Contents

Physical View

Lines may be added between the access layer and the aggregation layer as available for etherchannel connections. Recommended protocol is LACP.

Acronyms

CSU – Channel Service Unit
DSU – Data Service Unit
IDPS – Intrusion Detection/Prevention System
LACP – Link Aggregation Control Protocol
VRF – Virtual Routing and Forwarding

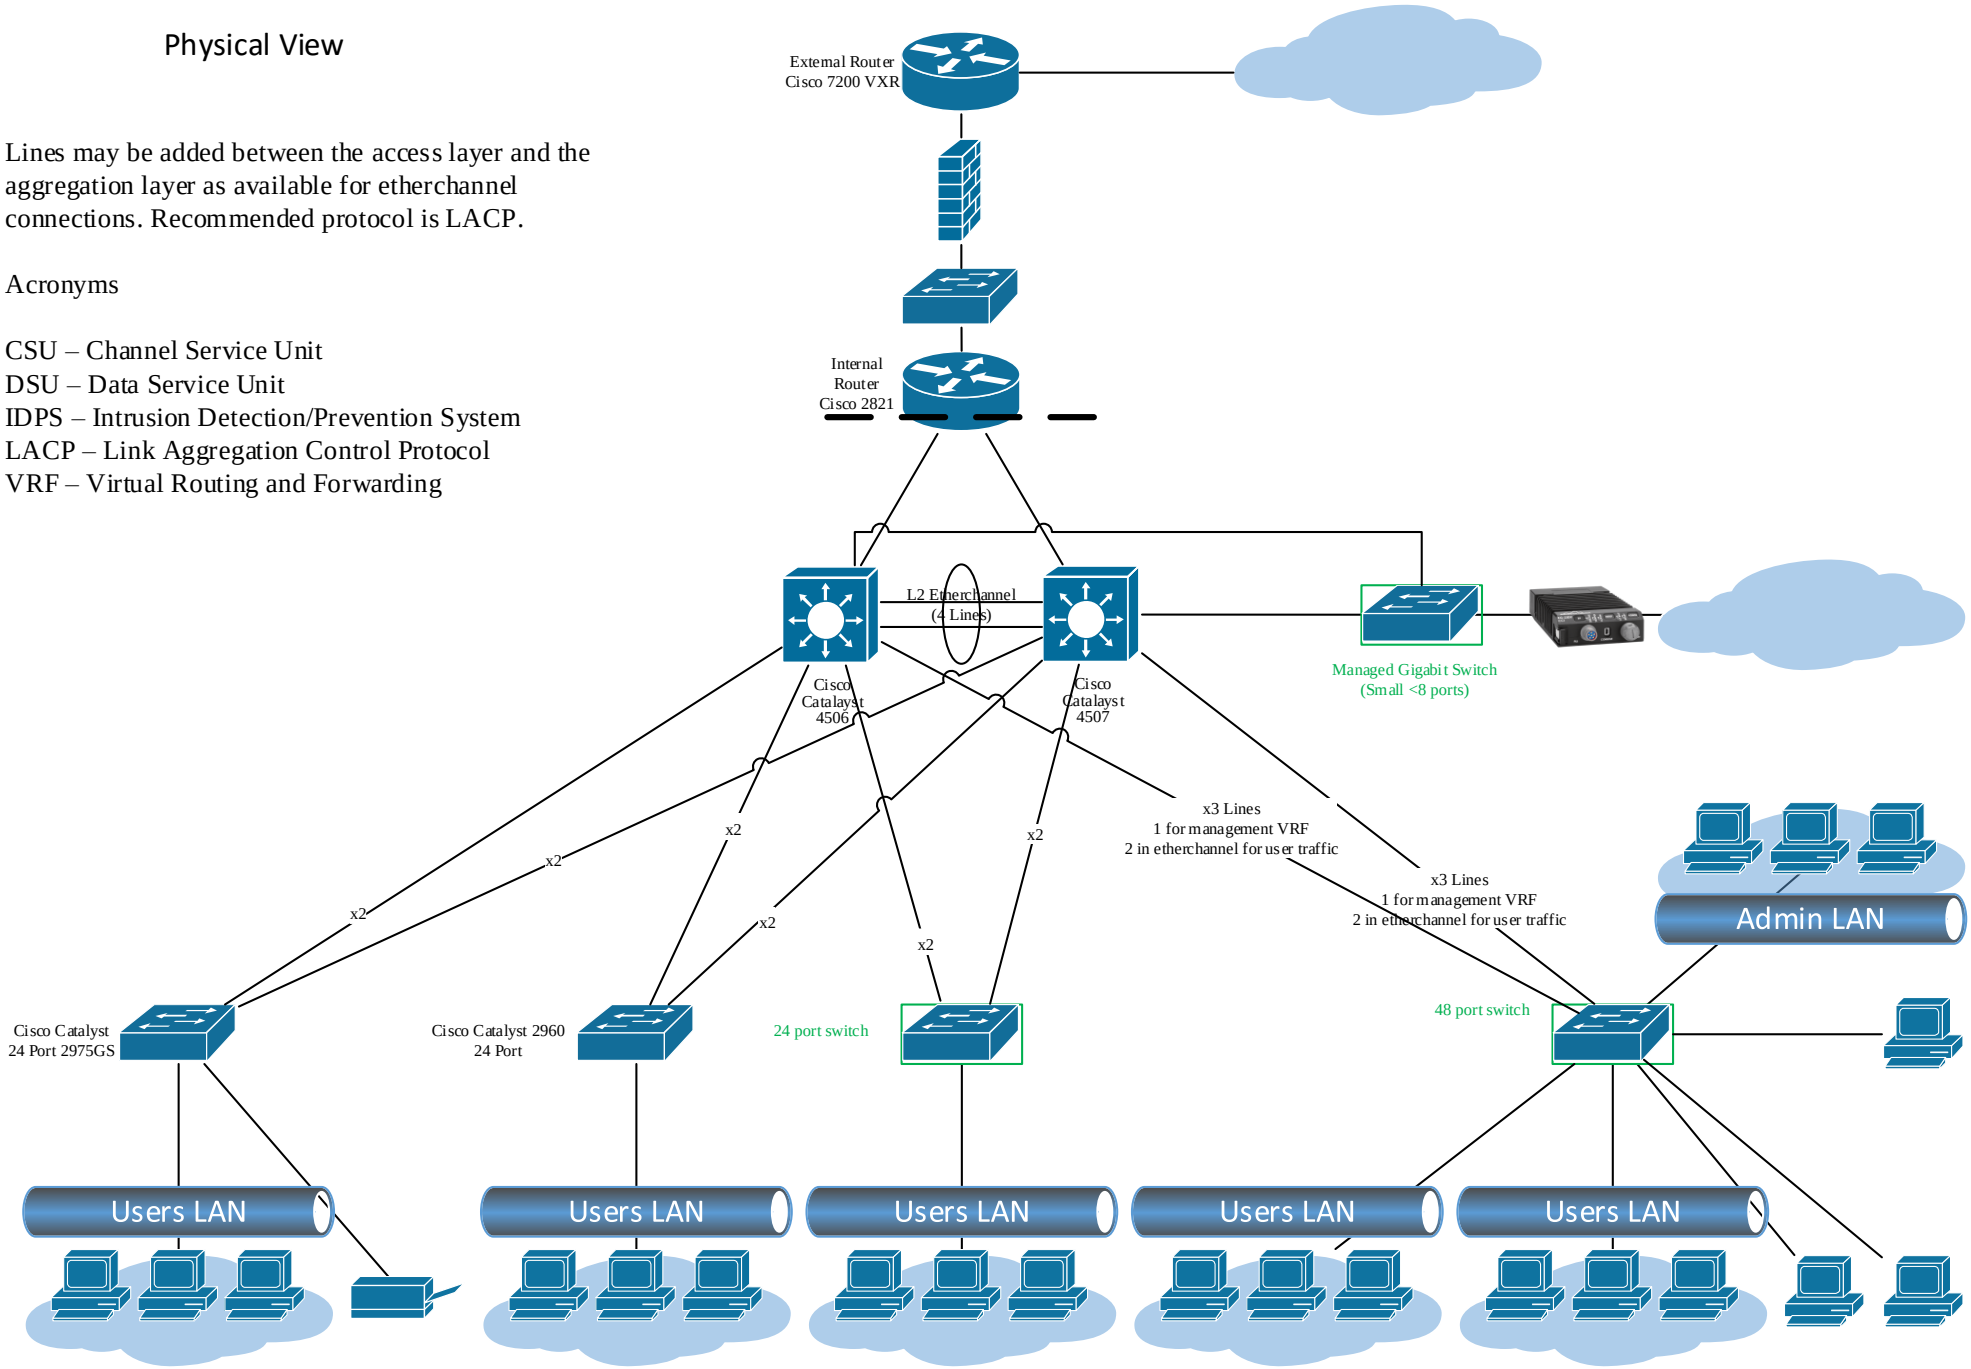

WARNING

I modified this document from its original format to meet releasability requirements. Most of it can stand on its own, but some wording may seem out of place/unusual or some context may be missing.

SECTION 1 - PHYSICAL DESIGN

This section covers the design aspects of the physical layer of the network.

1.1. Heirarchical Design

1.1.1. Creating a Distinct Access Layer

In a collapsed core model, end user devices connect to the access layer switches which in turn connect to both core switches. This permits either core switch to go down and user traffic can load balance immediately to the other core switch without incident. Without the access switches users must connect directly to the core which creates a single point of failure because a failure of the switch results in loss of connectivity to the users. The figure below illustrates the collapsed core model.

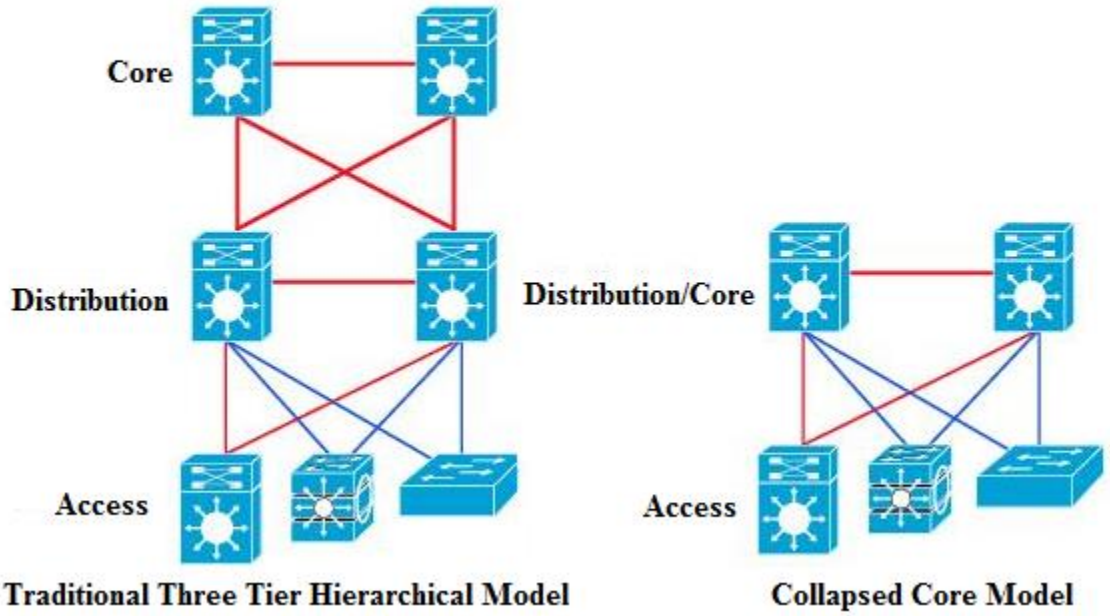

The traditional three tier model uses a separate distribution and core layer. This design is used in larger networks. The collapsed core removes the distribution layer and combines it with the core layer. The collapsed core model is used in smaller network topologies where the bandwidth requirements do not warrant a separate distribution layer.

Creating a collapsed core topology and adhering to a hierarchical network model carries with it several benefits:

- **Scalability:** Hierarchical networks provide modularity. Modularity allows you to replicate design elements as the network grows. The consistency of each instance of the module makes expansion easy to plan and implement.
- **Redundancy:** Redundancy at the core and distribution layers ensures path availability in case of any hardware failure at any one layer.
- **Performance:** Link aggregation between layers combined with high-performance core and distribution layer switches allows for near wire speed throughout the network.
- **Security:** Network heirarchy creates opportunities for a more sound security architecture. The access layer can provide port level security features, filtering, and tagging and the distribution layer can enforce security policies based on actions taken at the access layer.
- **Manageability:** Consistency between switches at each level makes management simpler. Each layer of the hierarchical design performs specific functions that are consistent throughout that layer. A network manager may easily make a change on one access layer switch and safely replicate it throughout the rest of the access layer.
- **Maintainability:** The modular nature of hierarchical networks allows them to scale easily while adhering to correct network design. This simplifies and standardizes maintenance. Without a heirarchy, manageability becomes increasingly complicated as the network grows. In the hierarchical model, switch functions are isolated at each layer making devices easier to track, replace, and upgrade.

1.1.2. Creating Redundancy

If configured according to the above design, Rapid Spanning-Tree Protocol (RSTP) and Hot Standby Routing Protocol (HSRP) will converge quickly enough to prevent disruption of user traffic due to the loss of the second core switch.

Network managers may assign each core switch primary responsibilities for a subset of Virtual Local Area Networks (VLAN). The traffic for that VLAN then flows only through that switch reducing network congestion. Each switch may also provide backup services to the opposite switch for the other switch’s VLANs. A failure of a single switch only causes its assigned VLANs to go down and the opposite switch may instantly assume responsibility of those VLANs.

1.2. Cabling

1.2.1. EtherChannel

EtherChannel allows a network device to utilize multiple cables as if they were a single cable attached to a single interface on either side. EtherChannel is critical for any switch with servers attached. In the organization's case, servers are attached to the network in multiple locations. At this juncture, a migration to a full server farm is not practical. EtherChannel provides the additional speed and bandwidth necessary to avoid network congestion to those access blocks housing servers. Additionally, a cable which fails in an EtherChannel bundle does not cause a failure of the entire link. The link continues to operate normally sans the failed cable. Link Aggregation Protocol (LACP) is the recommended EtherChannel protocol. Packet Aggregation Protocol (PAgP) is Cisco proprietary and only works with other Cisco devices. While the network currently utilizes a Cisco infrastructure, LACP provides flexibility for the future.

Layer 2 View

-Spanning Tree Type: RPVST+
-UDLD in aggressive mode on all fiber links
-BPDU Guard and Port fast on all client ports
-Blocking STP port for a given VLAN should be the interswitch link connected to the backup STP root.
-HSRP must have preemption enabled and a delay longer than the boot time should be added to prevent premature preempting
-Native VLAN must not be 1
-Servers should use a separate interface and cable for management whenever possible in addition to NIC teaming interfaces
-Additional VLANs should be put in place as necessary to further segregate server traffic
-Stateful Switchover and Non-Stop Forwarding should be configured on the Cisco 4506 and 4507 if two supervisors are available

Acronyms

BPDU – Bridge Protocol Data Unit
HSRP – Hot Standby Routing Protocol
STP – Spanning-Tree Protocol
UDLD – Unidirectional Link Detection
VLAN – Virtual Local Area Network
VoIP – Voice Over Internet Protocol

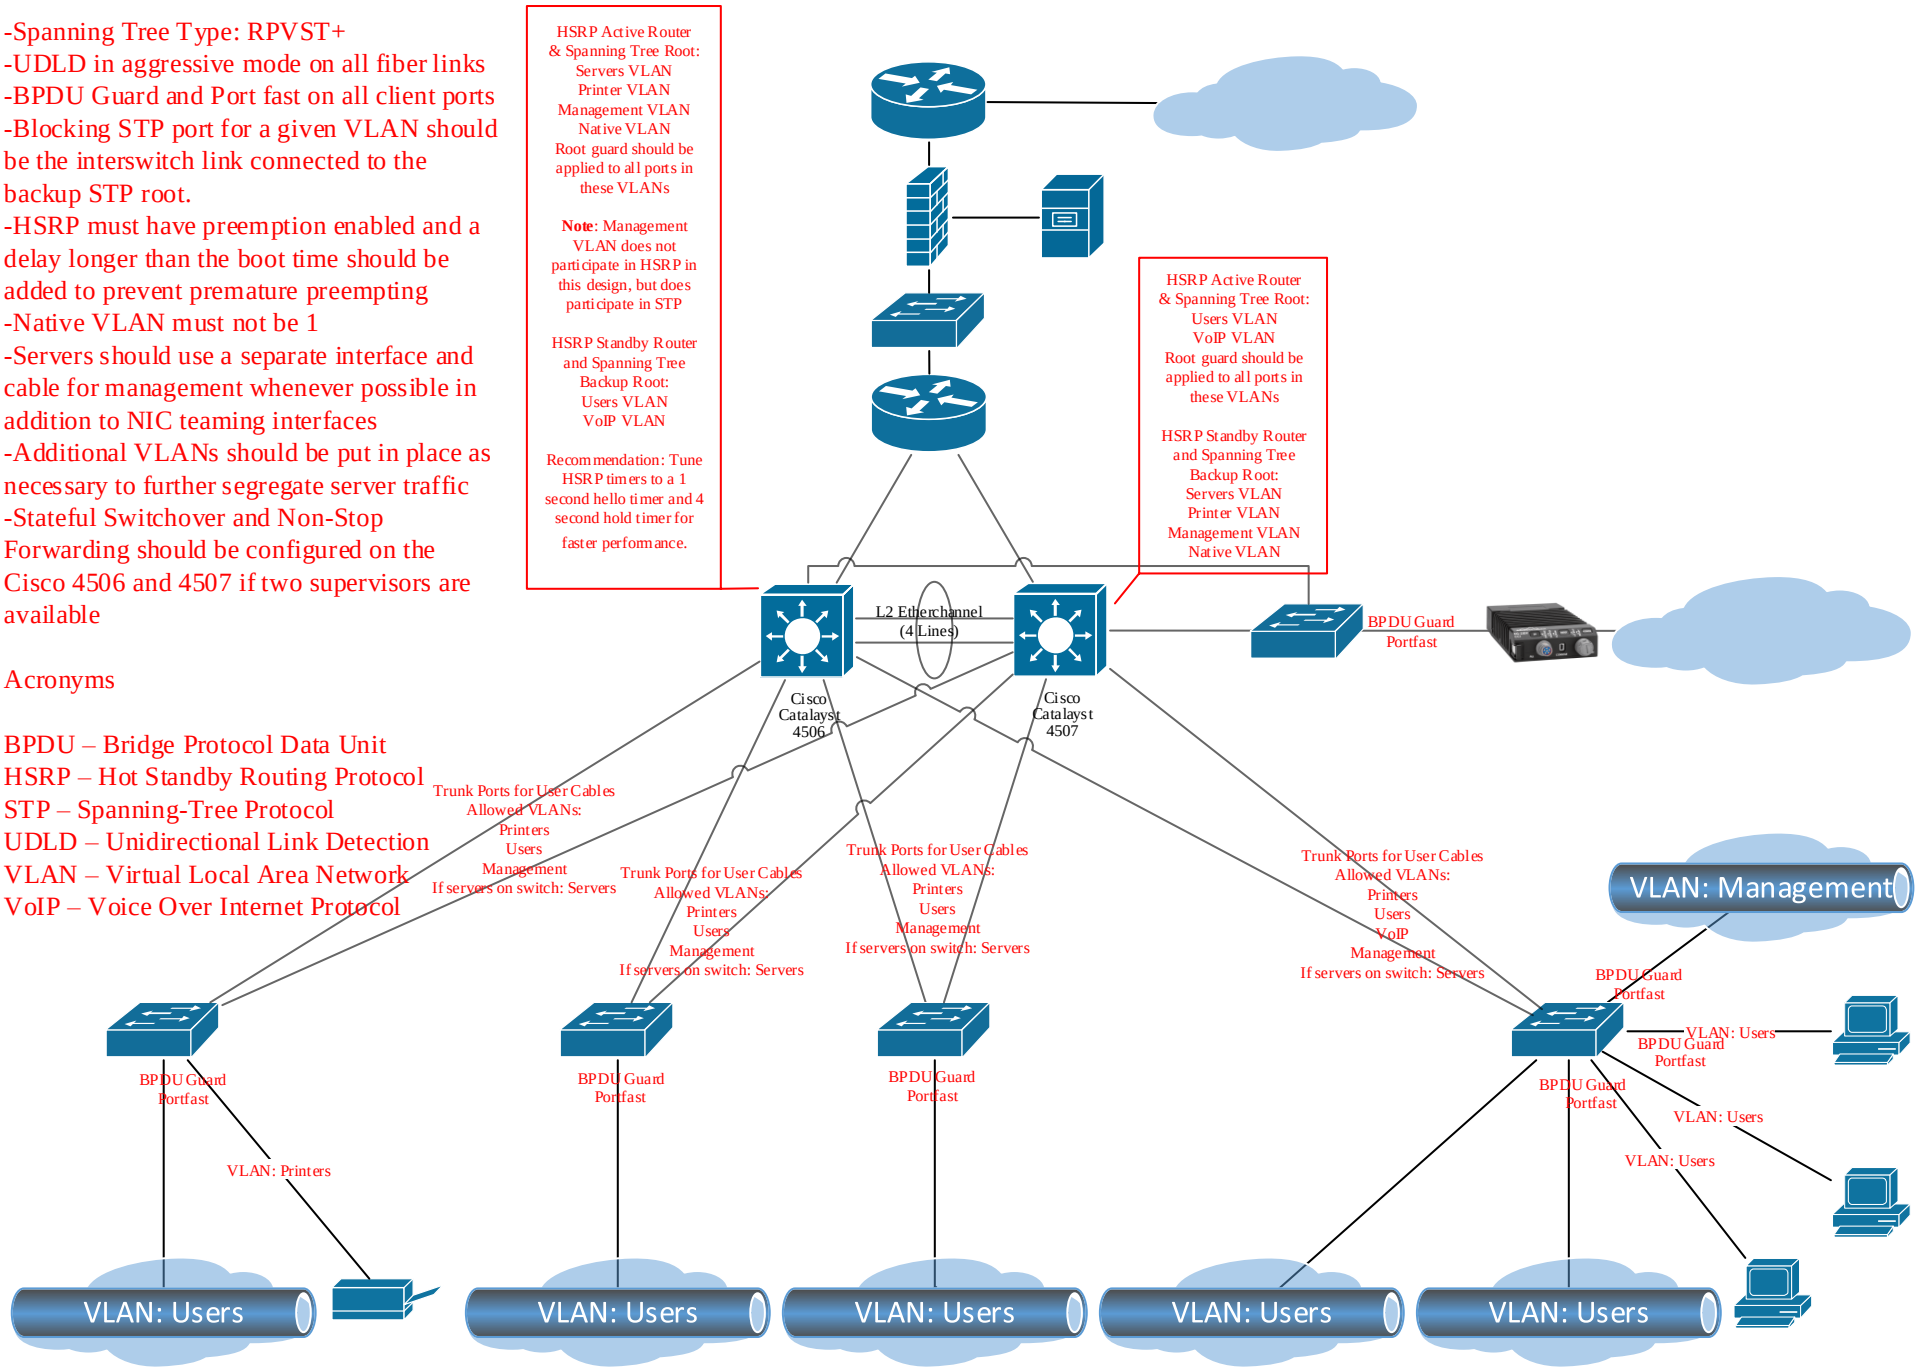

SECTION 2 – LAYER 2 DESIGN

This section covers the design aspects of the data link layer of the network.

2.1. Layer 2 Topology

The following sections discuss the layer 2 topology the design calls for.

2.1.1. Layer 2 Topology Type

The design specifies a layer 2 looped topology. This necessitates a layer 2 connection between the two core switches. This is not currently the industry recommended design because it introduces a reliance on spanning tree for convergence however, it is required in order to support VLANs spanning the access layer.

2.1.2. Spanning-Tree Type

PVST is an outdated protocol that takes up to 50 seconds to converge in the event of a failure. Such a long convergence time causes active user sessions to drop during a convergence event. RPVST+ converges in less than 6 seconds in most cases, which is less than the time it takes for the average TCP session to drop. As a result the network recovers quickly enough to prevent connectivity loss to the end users. Aside from measures to reduce convergence time RPVST+ also includes a number of spanning tree enhancements that potentially remove the need for any recalculation and drop the convergence time to near zero. For example, RPVST+ includes a variant of uplinkfast, which causes RPVST+ to immediately failover to a backup uplink in the event of a failure without additional configuration.

2.1.3 Spanning-Tree Tools

Spanning tree tools used include portfast, Bridge Protocol Data Unit guard (BPDU), and root guard. These tools perform a variety of tasks to include improving network speed, protecting the network topology, and performing security functions.

2.1.3.1 PortFast

Portfast allows ports to skip all phases of spanning-tree and move directly into the forwarding state. Network administrators should enable PortFast on any interface connecting to an end device to reduce the time required to obtain network connectivity. PortFast should not run on an interface that may connect to a switch or receive BPDUs.

2.1.3.2 BPDU Guard

BPDU Guard works in conjunction with PortFast. Any port running PortFast should not receive Spanning-Tree Protocol (STP) BPDUs. Received BPDUs on an interface configured with PortFast are usually an indicator of a rogue switch attaching to the network. An interface with BPDU Guard enabled moves to the disabled state if it receives a BPDU. BPDU guard prevents an adversary from masquerading as a switch in order to sniff traffic or a user connecting a rogue switch to the network.

2.1.3.3 Root Guard

Root Guard enforces the STP topology. RPVST+ runs on a per-VLAN basis and each of the core switches maintains the STP root for their respective VLANs. Root Guard exists to prevent a rogue switch with a lower bridge priority from joining the network and assuming root responsibility for a VLAN. A rogue switch acting in this way at a minimum causes poor traffic flow and could allow an adversary to sniff all traffic destined for a VLAN. Root Guard runs on any STP designated port and moves the port to an STP inconsistent state if a superior BPDU is received on the port. This prevents another switch from taking root from our specifically designated switches.

2.1.3.4 Unidirectional Link Detection

Unidirectional Link Detection (UDLD) shuts down any link only forwarding in a single direction. While a single uplink may fail in a UTP cable, UDLD is particularly useful for fiber connections. Many fiber connections have separate uplink and downlink cables. Wiring mistakes frequently cause these connections to only forward in a single direction. The interface may show as up and running, but in reality may not function. These errors are difficult to detect and UDLD provides the means to find them. UDLD uses bidirectional hello packets to detect links only forwarding in a single direction and shut them down. Once the interface is shut down a backup may take control and the incident reported either through a logging mechanism or Simple Network Management Protocol (SNMP).

2.2 Create VLANs

VLANs provide a means to create logical clusters of machines and limit broadcast propagation. The VLANs provided in the design are role based rather than geographically based. This allows network administrators to apply security controls against a role on the network. For example, network administrators may restrict access to the management interface of a server to only nodes residing within the management VLAN, which should only be accessible by network administrators. The included VLAN topology also provides the ability to prevent user to user communication. In most networks there are very few reasons one client should connect directly to another client. Limiting client to client communication improves network security by preventing an adversary from using one client machine to gather information from other client machines.

VLANs also provide a performance benefit by creating opportunities for load balancing. The displayed topology includes two core switches instead of one. Each switch takes responsibility for only a subset of the total VLANs. Traffic for different VLANs load balances across both switches reducing congestion.

The design shows a generic VLAN plan with a "Servers" VLAN. NET-VLAN-010 stipulates that, "Network traffic with differing security policies within the server farm should be logically grouped using multiple VLANs." Additional VLANs must be created to further split the servers VLAN to fufill this requirement.

2.3 Miscellaneous

A VTP domain should be configured on all switches and the VTP mode should be set to transparent. The VTP domain should be set in the event that a switch is inadvertently left in server or client mode. If the VTP domain is not set an attacker could push their own VLAN topology to the switch and the switch would accept it.

The switchport host macro should be run on all host ports. The switchport host command sets the DTP switchport mode to access, enables spanning-tree PortFast, and disables channel group (EtherChannel).

DTP mode should be set on trunking ports to either desirable on both sides or the distribution layer to desirable and the access layer to auto.

Required by STIGs: NET0986, NET0987, NET0988, NET0989, NET0991, NET0992, NET0994, NET0995, NET0996, NET1002, NET1003, NET1004, NET-VLAN-004, NET-VLAN-005, NET-VLAN-006, NET-VLAN-008, NET-VLAN-009, NET-VLAN-010, NET-SRVFRM-003, NET-SRVFRM-004

L3 Interfaces

Acronym List

SVI – Switch Virtual Interface

Note: Management VLAN does not have or require an SVI

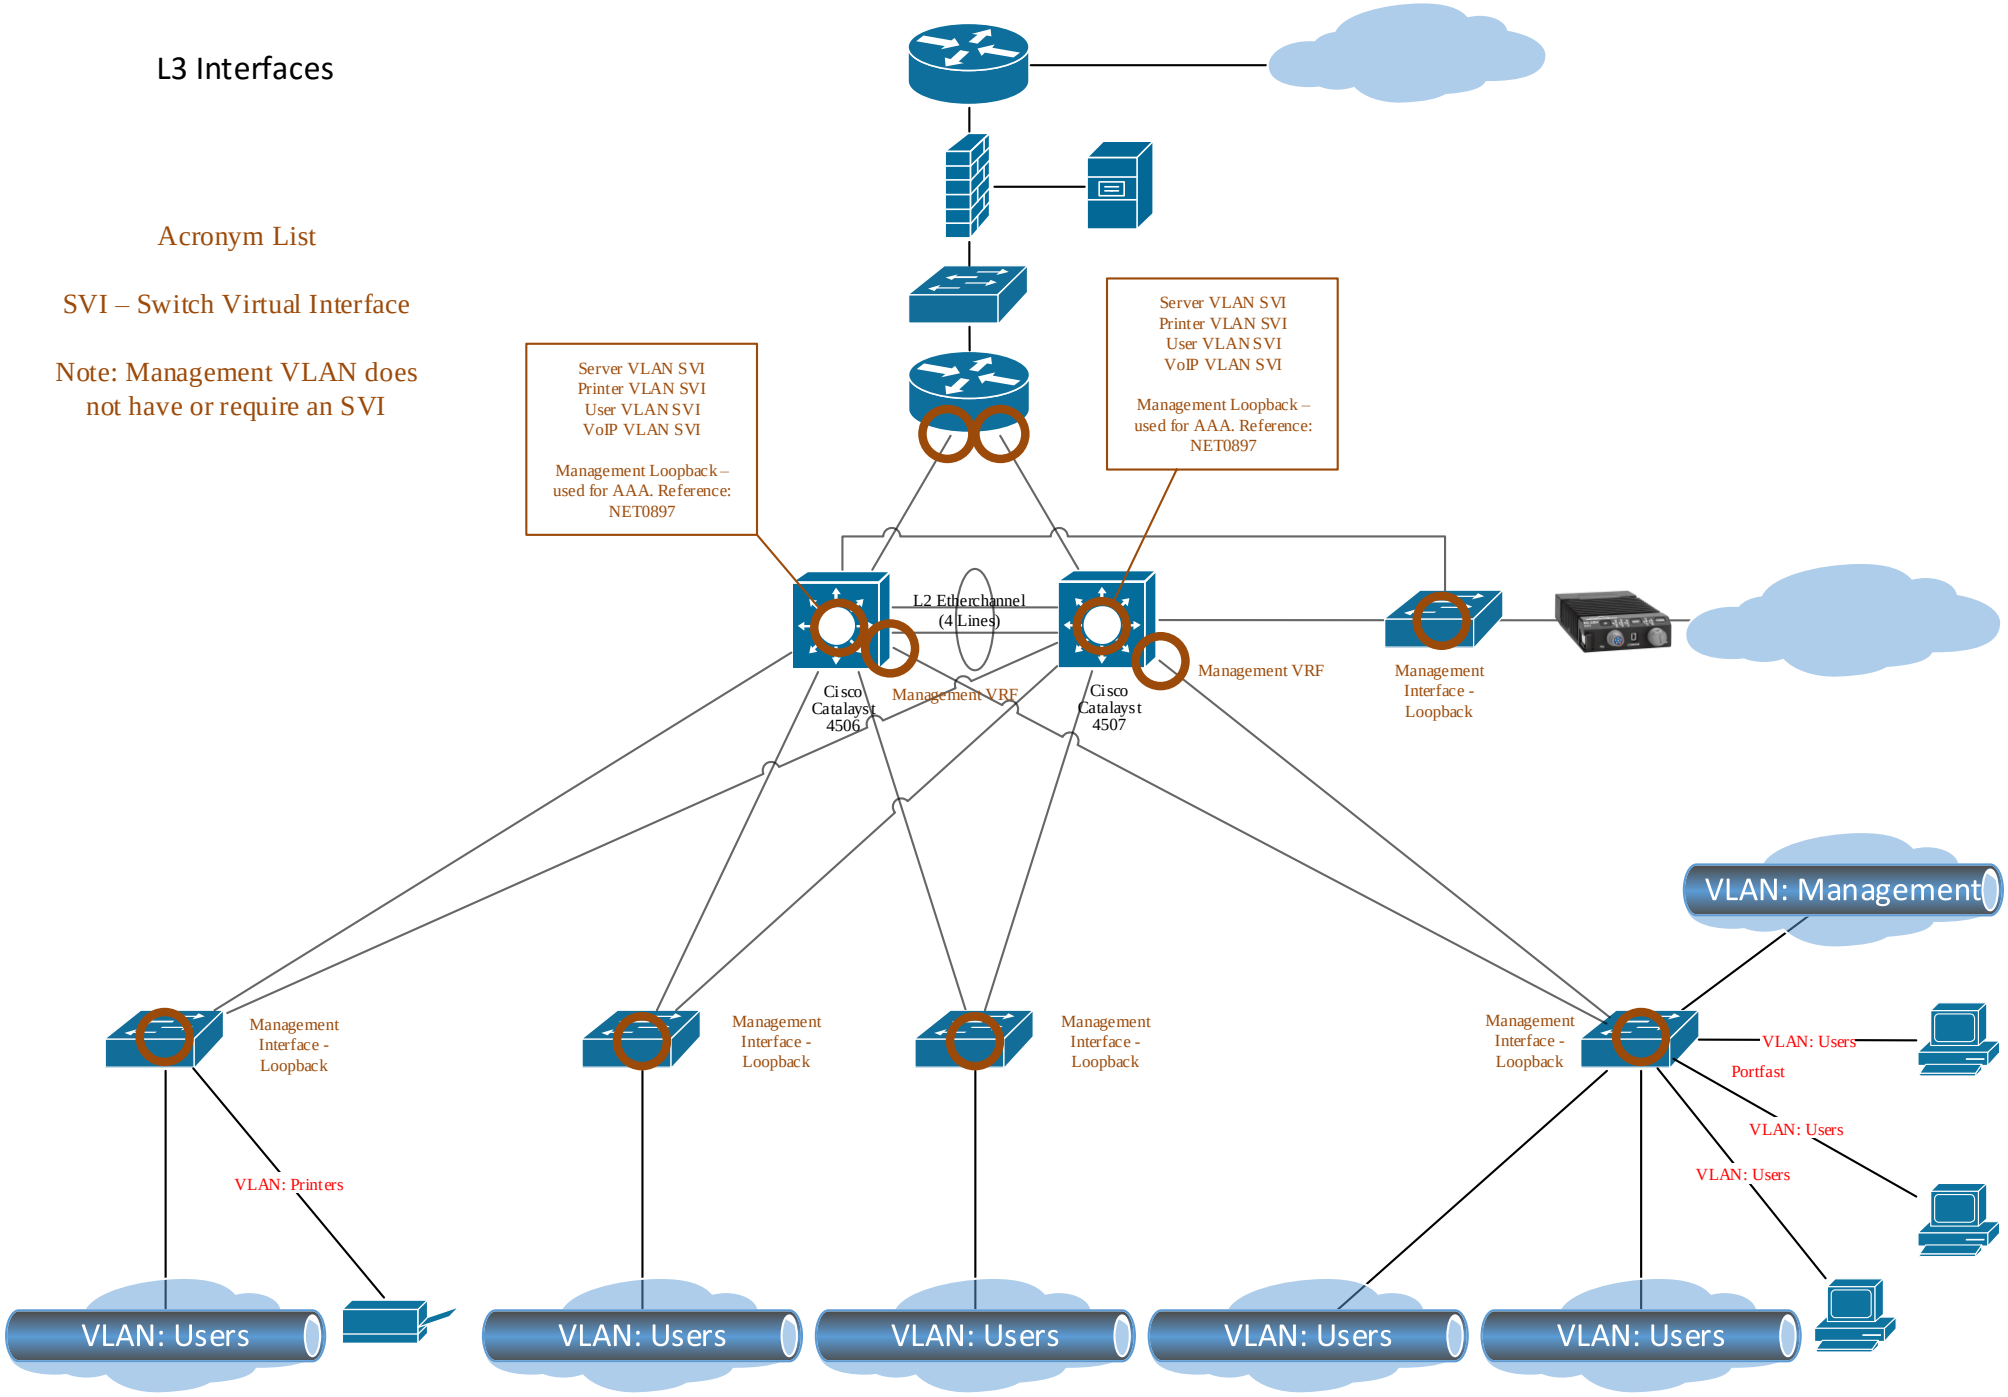

Layer 3 View – Administrative Traffic Flow

-Clarification: Management VRF lines are separate cables and interfaces leading to Fa 1 on the core switches. Fa 1 essentially acts as a L3 console port. Another cable must be used for the out of band management network

Suggestion: Source transit management traffic on the two core switches from the loopback interface.

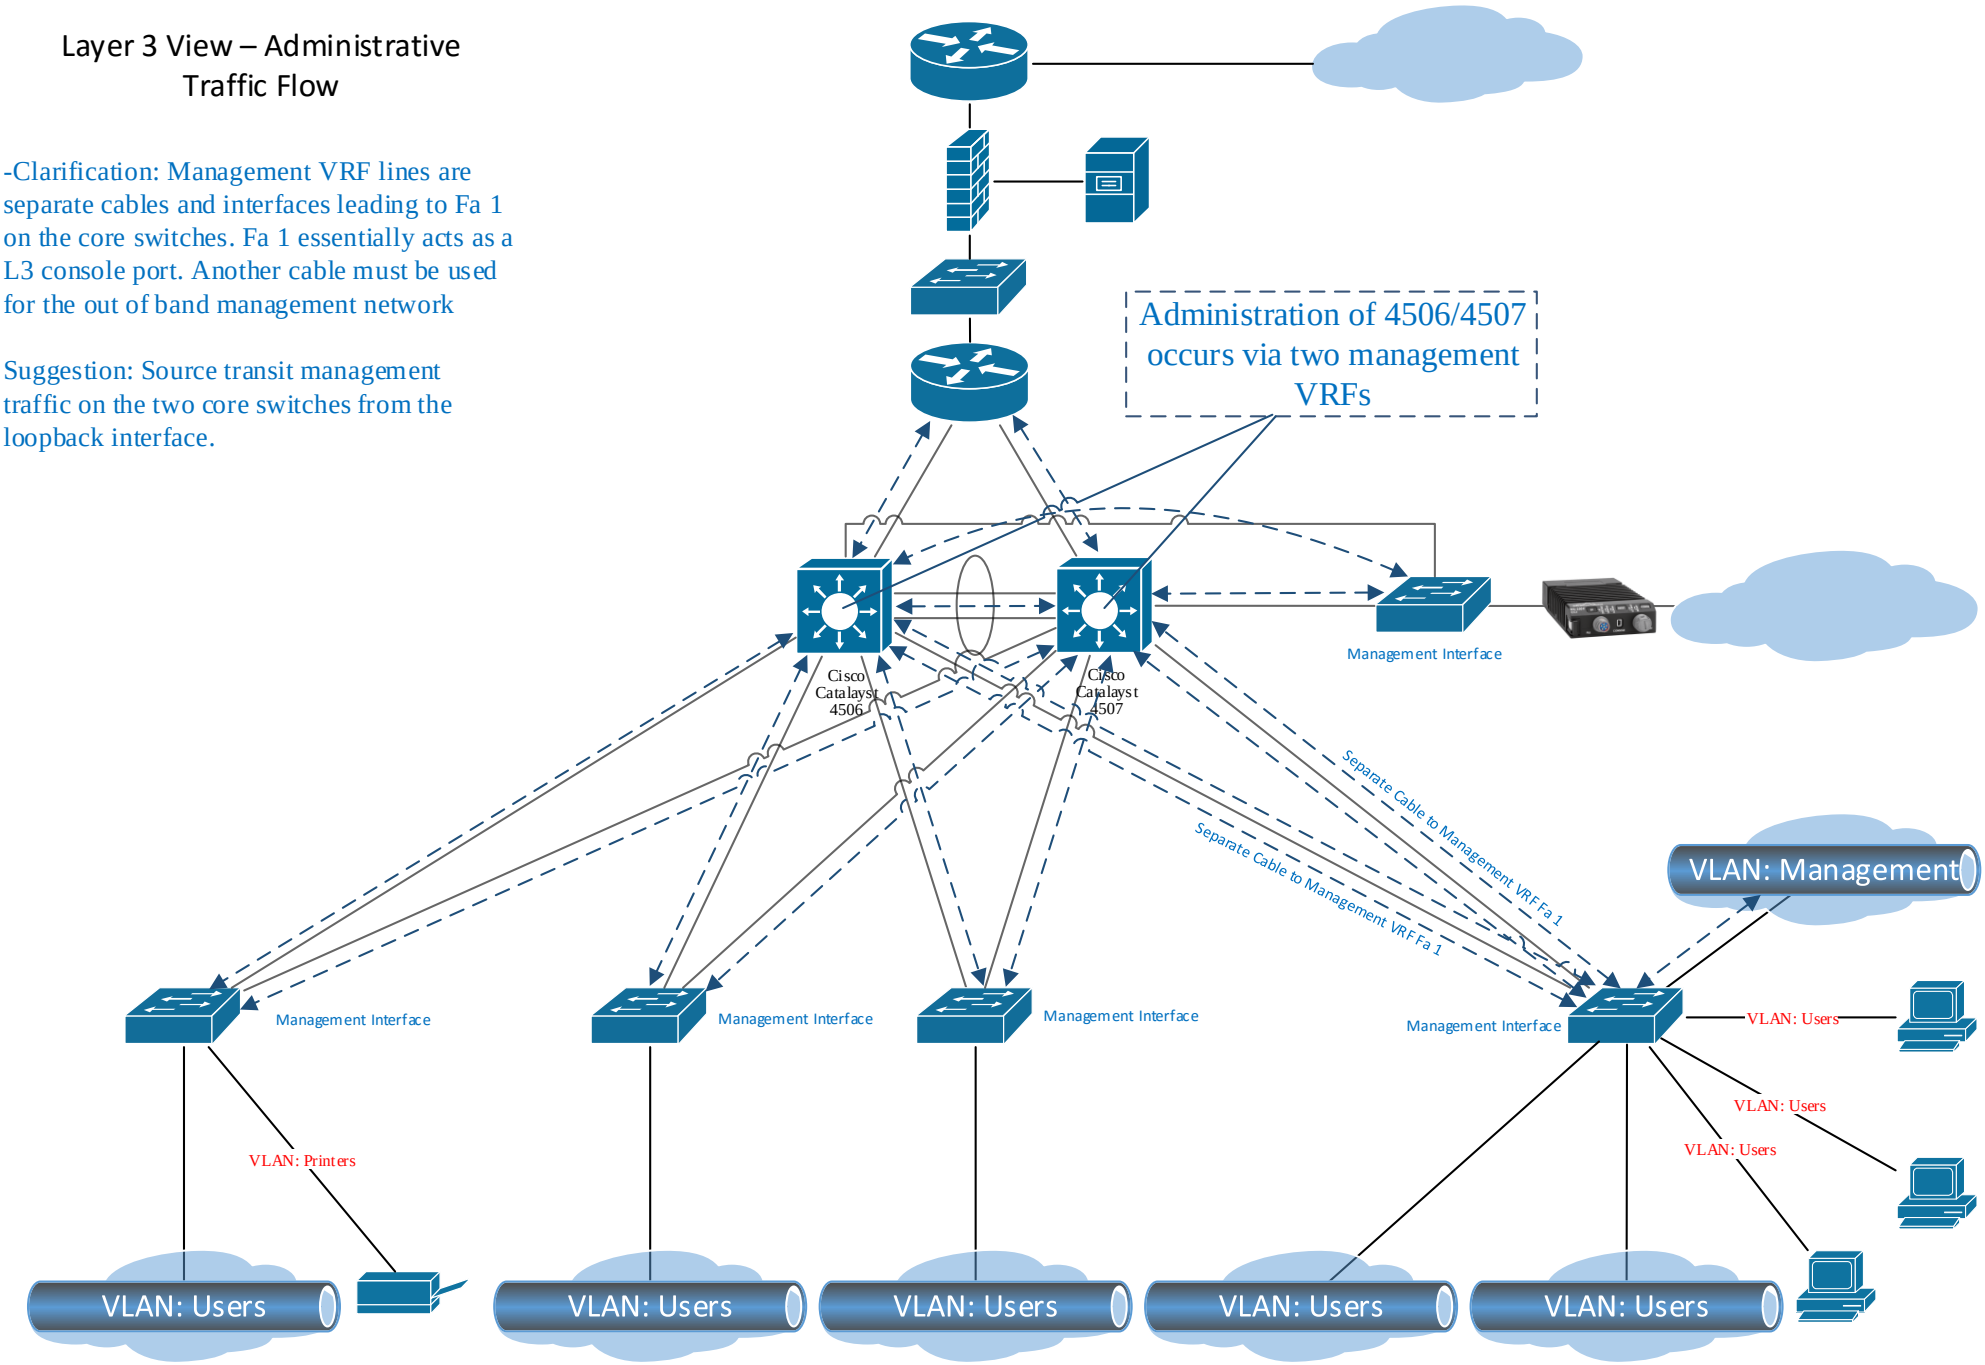

Layer 3 View – Access List Logic

-All switches should have ACLs in place to only allow bidirectional traffic to/from the server VLAN, the external network, and the printer VLAN to/from client machines. IE client machines should not be able to talk directly to each other.
-Network management nodes in the management VLAN may ONLY communicate with management VLAN. ACLs must enforce this policy
-As new VLANs are put in place to separate servers, ACLs must be installed to logically separate those VLANs
-All Cisco network devices MUST use control plane protection
-Cisco CEF should be enabled wherever possible (NET0949)

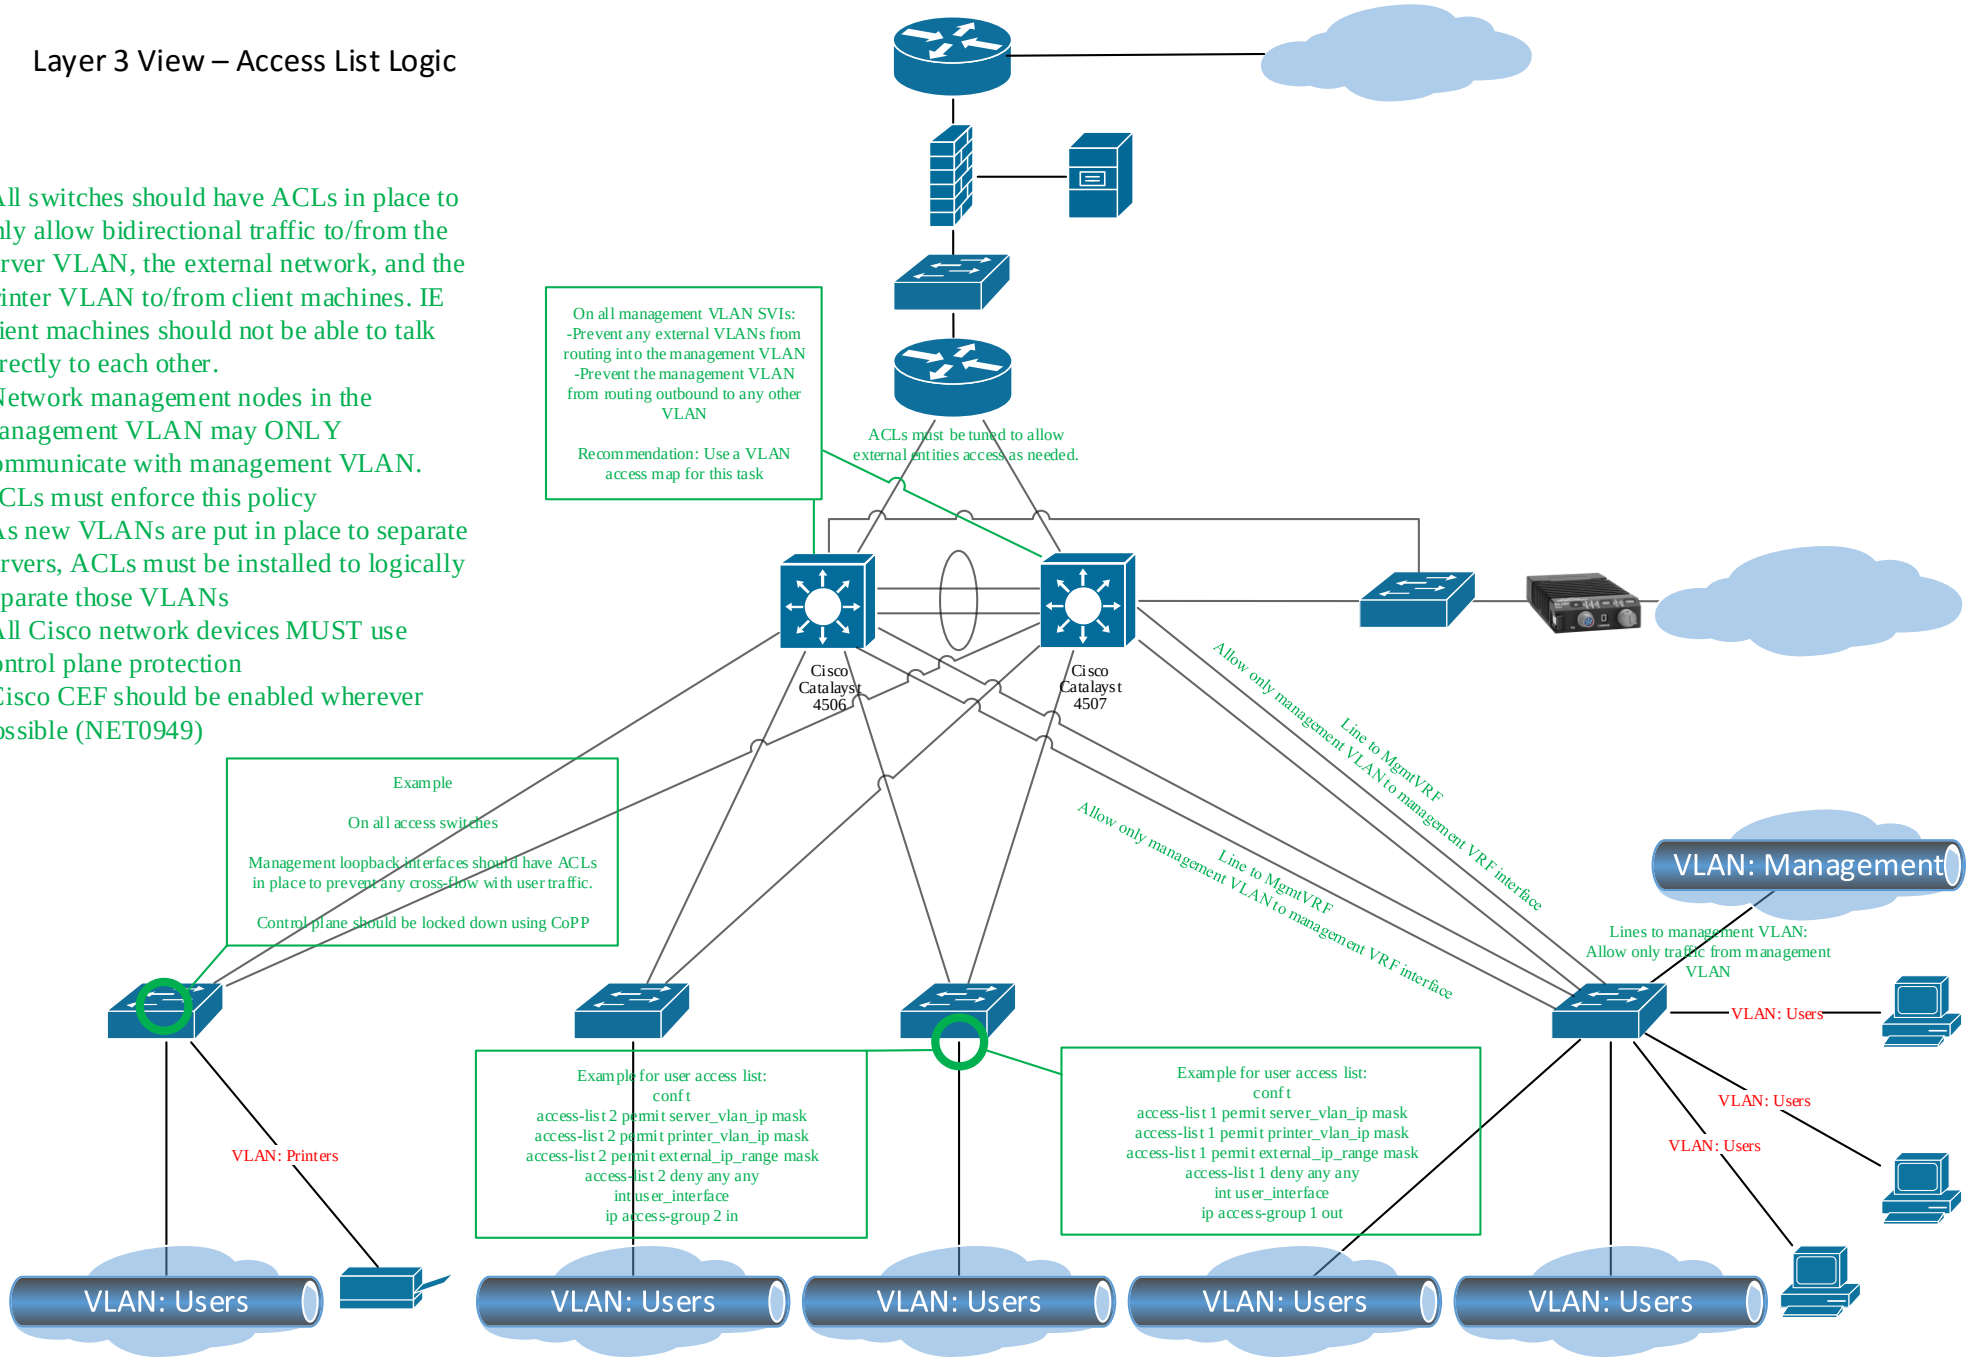

SECTION 3 – LAYER 3 DESIGN

The following sections discuss layer 3 design decisions.

3.1. HSRP

HSRP provides default gateway redundancy within the network. HSRP operates on a per VLAN basis. If either core switch goes down the other assumes the default gateway role for the VLANs of the down switch. Cisco defaults the hello and hold timers to very conservative values. Network administrators may safely tune HSRP timers to much more aggressive values to achieve faster convergence times. 1 second for the hello timer and 4 seconds for the hold time are the current industry recommendations. HSRP preemption is required otherwise a switch that goes down does not become the primary for its VLANs after coming back up until the other switch goes down.

3.2. Create In-Band Management Network

STIG NET0992 requires the use of an OOBM. An out of band management network includes separate infrastructure devices which only connect to management interfaces and management devices on the network. However, STIGs permit the use of designated management interfaces protected by ACLs, Cisco Control Plane Protection (CoPP), and Cisco Management Plane Protection to create an in-band management network. See NET0992 Alternate Solution text in the Infrastructure L3 Switch Secure Technical Implementation Guide STIG for details.

The design includes a separate in-band management network. A management network separate from the production network is critical to network security. Properly implemented, this heavily blunts the use of privilege escalation attacks. A compromise in any section of the network outside of the management network does not directly yield an attack avenue to the management network. For example, a server compromised via a userland SQL injection attack does not give the adversary any direct connection or access to the management functionality of any device on the network because that interface would be a member of the user VLAN.

Required by STIGs: NET0986, NET0987, NET0988, NET0989, NET0990, NET0991, NET0992, NET0994, NET0995, NET0996, NET0997

3.3 Implement Control Plane Protection

Infrastructure devices have three notional “planes’ consisting of the data plane, control plane, and management plane. The control plane is responsible for signaling on the network and is consulted whenever a decision must be made in the routing or switching process. Unnecessary traffic or a denial of service attack against the control plane can cause a device to go down. Control plane protection is a set of features Cisco provides to protect the control plane from flooding.

Required by STIGs: NET0966

3.4. Implement Access Control Lists

Access lists enforce policy between the different VLANs or interfaces. They are in place to provide logical separation between different security groups and zones. VLANs provide the ability to identify different roles on the network, but without access lists they do nothing other than break up the broadcast domain. ACLs are the primary way to create a logically separate, in-band management network.

Required by STIGs: NET0992, NET0994, NET0989, NET1000, NET1004, NET-SRVFRM-003, NET-SRVFRM-004

ANNEX A: ACRONYMS

ACL	Access Control List
BPDU	Bridge Protocol Data Unit
CoPP	Control Plane Protection
CS	Cyber Security
CSU	Channel Service Unit
DSU	Data Service Unit
HSRP	Hot Standby Routing Protocol
IDPS	Intrusion Detection/Prevention System
IOP	Information Operations
ITIL	Information Technology Infrastructure Library
LACP	Link Aggregation Control Protocol
OOBM	Out of Band Management Network
OSI	Open System Interconnect
PAgP	Packet Aggregation Protocol
PVST	Per-VLAN Spanning Tree
RPVST+	Rapid Per-VLAN Spanning-Tree
RADIUS	Remote Authentication Dial In User Service
RST	Rapid Spanning-Tree
SNMP	Simple Network Management Protocol
STIG	Security Technical Implementation Guide
STP	Spanning-Tree Protocol
SVI	Switch Virtual Interface
TACACS	Terminal Access Controller Access-Control System

VoIP	Voice Over Internet Protocol
VLAN	Virtual Local Area Network
VRF	Virtual Routing and Forwarding

ANNEX B: STIGS REFERENCED

Note: Individual items appear in multiple STIGs. Items are not listed multiple times even though they may appear in multiple STIG guides.

Infrastructure L3 Switch Secure Technical Implementation Guide - Cisco

NET0897	The router must use its loopback or OOB management interface address as the source address when originating TACACS+ or RADIUS traffic.
NET0949	Cisco Express Forwarding (CEF) must be enabled on all supported Cisco Layer 3 IP devices.
NET0966	The router must have control plane protection enabled.
NET0987	Traffic from the managed network is able to access the OOBM gateway router
NET0988	Traffic from the managed network will leak into the management network via the gateway router interface connected to the OOBM backbone.
NET0989	Management network traffic is leaking into the managed network.
NET0992	The management interface is not configured with both an ingress and egress ACL.
NET1004	The IAO will ensure that only authorized management traffic is forwarded by the multi-layer switch from the production or managed VLANs to the management VLAN.
NET-SRVFRM-003	Server VLAN interfaces must be protected by restrictive ACLs using a deny-by-default security posture.
NET-SRVFRM-004	The IAO will ensure the Server Farm infrastructure is secured by ACLs on VLAN interfaces that restrict data originating from one server farm segment destined to another server farm segment.

Layer 2 Switch Security Technical Implementation Guide - Cisco

NET0990	The OOBM access switch is not physically connected to the managed network element OOBM interface.
NET0991	The network element’s OOBM interface must be configured with an OOBM network address.

NET0994	The management interface is an access switchport and has not been assigned to a separate management VLAN.
NET0995	An address has not been configured for the management VLAN from space belonging to the OOBM network assigned to that site.
NET0996	The access switchport connecting to the OOBM access switch is not the only port with membership to the management VLAN.
NET0997	The management VLAN is not pruned from any VLAN trunk links belonging to the managed network's infrastructure.
NET1003	The management VLAN is not configured with an IP address from the management network address block.
NET-VLAN-004	The IAO/NSO will ensure VLAN1 is not used for user VLANs.
NET-VLAN-005	The IAO/NSO will ensure VLAN1 is pruned from all trunk and access ports that do not require it.
NET-VLAN-006	The IAO/NSO will ensure VLAN1 is not used for in-band management traffic. A dedicated management VLAN or VLANs will be defined to keep management traffic separate from user data and control plane traffic.
NET-VLAN-008	The IAS/NSO will ensure that the native VLAN is assigned to a VLAN ID other than the default VLAN for all 802.1q trunk links.
NET-VLAN-009	The IAO/NSO will ensure access switchports are not assigned to the native VLAN.

Network Policy Security Technical Implementation Guide

NET0998	A separate management subnet has not been implemented.
NET1002	The management station or server is not connected to the management VLAN.
NET-VLAN-010	The IAO will ensure the Server Farm is segmented by isolating business functions such as databases, applications, web, and email using VLAN provisioning.

ANNEX C: REFERENCES

"Campus Network for High Availability Design Guide." Cisco. N.p., 03 Dec. 2008. Web. 02 Dec. 2014. <http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html#wp1107687>.

"Cisco Guide to Harden Cisco IOS Devices." Cisco. N.p., n.d. Web. 02 Dec. 2014. <http://www.cisco.com/c/en/us/support/docs/ip/access-lists/13608-21.html#anc34>.

"Control Plane Policing Implementation Best Practices." Cisco. N.p., n.d. Web. 02 Dec. 2014. <http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html>.

Hucaby, Dave. CCNP SWITCH 642-813 Official Certification Guide. Indianapolis, IN: CISCO, 2010. Print.

"Management Plane Protection." Cisco. N.p., 20 Aug. 2008. Web. 02 Dec. 2014. <http://www.cisco.com/c/en/us/td/docs/ios/12_4t/12_4t11/htsecmpp.html>.

Tiso, John. Foundation Learning Guide Designing Cisco Network Service Architectures (ARCH). Indianapolis, IN: Cisco, 2012. Print.