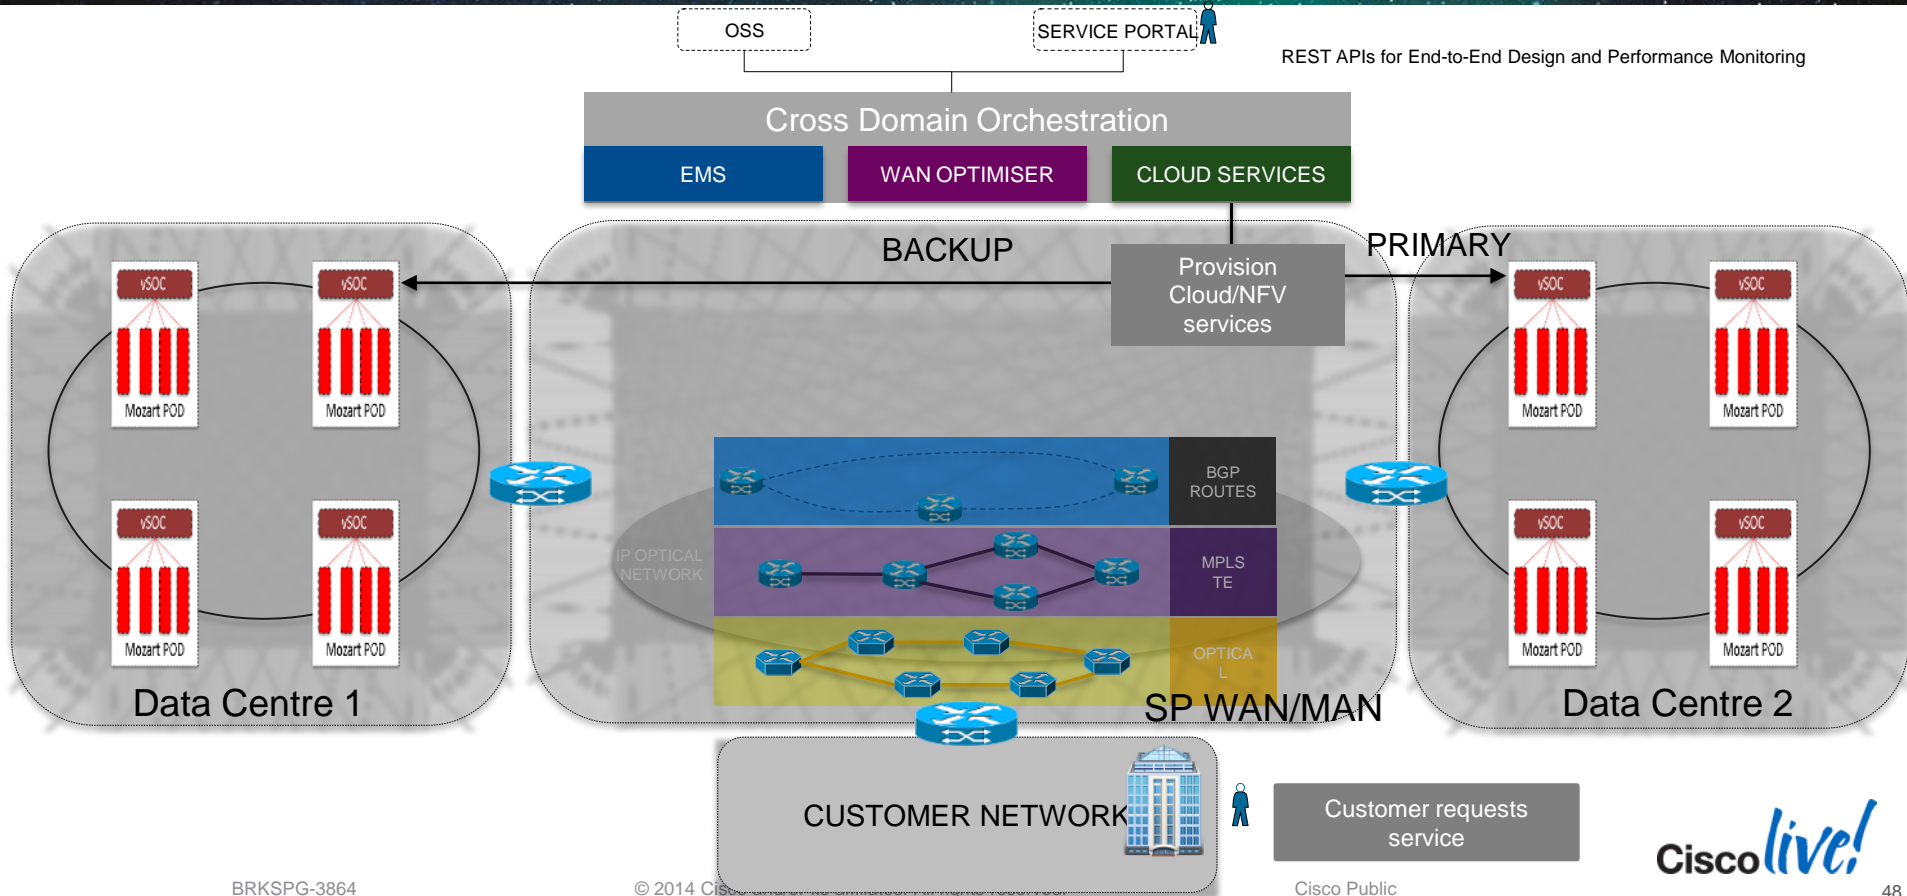


ESP PODs
Scale
Modular Construction
Fault Containment
Application Centricity
High Availability
and
Admin Separation

ESP SLA Aware End-to-End Service Provisioning



ESP Geo-Redundancy



Key Solution Highlights

End to end Solution offering

Based on Open, standards-based interfaces

Highest performance virtual forwarder

Virtual forwarder in a VM isolates network failure domain from compute

Overlay architecture independent of underlying fabric

Self Service model and automated network config enables zero touch provisioning

Service configuration integrated with Solution

Elastic Service management



Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm



Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.

www.CiscoLiveAPAC.com



CISCO TM