

NAC Layer 3 Out of Band Design Guide That Uses VRF-Lite for Traffic Isolation

Document ID: 108540

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Configure

- Infrastructure Configuration
- Topology
- Process Flows
- Configuration

NAC Configuration for Layer 3 OOB

- CAS Setup

Verify

- Appendix A: Switch Configurations

Troubleshoot

Related Information

Introduction

Note: Information in this document can change without notice. Confirm all recommendations if possible.

The purpose of this document is to describe a VRF-Lite based implementation of NAC in a Layer 3 Out of Band (OOB) deployment where the NAC server (CAS) is configured in Real IP Gateway (Routed) mode. Layer 3 Out of Band has rapidly become one of the most popular deployment methodologies for NAC. This shift in popularity is based on several dynamics. The first is better utilization of hardware resources. By the deployment of NAC in a Layer 3 OOB methodology, a single NAC Appliance can be made to scale to accommodate more users. It also allows the NAC Appliances to be centrally located rather than distributed across the campus or organization. Thus, Layer 3 OOB deployments are much more cost effective both from a Capital and Operational expense standpoint. There are two widely used approaches to deploy NAC in a Layer 3 OOB architecture.

1. Discovery-Host based approach Uses inherent ability within the NAC Agent in order to reach the NAC Server (CAS). ACLs applied on the access switch control traffic enforcement on the Dirty network. Refer to Connecting to the NAC Server (CAS) using the SWISS Protocol for more information.
2. VRF based approach Uses VRFs to route unauthenticated traffic to the CAS. Traffic policies configured on the NAC server (CAS) are used for enforcement on Dirty network. This approach has two sub-approaches. In the first approach, VRFs are pervasive throughout the infrastructure, in which case all Layer 3 devices participate in the tag switching. The second approach uses VRF-Lite and GRE tunnels to tunnel the VRFs through the Layer 3 devices that do not understand the tag switching. The benefit to the second approach is that minimal configuration changes are required to your core infrastructure.

Note: While Layer 3 OOB is one of the most common deployment methodologies, it cannot always be the optimal solution for every environment. There are other options to choose from that can be a more optimum

fit for your particular requirements. Refer to Planning Your Deployment for more information on these other NAC design options.

Prerequisites

Requirements

Ensure that you meet these requirements before you attempt this configuration:

- A basic understanding of Layer 2 and Layer 3 infrastructure operation and configuration
- A basic understanding of the Cisco NAC appliance, and the differences between the various implementation methodologies that are associated with it
- All NAC deployments and designs should be based on clear business requirements. These are the business requirement assumptions for this test setup:
 1. Users must be authenticated prior to being granted access to the network at large.
 2. Your access is limited based on who the users are. These privileges are mapped to Group Membership in Active Directory. The groups are Guests, Contractors, and Employees.
 3. Based on AD Group Membership, users are placed into a VLAN that has Network Access Privileges that are appropriate for each group.
 4. Guest User traffic continues to be isolated from the rest of the network even after authentication.
 5. After the user is admitted to the network, the NAC Appliance must no longer be in the traffic path. This prevents the NAC Appliance from becoming a bottleneck and allows the network to be used to its full potential by validated users.
- NAC has many capabilities that are not covered by this document. The purpose of this guide is to explore and document the design guidelines and configuration required for a VRF-Lite based Layer 3 Out of Band NAC deployment. This guide does not focus on Posture Assessment or Remediation. More information about the NAC Appliance and its full capabilities can be found at www.cisco.com/go/nac (registered customers only) .

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

Refer to the Cisco Technical Tips Conventions for more information on document conventions.

Configure

Infrastructure Configuration

Introduction:

When considering a VRF-Lite based Layer 3 OOB NAC deployment, there are several design principles that are very important to consider. These principles are listed here, and a brief discussion of their importance is included.

1. **Traffic Classification and Engineering** A key concept to realize and remember for this type of NAC design is that traffic classified as Dirty *must* flow into the UnTrusted side of the NAC Server (CAS). Always keep this principle top of mind during the design of a NAC implementation.

Additionally, Clean and Dirty networks should not be allowed to communicate directly with each other. In a Layer 3 OOB design with VRFs, the NAC server (CAS) acts as the enforcement point or controller that ensures segregation and secure communication between the Clean and Dirty networks.

2. **Traffic Isolation** It is important to be sure that an appropriate enforcement mechanism is selected to provide traffic and path isolation for all traffic sourced from non-authenticated and non-authorized hosts. VRF-Lite is used here to achieve complete data and control-plane isolation (VRF).
3. **Centralized Enforcement** Because the VRF-Lite methodology follows the natural path selection created by routing: topology changes, access control requirements, and/or address changes do not create the need to manipulate ACLs across the infrastructure. If you use a GRE tunnel in conjunction with VRF-Lite, this gives you the flexibility to drop the dirty traffic right in front of the NAC server without the need to configure multiple hops. VRF-Lite in conjunction with GRE only require configuration on Edge Layer 3 devices. This dramatically reduces the number of devices that must be touched in order to provide the path isolation requirement.
4. **Difficulty** Difficulty of implementation as well as ongoing maintenance. When you determine the approach that you are likely to use for NAC Layer 3 OOB in your network, it is important to consider the ease of implementation and ongoing operational cost and complexity of implementing that technology, particularly in a dynamic environment.

Note: The NAC Appliance is oblivious to how traffic is presented to it. In other words, the Appliance itself has no preference whether the traffic arrives through a GRE tunnel, or was re-directed through Policy Based Routing configuration, VRF Routed and so forth.

Note: For the best end-user experience possible, remember to use certificates that are trusted by the browser of the end-user. The use of Self-generated certificates on the NAC Server is not recommended for a production environment.

Note: Always generate the certificate for the NAC Server with the IP Address of its UNTRUSTED interface.

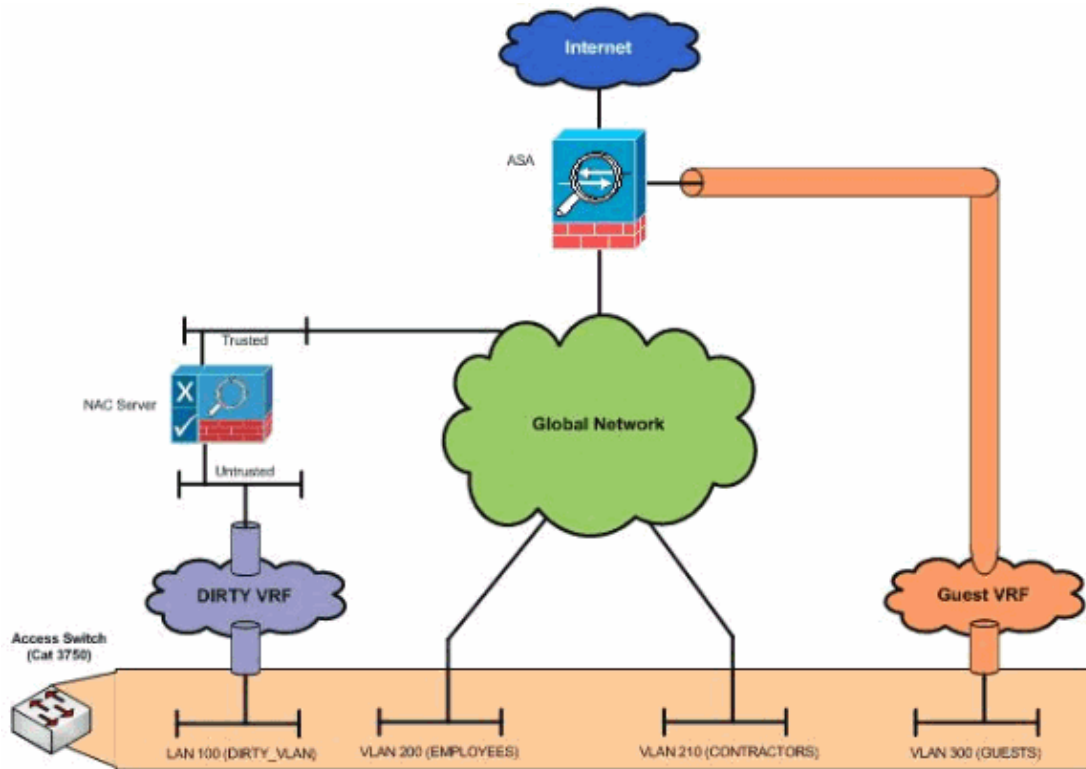
An illustration of device virtualization with VRFs can be seen here. This methodology provides Control Plane and Data Plane for path isolation.



Topology

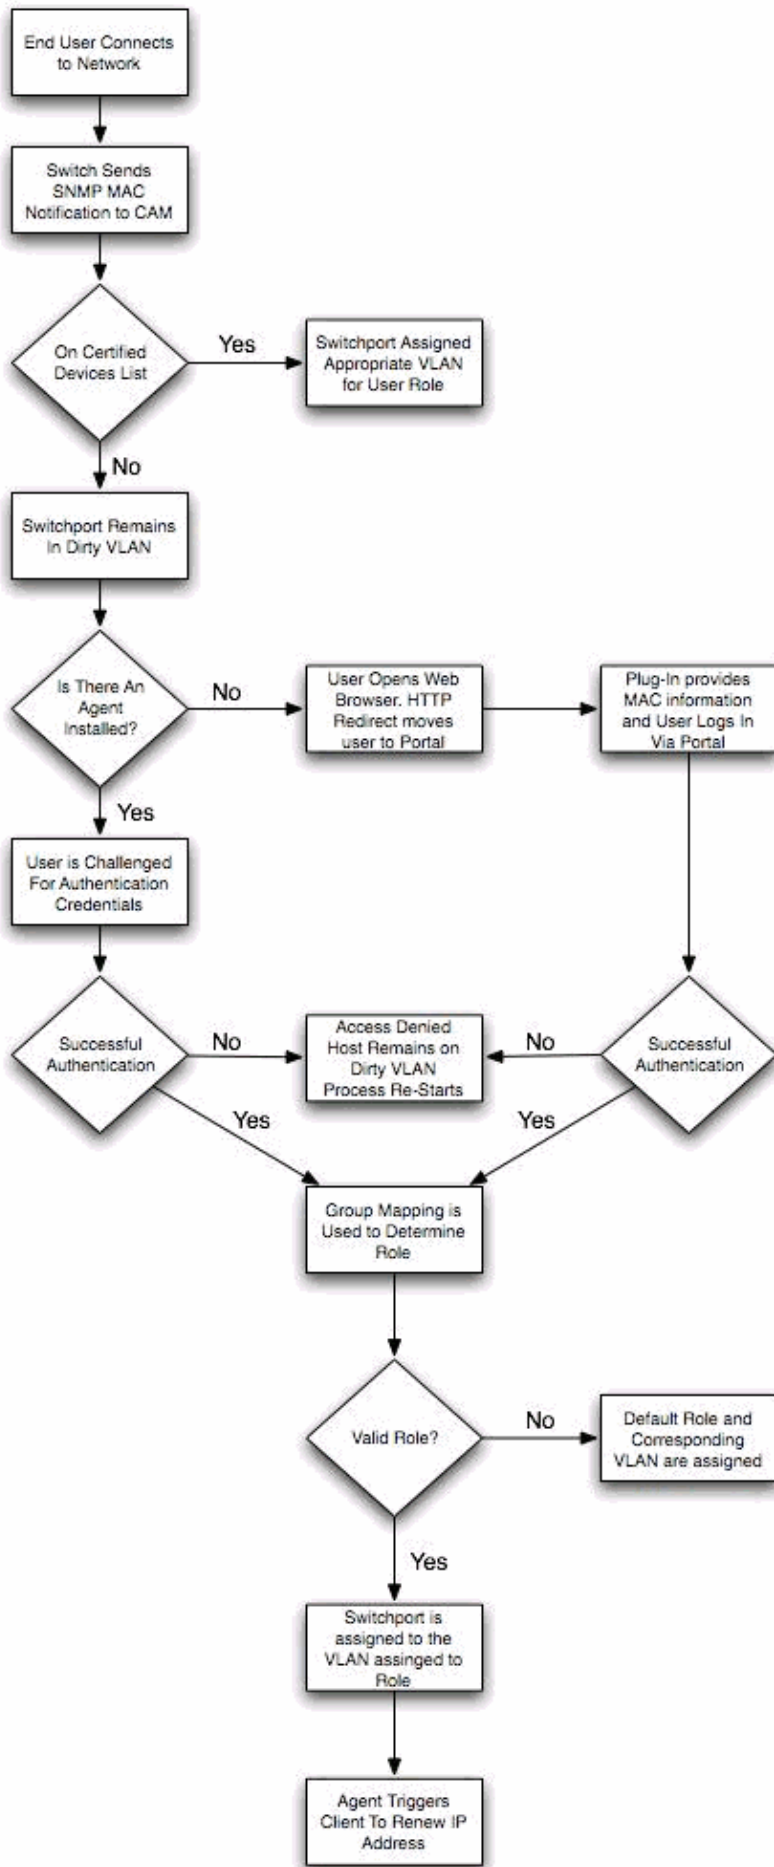
This diagram is representative of the topology used for the creation of this paper. The Internal Network is routing through the Global Routing Table and has no VRF associated with it. The DIRTY VRF contains only the Dirty_VLAN and the associated transit networks that are required to force all data sourced from the DIRTY_VLAN to flow through the Dirty Side of the NAC Appliances. The Guest VRF contains the GUEST_VLAN and associated transit networks required to terminate all data sourced from the

GUEST_VLAN on a separate Sub-Interface on the Firewall. Each of the three Virtual Networks are carried on the same physical infrastructure and provide complete traffic and path isolation respectively.



Process Flows

This section shows the basic process flow of what is required to gain network access both with, and without an agent installed. These process flows are macro-analytical in nature and contain only functional decision steps. They do not include every option or step that occurs and do not include authorization decisions that are based on endpoint assessment criteria.



Configuration

The configuration information details the steps required to configure your network for path isolation using VRF-Lite/GRE and the configuration required for the insertion of the NAC Appliance into your network as a Layer 3 OOB Real IP Gateway.

Note: VRF-lite is a feature that enables you to support two or more Virtual Networks. VRF-lite also allows for overlapping IP Addresses among the Virtual Networks. But, IP Address overlap is not recommended for a NAC implementation because while the infrastructure itself supports the overlapping addresses, it can create troubleshooting complexities and incorrect reporting.

VRF-lite uses input interfaces to distinguish routes for different Virtual Networks and forms virtual packet-forwarding tables by associating one or more Layer 3 interfaces with each VRF. Interfaces in a VRF can be either physical, such as Ethernet ports; or logical, such as sub-interfaces, Tunnel Interfaces or VLAN SVIs. Please note a Layer 3 interface cannot belong to more than one VRF at any time.

Important Considerations For VRF-Lite

- VRF-Lite is only locally significant to the switch where it is defined, and VRF membership is determined by the input interface. No packet header or payload manipulation is performed.
- A switch with VRF-lite is shared by multiple security domains, and all security domains have their own unique routing tables.
- VRF-Lite lets multiple Security Domains share the same physical link between network devices. Trunk ports with multiple VLANs or GRE tunnels provide traffic isolation which separates packets from each different security domain.
- All security domains must have their own VLANs.
- VRF-lite does not support all MPLS-VRF functionality: label exchange, LDP adjacency, or labeled packets.
- The Layer 3 TCAM resource is shared between all VRFs. In order to ensure that any one VRF has sufficient CAM space, use the **maximum routes** command.
- A Catalyst switch using VRF-Lite can support one global network and up to 64 VRFs. The total number of routes supported is limited by the size of the TCAM.
- Most routing protocols (BGP, OSPF, EIGRP, RIP and static routing) can be used between devices that run VRF-Lite.
- There is no need to run BGP with VRF-Lite unless you need to leak routes between VRFs.
- VRF-Lite does not affect the packet switching rate.
- Multicast and VRF-Lite cannot be configured on the same Layer 3 interface at the same time.
- The **capability vrf-lite** subcommand under **router ospf** should be used when you configure OSPF as the routing protocol between network devices.

Defining A VRF

In the design example, the requirements provide path isolation for both unauthenticated or DIRTY users as well as GUESTS. All other traffic is permitted to utilize the internal network. This requires the definition of two VRFs. This is the configuration:

```
!  
ip vrf DIRTY  
  
!--- Names the VRF and places you into VRF Configuration Mode  
  
description DIRTY_VRF_FOR_NAC  
  
!--- Gives the VRF a user friendly description field for documentation
```

```

rd 10:1

!--- Creates a VRF table by specifying a route distinguisher.
!--- Enter either an AS number and an arbitrary number (xxx:y) or an
!--- IP address and arbitrary number (A.B.C.D:y).

!
ip vrf GUESTS
description GUESTS_VRF_FOR_VISITORS
rd 30:1
!

```

Associate a VLAN or Interface with a VRF

After the VRF has been defined on the Layer 3 Switch or Router, the interfaces that participate in the VRF-Lite configuration must be associated with the VRF to which they belong. As mentioned earlier, either physical or virtual interfaces can be associated with a VRF. Included are examples of a physical interface, a switched virtual interface, a sub-interface and a tunnel interface which are associated with a VRF.

```

!
interface FastEthernet0/1
ip vrf forwarding GUESTS
!!Associates the interface with the appropriate VRF defined in Step 1!!
ip address 192.168.39.1 255.255.255.252
!
interface FastEthernet3/1.10
encapsulation dot1Q 10
ip vrf forwarding DIRTY
ip address 192.168.10.1 255.255.255.252
!
interface Vlan100
ip vrf forwarding DIRTY
ip address 192.168.100.1 255.255.255.0
!
interface Tunnel0
ip vrf forwarding GUESTS
ip address 192.168.38.2 255.255.255.252
tunnel source Loopback0
tunnel destination 192.168.254.1
!

```

Extend a VRF Between Two Devices

There are several acceptable methodologies for the extension of a VRF between two pieces of infrastructure. The method you choose should be based on this criteria:

1. Capabilities of Platform In regards to platform capabilities, all current Cisco Layer 3 capable Enterprise Switching and Routing platforms support VRF-Lite. This includes but is not limited to the Catalyst 6500, 4500, 3750, and 3560 platforms.
2. Any routing platform that runs appropriate Cisco IOS®, which include but are not limited to the 7600, 3800, 2800, 1800, and 800 series ISRs.
3. Number of Layer 3 Hops Between Relevant Pieces of Infrastructure Determining the number of Layer 3 hops is critical to keep the deployment as simple as possible. For example, if there were five Layer 3 hops between the infrastructure that host the CAS devices and the clients, it can create administrative overhead.

With the incorrect solution:

1. Layer 2 trunking creates a very suboptimal Layer 2 topology.

2. Layer 3 Sub-Interfaces create many additional interfaces to configure. As a result this can create additional management overhead and potential IP Addressing issues. This is illustrated in the diagram. If you assume that there is no redundancy in the infrastructure, each Layer of the Network shown have both an ingress and egress physical interface. The computation for number of sub-interfaces is then $(2 * \text{number of tiers in the network} * (\text{number of VRFs}))$. In this example there are two VRFs so the formula is $((2*5)*2)$ or 20 Sub-Interfaces. Once redundancy is added this number more than doubles. Compare this to GRE extension, where only four interfaces are required with the same end result. This illustrates plainly how GRE dramatically reduces the configuration impact.

Layer 2 Trunking

Layer 2 trunking is preferred in scenarios where Layer 3 closets are not deployed or where the Network Devices do not support GRE or Sub-Interfaces. It should be noted that the Catalyst 3560, 3750 and 4500 platforms do not support Sub-Interfaces. The Catalyst 3560, and 3750 also do not support GRE. The Catalyst 4500 supports GRE in software, and the Catalyst 6500 supports GRE in hardware.

In a Layer 3 closet model where you connect a platform that does not support Sub-Interfaces or GRE to a platform that does, it is preferred to only use Layer 2 trunking on one side, and to use Sub-Interfaces on the other side. This allows you to maintain all the benefits of a Layer 3 closet architecture, and still overcome the limitation of no GRE or Sub-Interface support on some platforms. One of the primary advantages of the configuration of a Layer 2 trunking only on one side of the link is that Spanning Tree is not introduced back into the Layer 3 environment. See the example where a 3750 Access Switch (No GRE or Sub-Interface Support) is connected to a 6500 Distribution Switch, which does support GRE and Sub-Interfaces.

3750 Relevant Configuration:

In this configuration, note that on FastEthernet 1/0/1 the default setting for the NATIVE VLAN is VLAN 1. This configuration has not been changed. You also notice, however, that VLAN 1 is not allowed to be trunked across the link. The allowed VLANs is limited to only the VLANs that are tagged. Because in this Layer 3 topology there is no need for trunk negotiation, or VTP traffic to go from switch to switch, there is also no need for unencapsulated traffic to transit this link. This configuration increases the security posture of the architecture since it doesn't open up unnecessary Layer 2 security holes.

```
!  
ip vrf DIRTY  
description DIRTY_VRF_FOR_NAC  
rd 10:1  
!  
ip vrf GUESTS  
description GUESTS_VRF_FOR_VISITORS  
rd 30:1  
!  
!  
interface FastEthernet1/0/1  
description CONNECTION_TO_DISTRIBUTION_6504  
switchport trunk encapsulation dot1q  
switchport trunk allowed vlan 10,20,30  
switchport mode trunk  
speed 100  
duplex full  
!  
!  
interface Vlan10  
description DIRTY_VRF_TRANSIT  
ip vrf forwarding DIRTY  
ip address 192.168.10.2 255.255.255.252  
!
```



```

interface Vlan20
description CLEAN_TRANSIT
ip address 192.168.20.2 255.255.255.252
!
interface Vlan30
description GUESTS_VRF_TRANSIT
ip vrf forwarding GUESTS
ip address 192.168.30.2 255.255.255.252
!

```

6500 Relevant Configuration:

In this configuration, note that dot1q encapsulation is used and the frames with VLAN 10, 20 and 30 are tagged. When you choose the VLAN tags to use, you cannot use a VLAN number that is already defined locally in the VLAN database on the switch.

```

!
ip vrf DIRTY
description DIRTY_VRF_FOR_NAC
rd 10:1
!
ip vrf GUESTS
description GUESTS_VRF_FOR_VISITORS
rd 30:1
!
interface FastEthernet3/1
description CONNECTION_TO_3750_ACCESS
no ip address
speed 100
duplex full
!
!
interface FastEthernet3/1.10
description DIRTY_VRF_TRANSIT
encapsulation dot1Q 10
ip vrf forwarding DIRTY
ip address 192.168.10.1 255.255.255.252
!
interface FastEthernet3/1.20
description CLEAN_TRANSIT
encapsulation dot1Q 20
ip address 192.168.20.1 255.255.255.252
!
interface FastEthernet3/1.30
description GUESTS_VRF_TRANSIT
encapsulation dot1Q 30
ip vrf forwarding GUESTS
ip address 192.168.30.1 255.255.255.252
!

```

Layer 3 Sub-Interfaces

Layer 3 Sub-Interfaces are a good option when you only need to extend the VRF over one Layer 3 Hop in the network. Either GRE or Sub-Interfaces can be chosen based on your level of comfort with each configuration. This is a sample configuration for a Layer 3 Sub-Interface:

```

!
interface FastEthernet3/1
description CONNECTION_TO_3750_ACCESS
no ip address
speed 100
duplex full

```

```

!
!
interface FastEthernet3/1.10
description DIRTY_VRF_TRANSIT
encapsulation dot1Q 10
ip vrf forwarding DIRTY
ip address 192.168.10.1 255.255.255.252
!

```

GRE Tunnels

GRE Tunnels are the preferred method to extend a VRF-Lite VRF when there are multiple Layer 3 hops between the clients who need to access the VRF. This type of design is more common with remote branch office NAC where the remote clients want to access a centrally located NAC server. For example, in a typical Core, Distribution, Access network model clients are not directly connected to the Distributions or to the Core. Therefore, there is no need to add the complexity of VRF definition on the Distribution or Core devices. GRE can be used to simply transport the traffic that needs to be isolated to the point in the Network where the NAC servers are connected. This is an example of a GRE tunnel Interface.

```

!
interface Tunnel0
ip vrf forwarding GUESTS
ip address 192.168.38.2 255.255.255.252
tunnel source Loopback0
tunnel destination 192.168.254.1
!

```

Configuring Routing for the VRF

As discussed earlier in the document, VRF-Lite supports BGP, OSPF, and EIGRP. In this configuration example, EIGRP is chosen because it is typically the Cisco recommended routing protocol implemented on Campus networks where fast convergence is required.

It should be noted, that OSPF works equally well with VRF-Lite, as does BGP.

It should also be noted that if the design requires that traffic should be leaked between VRFs, then BGP is required.

This is an example of the configuration of the Routing for a VRF with EIGRP.

```

!
!--- As with any configuration this is base routing protocol
!--- configuration which handles the routing for the Global Routing Table.

router eigrp 1
network 192.168.20.0 0.0.0.3
network 192.168.21.0
network 192.168.22.0
network 192.168.28.0 0.0.0.3
network 192.168.29.0 0.0.0.3
network 192.168.254.1 0.0.0.0
no auto-summary
!

!--- An Address Family must be defined for each VRF
!--- that is to be routing through the routing protocol.
!--- Routing Protocol options such as auto-summarization,
!--- autonomous system number, router id, and so forth are all

```

```

!--- configured under the address family. Note that EIGRP does not
!--- neighbor without the autonomous system specified under
!--- the address family. Also note, that this autonomous system
!--- number should be unique for each VRF and should not be
!--- the same as the Global AS number.

!
address-family ipv4 vrf GUESTS
network 192.168.30.0 0.0.0.3
network 192.168.38.0 0.0.0.3
no auto-summary
autonomous-system 30
exit-address-family
!
address-family ipv4 vrf DIRTY
network 192.168.10.0 0.0.0.3
network 192.168.11.0
no auto-summary
autonomous-system 10
exit-address-family
!


```

Routing Traffic Between The Global Routing Table and the Dirty VRF

It depends on the NAC deployment requirements if it can be necessary to pass traffic from the Untrusted or Dirty side of the network to the Trusted or Clean side of the network. For example, remediation services can potentially live on the Trusted side of the NAC Appliance. In the case of Active Directory Single Sign on deployments, it is necessary to pass a subset of traffic to Active Directory to allow Interactive Logons, Kerberos Ticket Exchange, and so forth. In any event, it is very important that the Global Routing Table knows how to reach the Dirty VRF, and that the DIRTY VRF knows how to reach the Global Routing Table if any data needs to pass between the two. This is typically handled by this methodology.

The Dirty VRF defaults to the Untrusted or Dirty interface of the NAC Appliance. The Global has Static Routes *only* to the subnets that are considered DIRTY VLAN s.

Consider this drawing.



The first Layer 3 hop on the Untrusted or Dirty side of the NAC Appliance redistributes a default route into routing process that points to the NAC Appliance. The first Layer 3 hop on the Trusted or Clean side of the NAC Appliance redistributes a static route for the subnet that belongs to VLAN 100, which in this case is 192.168.100.0/24.

Note: The first Layer 3 hop on opposite sides of the NAC Appliance can be on the same physical device, but in different VRFs. In the next example, the Untrusted or Dirty side of the NAC server is in a VRF, while the Trusted or Clean side of the NAC Appliance remains in the Global Routing table.

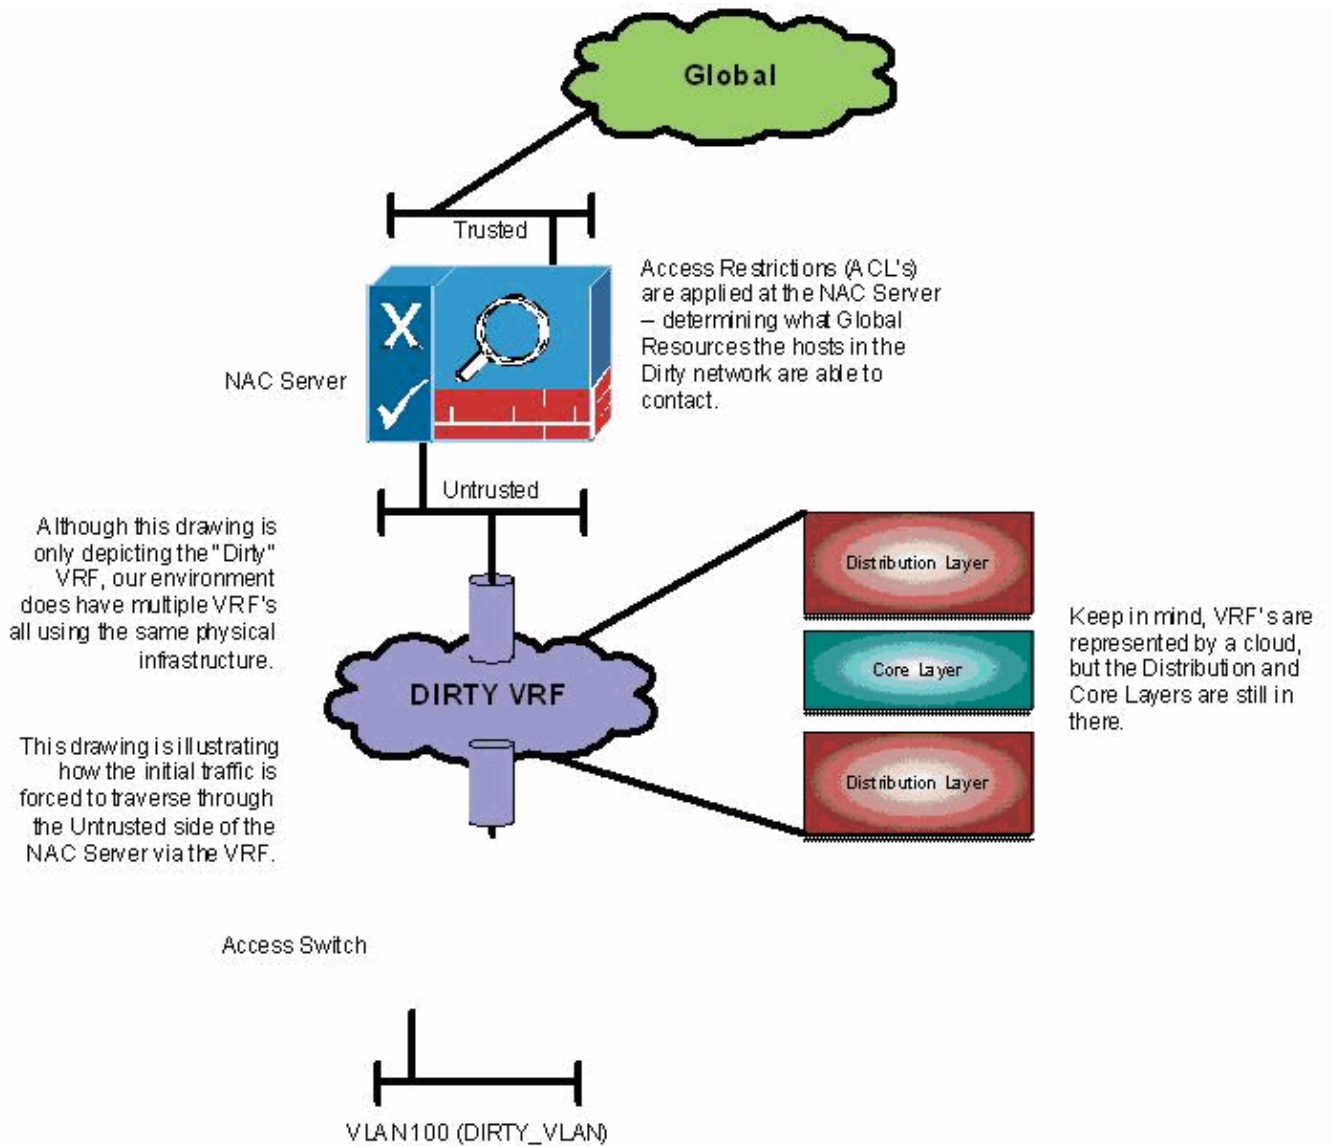
The configuration is as follows:

```
!  
router eigrp 1  
  redistribute static  
  network 192.168.20.0 0.0.0.3  
  network 192.168.21.0  
  network 192.168.22.0  
  network 192.168.28.0 0.0.0.3  
  network 192.168.29.0 0.0.0.3  
  network 192.168.254.1 0.0.0.0  
  no auto-summary  
!  
address-family ipv4 vrf GUESTS  
  network 192.168.30.0 0.0.0.3  
  network 192.168.38.0 0.0.0.3  
  no auto-summary  
  autonomous-system 30  
exit-address-family  
!  
address-family ipv4 vrf DIRTY  
  redistribute static  
  network 192.168.10.0 0.0.0.3  
  network 192.168.11.0  
  no default-information out  
  no auto-summary  
  autonomous-system 10  
exit-address-family  
!  
ip classless  
ip route 192.168.100.0 255.255.255.0 192.168.21.10  
ip route vrf DIRTY 0.0.0.0 0.0.0.0 192.168.11.2  
!  
!!
```

NAC Configuration for Layer 3 OOB

CAS Setup

Remember our number one principle from the Introduction section: The Trick to a successful NAC design is to always remember that traffic classified as Dirty *must* flow into the UnTrusted side of the NAC Server (CAS).



In the first screen shot, pay attention to the NAC Server network setup. You notice the Server is deployed as an Out-of-Band Real-IP Gateway. Note that the default route of the NAC Server is pointed to the TRUSTED side.

Device Management > Clean Access Servers > 192.168.21.10

Status Network Filter Advanced Authentication Misc

IP DHCP DNS Certs

Clean Access Server Type: Out-of-Band Real-IP Gateway

Enable L3 support

Enable L3 strict mode to block NAT devices with Clean Access Agent

Enable L2 strict mode to block L3 devices with Clean Access Agent

Platform: APPLIANCE

Trusted Interface (to protected network)		Untrusted Interface (to managed network)	
IP Address	192.168.21.2	IP Address	192.168.11.2
Subnet Mask	255.255.255.0	Subnet Mask	255.255.255.0
Default Gateway	192.168.21.1	Default Gateway	192.168.21.1
<input type="checkbox"/> Set management VLAN ID:	0	<input type="checkbox"/> Set management VLAN ID:	0

(Make sure the Clean Access Server is on VLAN n before you set its management VLAN ID to n.)

Update Reboot

The Server needs to be configured with Static Routes for each of the DIRTY VLANS that exist on the UNTRUSTED side. See the second screen shot.

Device Management > Clean Access Servers > 192.168.21.10

Status	Network	Filter	Advanced	Authentication	Misc	
Managed Subnet	VLAN Mapping	NAT	1:1 NAT	Static Routes	ARP	Proxy

Dest. Subnet Address/Mask: /

Gateway (optional):
(gateway should be the address of an external gateway for the dest. subnet, not of the Clean Access Server)

Link:

Description:

Subnet	Gateway	Link	Description	Delete
192.168.100.0 / 255.255.255.0	192.168.11.1	Untrusted		X

Verify

Find the documented process of the user NAC–Employee logging into our network. Cisco has captured the activity from the Access Switch, the workstation, and shows information from the routing tables of the Distribution Switches.

Use this section to confirm that your configuration works properly.

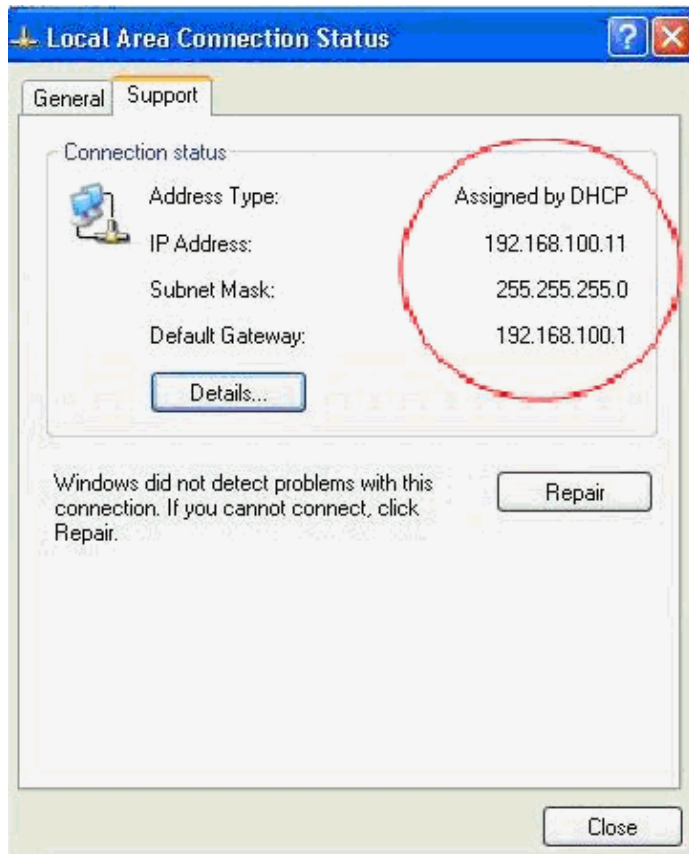
The Output Interpreter Tool (registered customers only) (OIT) supports certain **show** commands. Use the OIT to view an analysis of **show** command output.

Stage 1 You have not connected to the network yet, and the switchport on the Access Switch is down.

```
! Catalyst 3750 Access Switch
!
!--- Note: Client machine is off the network at this point.
!
3750-Access#show int status | i Fa1/0/13
Fa1/0/13 CLIENT_CONNECTION notconnect 100 auto auto 10/100BaseTX
!
!
3750-Access#!Notice it is in the "notconnect" state.
!
```

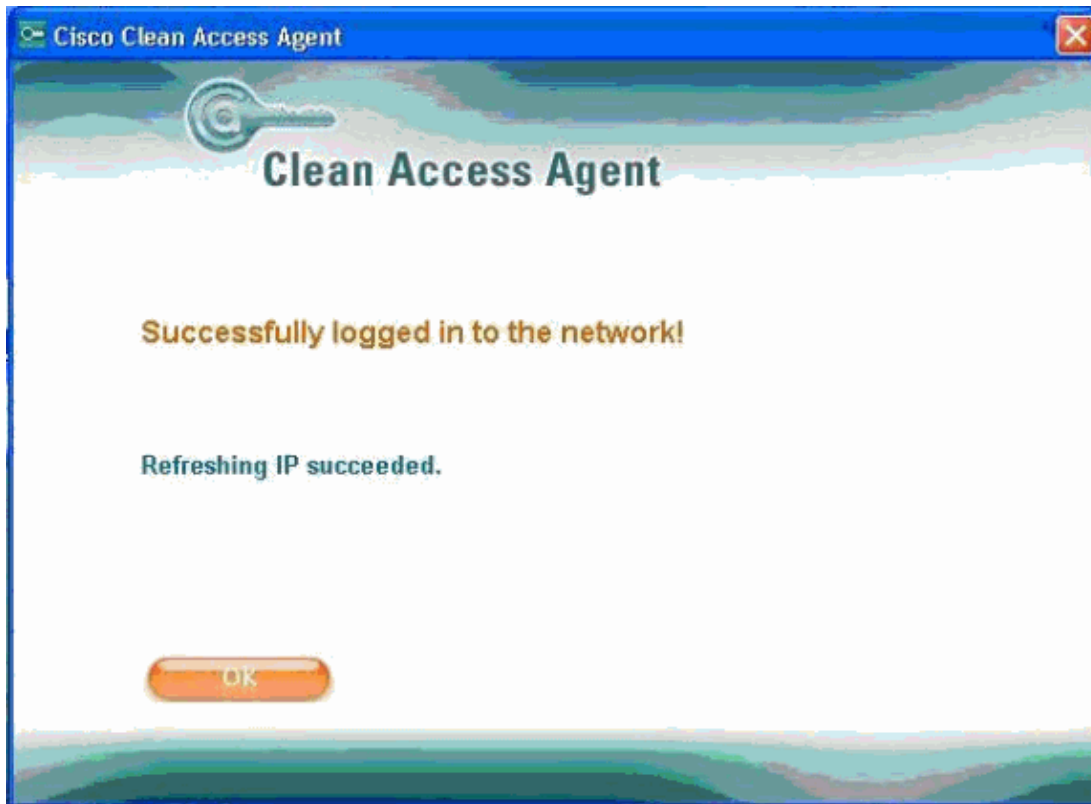
Stage 2 The Windows Client plugs into the network, and the initial VLAN on the Switch is VLAN 100 (the Dirty VLAN). An IP Address is assigned to the host, as you can see in this screen shot.

```
! Catalyst 3750 Access Switch
!
!--- Note: Client just connected to the network.
!
2w5d: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan100, changed state to up
2w5d: %LINK-3-UPDOWN: Interface FastEthernet1/0/13, changed state to up
2w5d: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet1/0/13, changed state to up
!
!
3750-Access#show int status | i Fa1/0/13
Fa1/0/13 CLIENT_CONNECTION connected 100 a-full a-100 10/100BaseTX
```



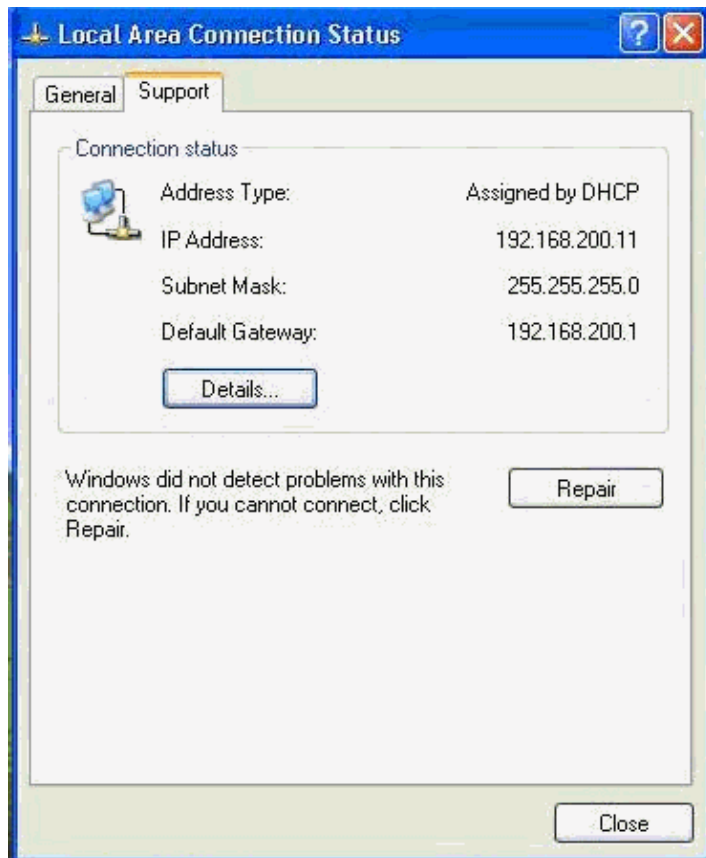
Stage 3 Within a few seconds, the NAC Agent begins its logon process. In this example, Active Directory Single-Sign-On is configured, so you are not prompted for a username and password. Instead, you see a pop-up window that describes that Single-Sign-On occurs.

After the Authentication and Posture Assessment has been completed, a Success message is displayed, the switchport is moved from the Dirty VLAN to the Employee VLAN and the NAC Agent refreshes the IP Address of the PC.




```
! Catalyst 3750 Access Switch
2w5d: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan100, changed state to down
2w5d: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan200, changed state to up
!
!--- Note: As you can tell from the previous messages,
!--- the switchport was just moved from VLAN 100 to VLAN 200.
!
3750-Access#show int status | i Fa1/0/13
Fa1/0/13 CLIENT_CONNECTION connected 200 a-full a-100 10/100BaseTX
!
!
```

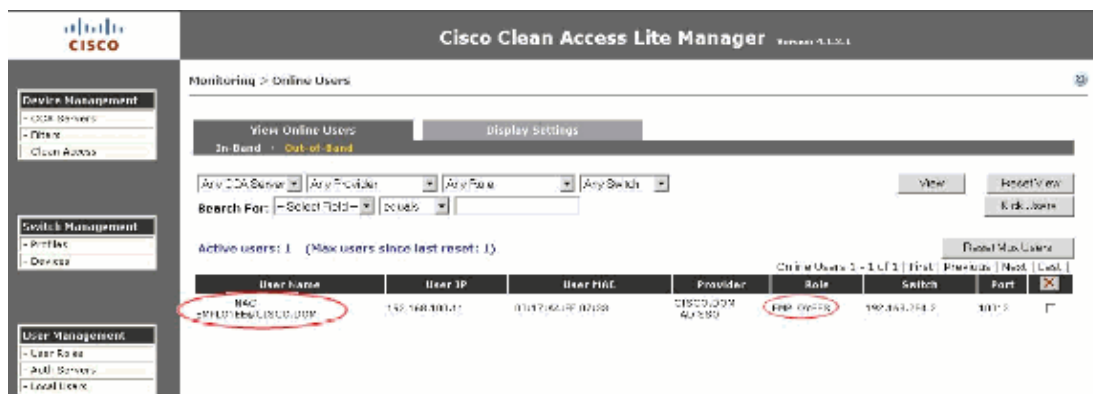
This screen shot shows the Final IP address, which is in the Employee VLAN (VLAN 200).



This screen shot shows the device of the NAC–Employee user as listed in the Certified Devices List. The Role is assigned to *EMPLOYEES* and the VLAN is 200.



This screen shot shows the Online Users list on the NAC Manager.



This is the NAC Manager event log, which shows the successful login of the out-of-band user.



In this section, the routing tables of the Global Route Table and the DIRTY VRF are examined. In the first screen capture, note the **show ip route** command. This indicates that you see the routing table for the Global Routes.

```
6504-DISTRIBUTION#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 192.168.28.2 to network 0.0.0.0

192.168.29.0/30 is subnetted, 1 subnets
D    192.168.29.0 [90/30720] via 192.168.28.2, 2w5d, FastEthernet3/48
192.168.28.0/30 is subnetted, 1 subnets
C    192.168.28.0 is directly connected, FastEthernet3/48
D EX 192.168.31.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
D EX 192.168.30.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
D    192.168.200.0/24 [90/28416] via 192.168.20.2, 6d19h, FastEthernet3/1.20
D EX 192.168.38.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
C    192.168.21.0/24 is directly connected, Vlan21
D EX 192.168.39.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
192.168.20.0/30 is subnetted, 1 subnets
C    192.168.20.0 is directly connected, FastEthernet3/1.20
D EX 192.168.36.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
C    192.168.22.0/24 is directly connected, Vlan22
D EX 192.168.37.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
D EX 192.168.34.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
192.168.254.0/32 is subnetted, 3 subnets
D    192.168.254.2 [90/156160] via 192.168.20.2, 2w5d, FastEthernet3/1.20
D    192.168.254.3 [90/156160] via 192.168.28.2, 2w5d, FastEthernet3/48
C    192.168.254.1 is directly connected, Loopback0
D EX 192.168.35.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
D EX 192.168.32.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
S    192.168.100.0/24 [1/0] via 192.168.21.10
D EX 192.168.33.0/24 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
D*EX 0.0.0.0/0 [170/30976] via 192.168.28.2, 2w5d, FastEthernet3/48
```

Note: The 192.168.100.0/24 network (the Dirty Network) is in the routing table as a static route, with the next-hop being the Trusted Interface of the NAC Server.

Note the **show ip route vrf DIRTY** command. This indicates that you see the routing table for the DIRTY virtual network only.

```

6504-DISTRIBUTION#show ip route vrf DIRTY

Routing Table: DIRTY
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 192.168.11.2 to network 0.0.0.0

    192.168.10.0/30 is subnetted, 1 subnets
C       192.168.10.0 is directly connected, FastEthernet3/1.10
C       192.168.11.0/24 is directly connected, Vlan11
D       192.168.100.0/24
          [90/28416] via 192.168.10.2, 01:03:19, FastEthernet3/1.10
S*     0.0.0.0/0 [1/0] via 192.168.11.2

```

Note: Note the Dirty Access VLAN (192.168.100.0/24) is learned in the distribution through EIGRP from the 3750 Access Switch, only in the DIRTY VRF Routing Table. This route does not exist in the Global table.

Appendix A: Switch Configurations

Access Switch Running Configuration

```

service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname 3750-Access
!
!
no aaa new-model
clock timezone EST -5
clock summer-time EST recurring
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip vrf DIRTY
  description DIRTY_VRF_FOR_NAC
  rd 10:1
!
ip vrf GUESTS
  description GUESTS_VRF_FOR_VISITORS
  rd 30:1
!
!
!
crypto pki trustpoint TP-self-signed-819048320
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-819048320
  revocation-check none
  rsa-keypair TP-self-signed-819048320
!
!
crypto ca certificate chain TP-self-signed-819048320
  certificate self-signed 01
!
!
no file verify auto

```

```
spanning-tree mode pvst
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
!
interface Loopback0
 ip address 192.168.254.2 255.255.255.255
!
!
interface FastEthernet1/0/1
 description CONNECTION_TO_DISTRIBUTION_6504
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,20,30
 switchport mode trunk
 speed 100
 duplex full
!
interface range FastEthernet1/0/2 - 24
 description CLIENT_CONNECTION
 switchport access vlan 100
 switchport mode access
 spanning-tree portfast
 spanning-tree bpduguard enable
!
!- SNIP -
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan10
 description DIRTY_VRF_TRANSMIT
 ip vrf forwarding DIRTY
 ip address 192.168.10.2 255.255.255.252
!
interface Vlan20
 description CLEAN_TRANSIT
 ip address 192.168.20.2 255.255.255.252
!
interface Vlan30
 description GUESTS_TRANSIT
 ip vrf forwarding GUESTS
 ip address 192.168.30.2 255.255.255.252
!
interface Vlan100
 description DIRTY_VLAN
 ip vrf forwarding DIRTY
 ip address 192.168.100.1 255.255.255.0
 ip helper-address 192.168.22.11
!
interface Vlan200
 description EMPLOYEES_VLAN
 ip address 192.168.200.1 255.255.255.0
 ip helper-address 192.168.22.11
!
interface Vlan210
 description CONTRACTORS_VLAN
 ip address 192.168.210.1 255.255.255.0
 ip helper-address 192.168.22.11
!
!
interface Vlan300
 description GUESTS
```

```

ip vrf forwarding GUESTS
ip address 192.168.31.1 255.255.255.0
!
router eigrp 1
network 192.168.20.0 0.0.0.3
network 192.168.200.0
network 192.168.254.2 0.0.0.0
no auto-summary
!
address-family ipv4 vrf GUESTS
network 192.168.30.0 0.0.0.3
network 192.168.31.0
no auto-summary
autonomous-system 30
exit-address-family
!
address-family ipv4 vrf DIRTY
network 192.168.10.0 0.0.0.3
network 192.168.100.0
no auto-summary
autonomous-system 10
exit-address-family
!
router bgp 1
no synchronization
bgp log-neighbor-changes
neighbor 192.168.254.3 remote-as 1
neighbor 192.168.254.3 update-source Loopback0
no auto-summary
!
ip classless
ip route 192.0.2.1 255.255.255.255 Null0
ip http server
ip http secure-server
!
!
snmp-server community NIC-NAC-PADDYWHACK RW
snmp-server user NIC-NAC-PADDYWHACK NIC-NAC-PADDYWHACK v1
snmp-server user NIC-NAC-PADDYWHACK NIC-NAC-PADDYWHACK v2c
snmp-server trap-source Loopback0
snmp-server host 192.168.22.5 version 2c NIC-NAC-PADDYWHACK
!
!- SNIP
!
ntp clock-period 36028450
ntp source Loopback0
ntp server 192.168.254.1 version 2 prefer
end

```

Distribution Switch Running Configuration

```

! SNIP -
!
hostname 6504-DISTRIBUTION
!
boot-start-marker
boot system disk0:s72033-advipservicesk9_wan-mz.122-33.SXH2a.bin
boot-end-marker
!
!
no aaa new-model
clock timezone EST -5
clock summer-time EST recurring
!
!- SNIP -

```

```
!  
ip vrf DIRTY  
  description DIRTY_VRF_FOR_NAC  
  rd 10:1  
!  
ip vrf GUESTS  
  description GUESTS_VRF_FOR_VISITORS  
  rd 30:1  
!  
ipv6 mfib hardware-switching replication-mode ingress  
vtp domain cmpd  
vtp mode transparent  
no mls acl tcam share-global  
mls netflow interface  
no mls flow ip  
no mls flow ipv6  
mls cef error action freeze  
!  
!  
redundancy  
  keepalive-enable  
  mode sso  
  main-cpu  
    auto-sync running-config  
spanning-tree mode pvst  
no spanning-tree optimize bpdu transmission  
spanning-tree extend system-id  
diagnostic cns publish cisco.cns.device.diag_results  
diagnostic cns subscribe cisco.cns.device.diag_commands  
!  
vlan internal allocation policy ascending  
vlan access-log ratelimit 2000  
!  
!  
!  
!  
vlan 11  
  name CAS_DIRTY  
!  
vlan 21  
  name CAS_CLEAN  
!  
vlan 22  
  name SERVER_VLAN  
!  
interface Tunnel0  
  ip vrf forwarding GUESTS  
  ip address 192.168.38.1 255.255.255.252  
  tunnel source Loopback0  
  tunnel destination 192.168.254.3  
!  
interface Loopback0  
  ip address 192.168.254.1 255.255.255.255  
!  
!- SNIP -  
!  
interface FastEthernet3/1  
  description CONNECTION_TO_3750_ACCESS  
  no ip address  
  speed 100  
  duplex full  
!  
interface FastEthernet3/1.10  
  description DIRTY_VRF_TRANSIT  
  encapsulation dot1Q 10
```

```
ip vrf forwarding DIRTY
ip address 192.168.10.1 255.255.255.252
ip verify unicast source reachable-via rx allow-default
!
interface FastEthernet3/1.20
description CLEAN_TRANSIT
encapsulation dot1Q 20
ip address 192.168.20.1 255.255.255.252
!
interface FastEthernet3/1.30
description GUESTS_TRANSIT
encapsulation dot1Q 30
ip vrf forwarding GUESTS
ip address 192.168.30.1 255.255.255.252
!
!
!
!
!
interface FastEthernet3/2
description CAS1_DIRTY
switchport
switchport access vlan 11
switchport mode access
speed 100
duplex full
spanning-tree portfast
spanning-tree bpduguard enable
!
interface FastEthernet3/3
description CAS2_DIRTY
switchport
switchport access vlan 11
switchport mode access
speed 100
duplex full
spanning-tree portfast
spanning-tree bpduguard enable
!
interface FastEthernet3/4
description CAS1_CLEAN
switchport
switchport access vlan 21
switchport mode access
speed 100
duplex full
spanning-tree portfast
spanning-tree bpduguard enable
!
interface FastEthernet3/5
description CAS2_CLEAN
switchport
switchport access vlan 21
switchport mode access
speed 100
duplex full
spanning-tree portfast
spanning-tree bpduguard enable
!
interface FastEthernet3/6
description CAM
switchport
switchport access vlan 22
switchport mode access
speed 100
```

```
duplex full
spanning-tree portfast
spanning-tree bpduguard enable
!
!
!- SNIP -
!
!
!
interface FastEthernet3/48
description CONNECTION_TO_THE_WORLD
ip address 192.168.28.1 255.255.255.252
speed 100
duplex full
!
interface Vlan1
no ip address
shutdown
!
interface Vlan11
description NAC_DIRTY
ip vrf forwarding DIRTY
ip address 192.168.11.1 255.255.255.0
!
interface Vlan21
description NAC_CLEAN
ip address 192.168.21.1 255.255.255.0
!
interface Vlan22
description SERVER_VLAN
ip address 192.168.22.1 255.255.255.0
!
router eigrp 1
redistribute static
network 192.168.20.0 0.0.0.3
network 192.168.21.0
network 192.168.22.0
network 192.168.28.0 0.0.0.3
network 192.168.29.0 0.0.0.3
network 192.168.254.1 0.0.0.0
no auto-summary
!
address-family ipv4 vrf GUESTS
network 192.168.30.0 0.0.0.3
network 192.168.38.0 0.0.0.3
no auto-summary
autonomous-system 30
exit-address-family
!
address-family ipv4 vrf DIRTY
redistribute static
network 192.168.10.0 0.0.0.3
network 192.168.11.0
no default-information out
no auto-summary
autonomous-system 10
exit-address-family
!
!
!
!
!
!
router bgp 1
```



```
no synchronization
bgp log-neighbor-changes
neighbor 192.168.254.3 remote-as 1
neighbor 192.168.254.3 update-source Loopback0
no auto-summary
!
ip classless
ip route 192.0.2.1 255.255.255.255 Null0
ip route 192.168.100.0 255.255.255.0 192.168.21.10
ip route vrf DIRTY 0.0.0.0 0.0.0.0 192.168.11.2
!
!
!- SNIP -
!
ntp source Loopback0
ntp master 2
!
end
```

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2013 – 2014 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Dec 03, 2008

Document ID: 108540
