

TOMORROW starts here.



Cisco *live!*

CCDE: The Cisco Certified Design Expert

BRKCRT- 8001

BRKCRT- 8002

Bruce Pinsky – Distinguished Engineer

Alvaro Retana – Distinguished Engineer

Yusuf Bhajji – Learning and Development Manager

The CCDE

- What's the CCDE?
- CCDE Exam Topics
- CCDE Domains and Drivers
- Technologies Overview
- CCDE Written Examples
- CCDE Practical Examples
- Conclusion / Review





What is the CCDE?

What is the CCDE?



Industry



Candidate



Business

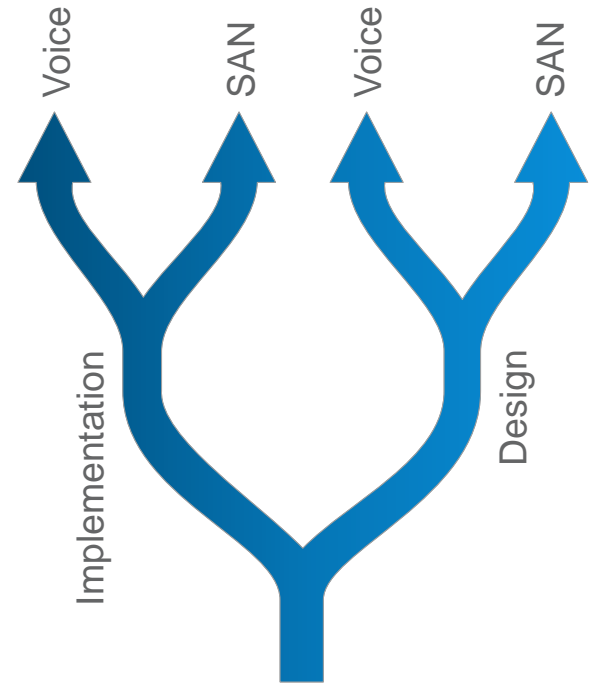


Certifications

Perspectives on Design

Industry

- Network engineering is specialising
- Over time, people move from operations to design work
- We seem to have lost our “roots”
 - Focus on technologies (Voice, WAN Acceleration, Security, etc.)
 - Focus on places in the network (Data Centre, WAN, Campus, etc.)



Network Infrastructure
Cisco *live!*

Perspectives on Design

Industry

- Basic design is no longer widely taught or practiced
 - It just works
 - It's not exciting, it's just plumbing
 - It's easy
 - It's not new
- And yet...
 - Basic design problems still crop up
 - “Exciting stuff” doesn't work without a solid routed design
 - We spend far too much time and money managing failures in basic design work

Perspectives on Design

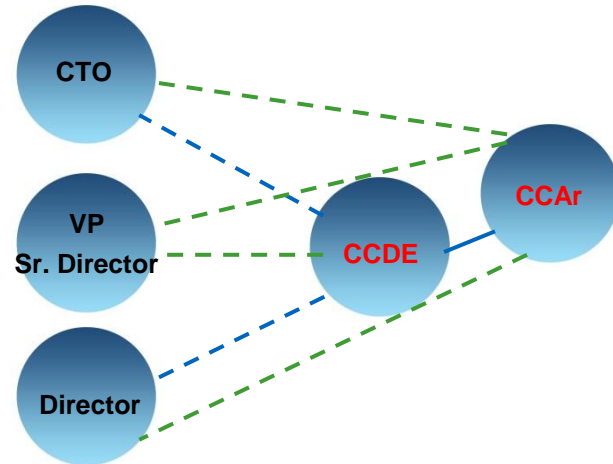
Candidate

- Create a certification relevant to designers
- Provide a baseline set of skills
 - Builds on top of these skills, not in lieu of them
 - Encourages picking up “skipped skills”
- Encourage end-to-end (big picture) design thinking
 - Places in the Network are great, within the framework of a running network
 - Technologies are great, but need to be used for larger design goals

Perspectives on Design

Business

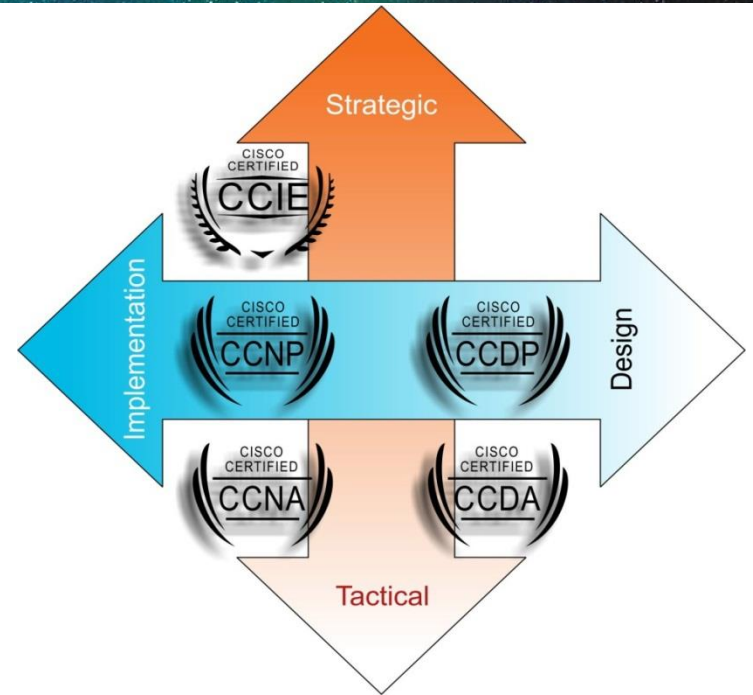
- Interact horizontally with middle level business management
- Understand impact of business decisions and direction on design



Perspectives on Design

Certification

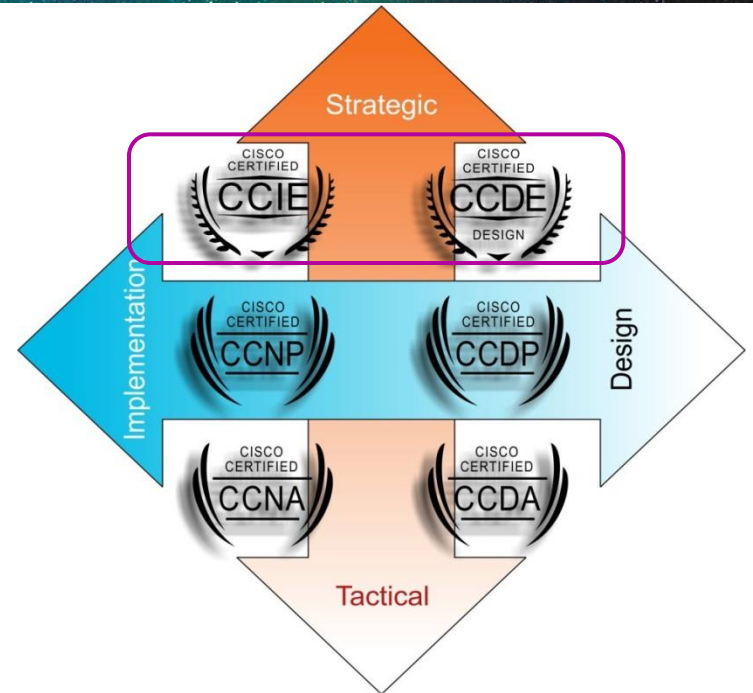
- Certifications can be seen in two dimensions
 - What does it certify?
 - Implementation or Design
 - How does it relate to the business?
 - Tactical or Strategic



Perspectives on Design

Technical

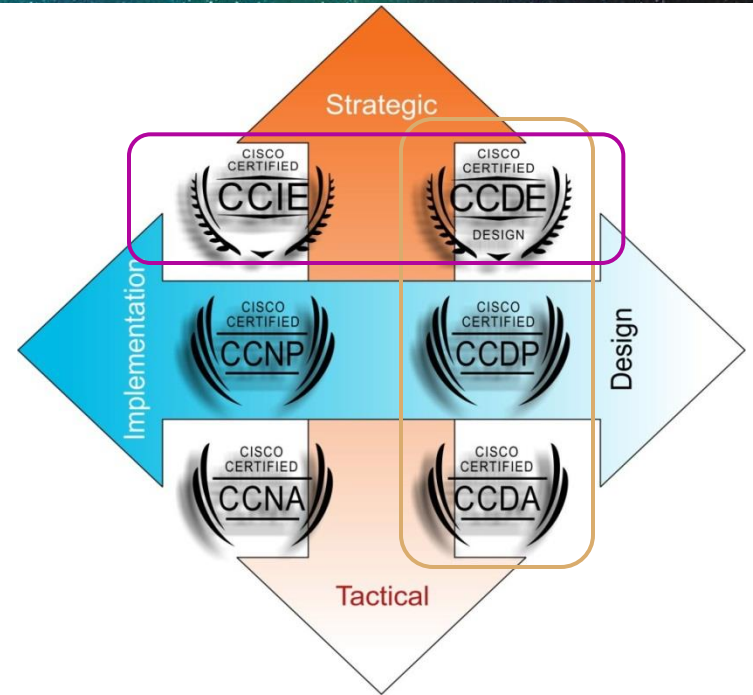
- Where is this network now?
- What changes do I need to make to....
 - Merge these networks?
 - Implement this application?
 - Provide security?
 - Prepare for the next five years?



Perspectives on Design

Business

- How do I transition the network?
 - Business hurdles?
 - Technical hurdles?
 - People hurdles?



Cisco Certified Design Expert

What the CCDE “IS”

- Validates skills in job roles, including:
 - Senior Network Designer
 - Network Lead for Enterprise IT Infrastructure Team
 - Network Lead for Enterprise Architecture Team
- Emphasises network design principles at the routing layer as well as the ability to assess and translate network business requirements into technical designs



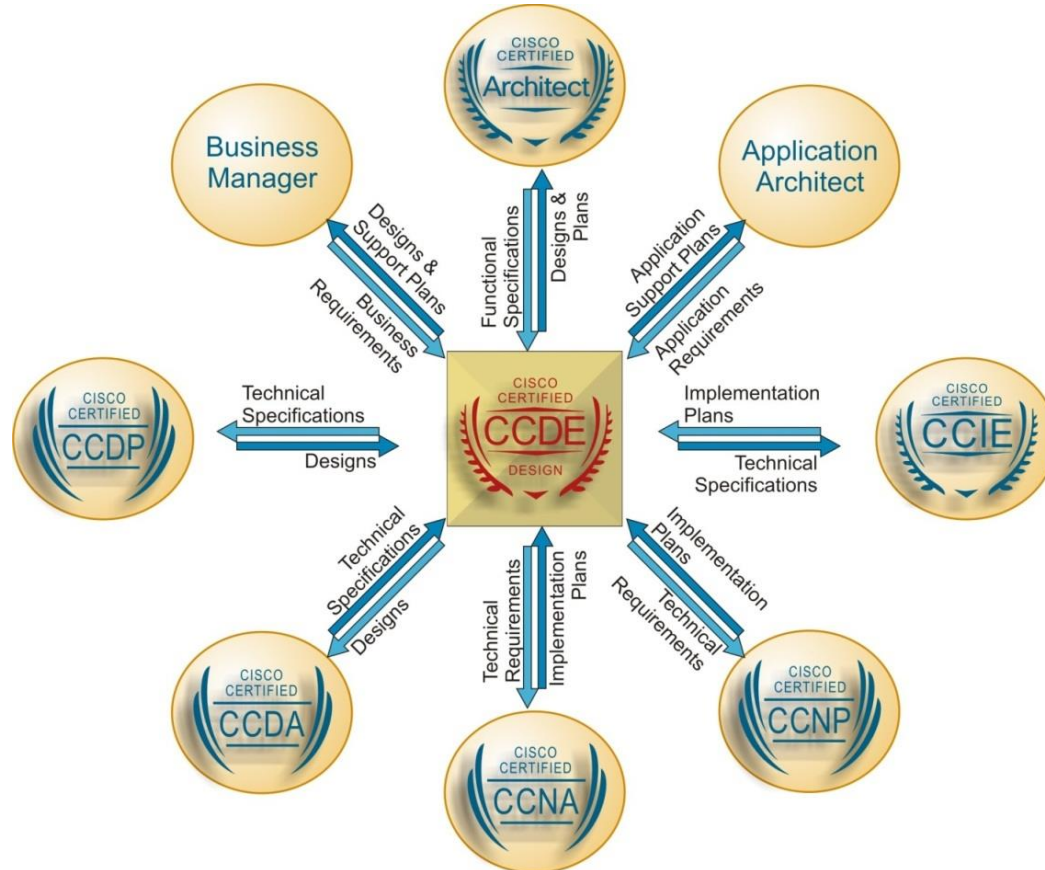
Cisco Certified Design Expert

What the CCDE “IS NOT”

- This is not a business test
 - There is no “budget” for any given problem
- But—there are business problems on the test
 - Business problems provide the primary structure
 - Business problems provide the primary driver towards specific technology solutions



Relationship to Other Certifications



Design Certifications

Comparison Within the Track



Senior Network Designer

- Large scale network design
- Variety of business environments
- Troubleshoots and resolves design level issues



Network Designer

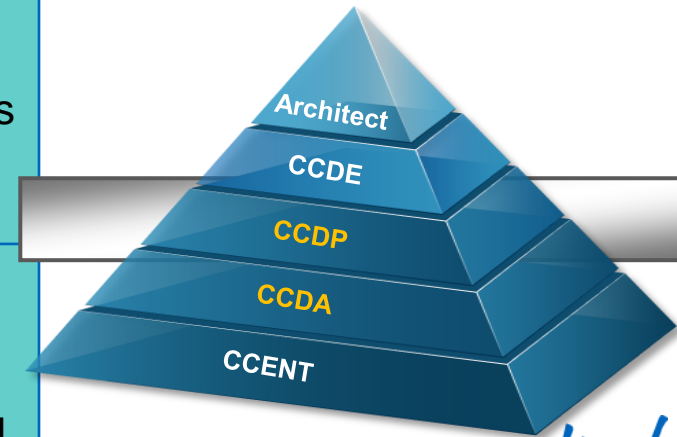
- Moderate scale networks
- Narrow set of business environments
- Designs components of larger networks



Junior Network Designer

- Understands network design fundamentals
- Designs components of medium and large scale networks

BRKCRT-8002

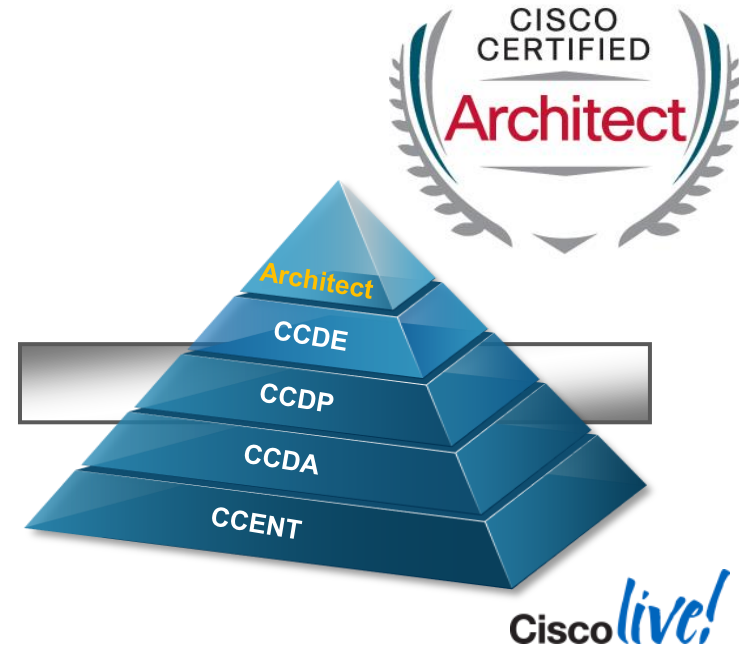


Cisco Public

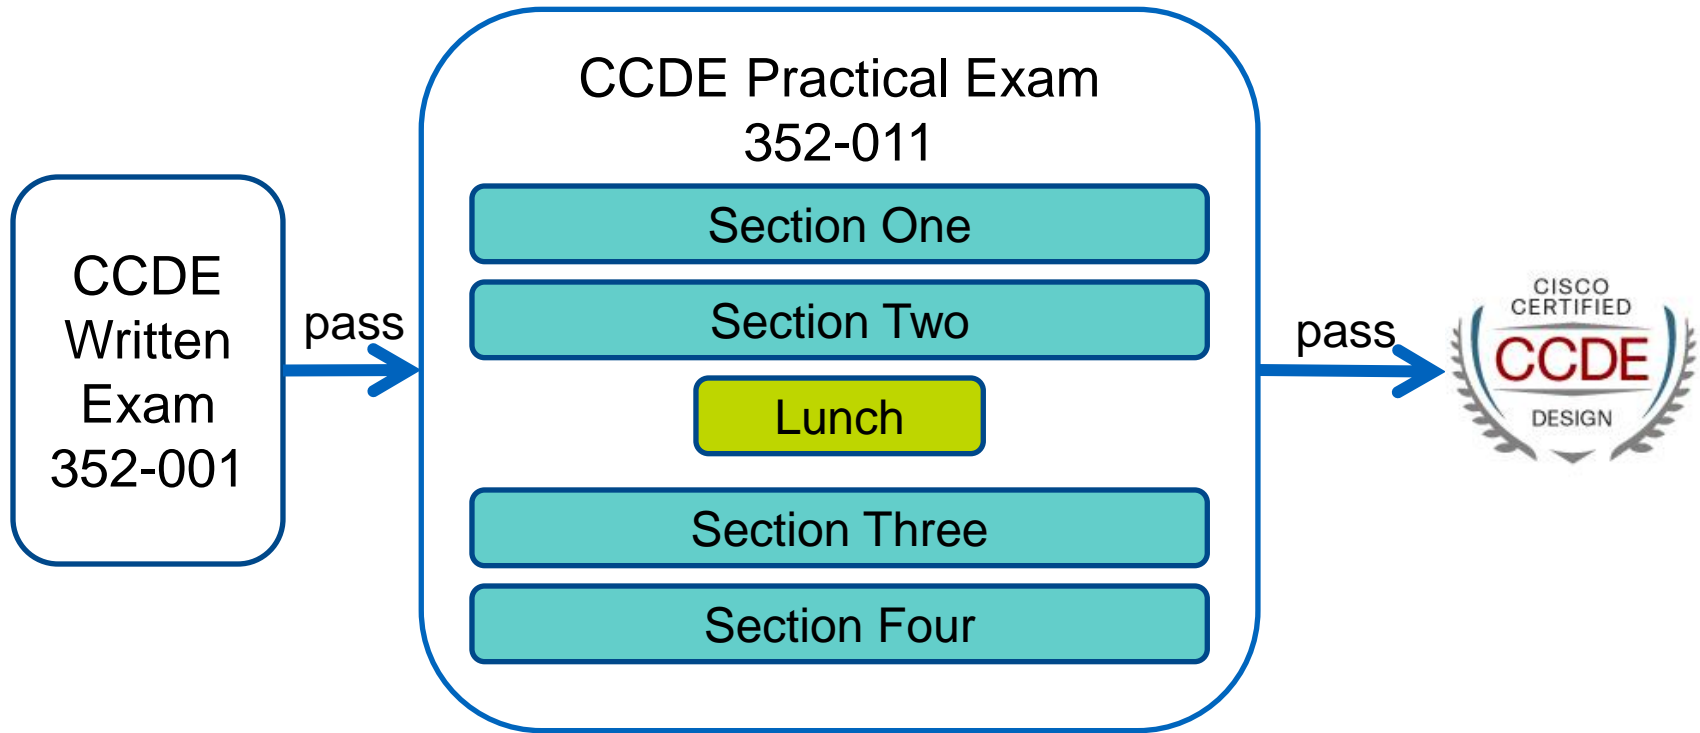
Cisco *live!*


Cisco Certified Architect

- Highest-level Cisco certification
- Valid CCDE → CCAr Application and Interview → Board Exam → CCAr
- Focused on
 - Infrastructure (Technology)
 - Business requirements and Long Term Architecture
 - Conflict resolution and communication skills
- Not focused on
 - Technology implementation
 - Detailed deployment
 - Processes and standards
 - Application, enterprise, IT or solutions architect



Certification Process





CCDE Exam Topics

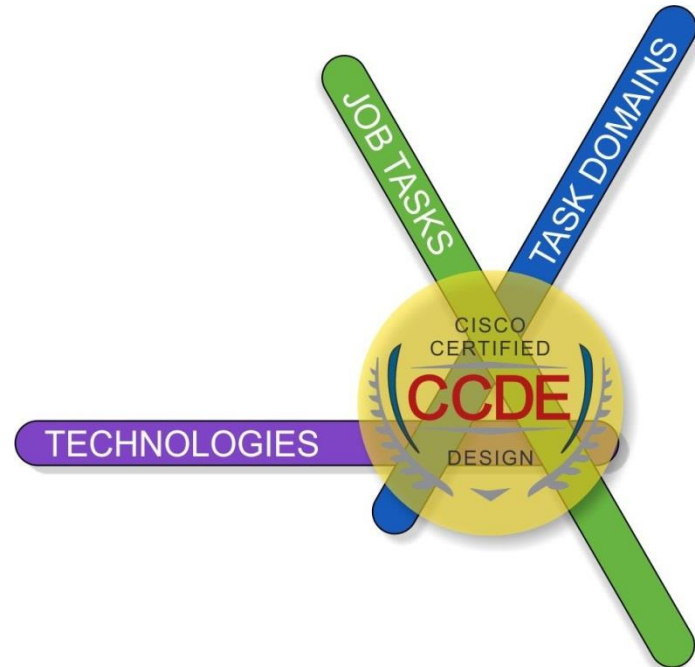
CCDE Exam Topics v2.0

- CCDE Exam Topics v2.0:
 - Represent important high-level design-specific topics
 - Assessed by CCDE Written and CCDE Practical exams
- Exam questions are aligned to Exam Topics in order to assess candidates' expert-level design-related skills, knowledge, and abilities

Exam	URL
CCDE Written	https://learningnetwork.cisco.com/docs/DOC-13054
CCDE Practical	https://learningnetwork.cisco.com/docs/DOC-13059
CCDE Practical expanded	https://learningnetwork.cisco.com/docs/DOC-11798

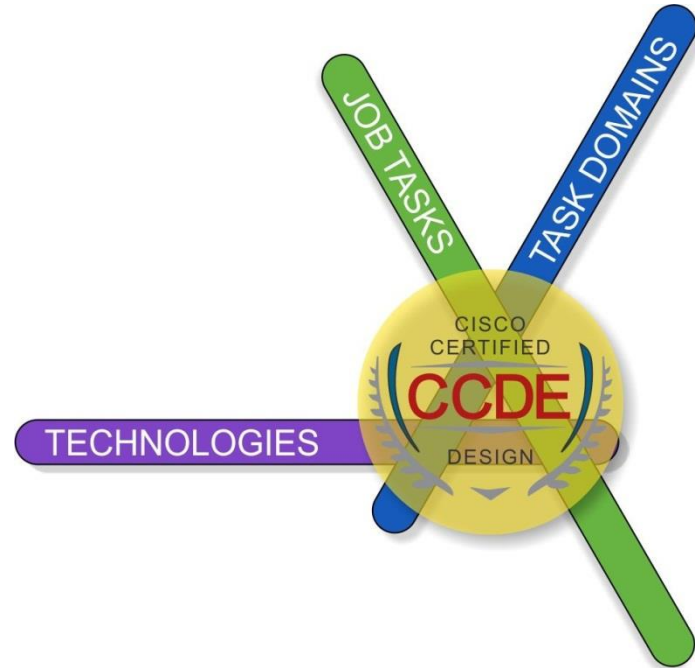
The CCDE Exam Topics

- Task domains are large, overarching, and related to design
- Job tasks are smaller, individual, and related to day-to-day functions
- Technologies describe areas where a designer must be proficient



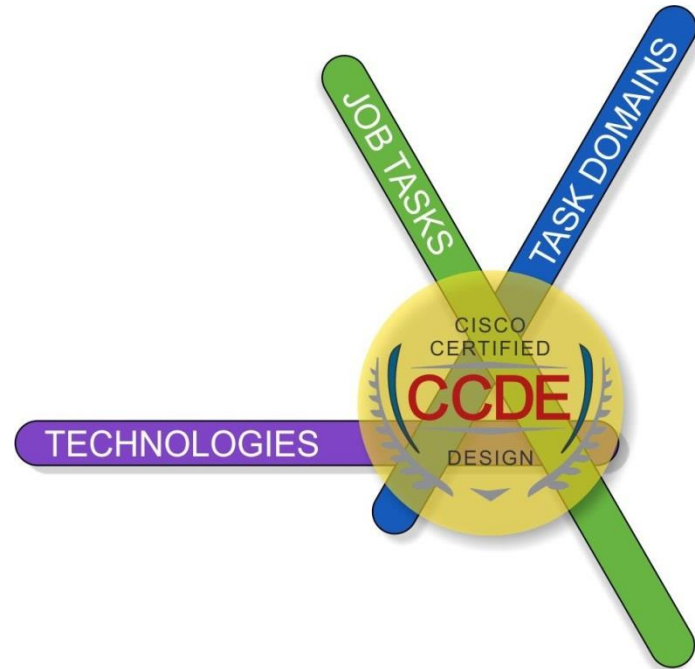
The CCDE Exam Topics

- The written focuses only on the technology exam topics
- The practical is multidimensional
 - Each question or problem must map to all three areas of expertise



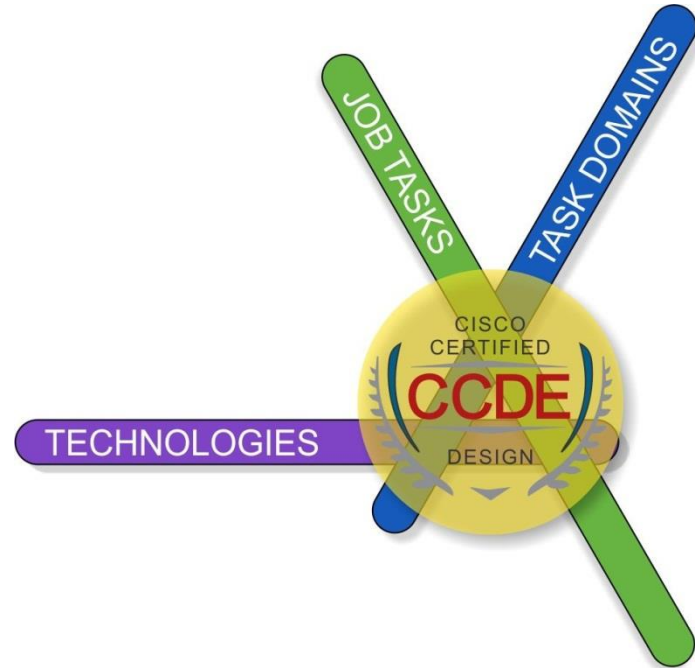
Task Domains

- Merge/Divest
- Add Technology
- Replace Technology
- Scaling



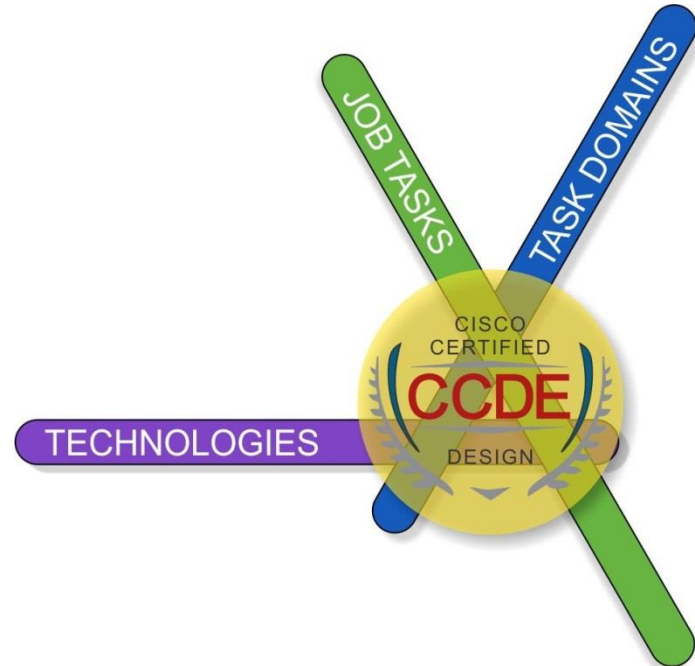
Job Tasks

- Analyse
- Design
- Implement & Deploy
- Validate & Optimise



Technologies

- Technologies are covered in the *Technologies* sections




Essential Skills

- Reading
 - Read to find information that's important
 - Skimming when possible, rather than reading completely
 - Time is an issue; taking in a lot of information at a solid pace is important
- Abstracting
 - See sections of the network as “black boxes”
 - But, be able to see into the black box as needed
 - Don't get overwhelmed with information!
 - Any design situation involves a lot of information
 - It takes a lot of skill to categorise it, and know where it is when you need it

Essential Skills

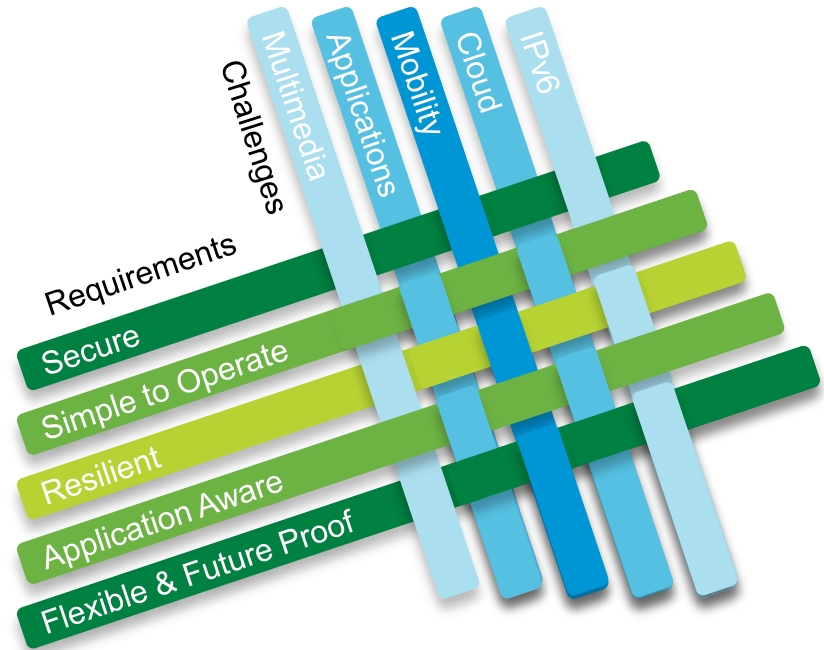
- Analysing Information
 - If you see 2 and 2, can you add them to make 4?
 - If you see 1 and 1, can you add them to make 10?
 - Binary!
 - Do you know when things relate, and when they don't?
 - Do you know what to apply where?
- Knowing why you would do this
 - Many questions are in two stages
 - What would you do
 - Why would you do it
 - This goes farther than “rules of thumb”



Domains & Drivers


Domains & Drivers

- The next generation challenge won't be won...
 - ...in the data centre
 - ...in the cloud
 - ...or even in the core, or campus, or...



Domains & Drivers

- The next generation challenge is to put the right information in the right places to meet business requirements
 - Right sizing compute services
 - Asking business questions, replying with design and technology answers
- Network technologies are a pendulum...
 - Layer 2/layer 3, centralised/decentralised, virtualise/back to basics
 - To steady the ship, focus on the big picture!



Domains & Drivers

- The CCDE breaks the problem into two pieces
 - Domains
 - Large common problem areas in designing a network
 - Drivers
 - What problem are you trying to solve?
 - What tools do you have to solve the problem?

Domains

Availability

- “Five 9’s” is only the beginning...
- What is the mean time between failures?
- Can you choose the simplest solution for the problem?
- What do the applications require to run?
 - Delay?
 - Jitter?
 - bandwidth profile?
 - loss of connection tolerance?

Domains

Serviceability

- What is the “mean time between mistakes?”
- What is the mean time to repair?
- Can you balance complexity against services?
- Could you explain this solution at 2am to a sleepy TAC engineer through the phone?

Domains

Scalability

- Scale is more than size
- How many services, where are they located, how do they work?
- Is the network fast enough?
- Will that solution work if I double the size of the deployment?

Domains

Security

- Where is this network vulnerable to attack?
- What mechanisms can you use to prepare for and repel those attacks?
- How can the network devices be protected against attacks?
- How can the network be protected against attacks?
- How can the network be used to protect connected devices against attacks?

Domains

Flexibility

- Where will this business be in five years?
- What if you need to double the size of this network tomorrow?
- Can you tell me when the network will need to grow, and where?

Design Drivers

- It's not...
 - ... just hierarchical design...
 - ... just a bunch of design problems tossed into a big hat...
- There is an underlying design philosophy
 - What are the network design drivers?
 - What are the tools designers have to meet the challenges posed by these drivers?
- What technology fits where?
 - Why would you use this here?
 - What combination of business and technical problems does this solve?
- How would you deploy this?

Business Drivers

Key Questions

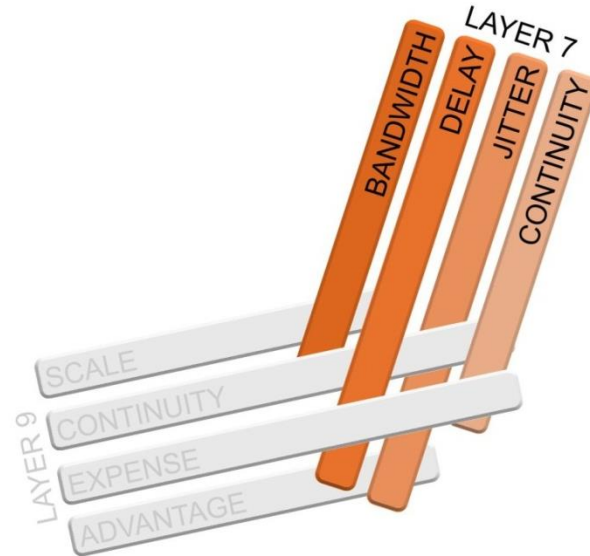
- How rapidly does the business change?
- What is the business sizing model (growth/downsize/steady)?
- What are the key business processes?
- What is the business' attitudes towards the network?



Application Drivers

Key Questions

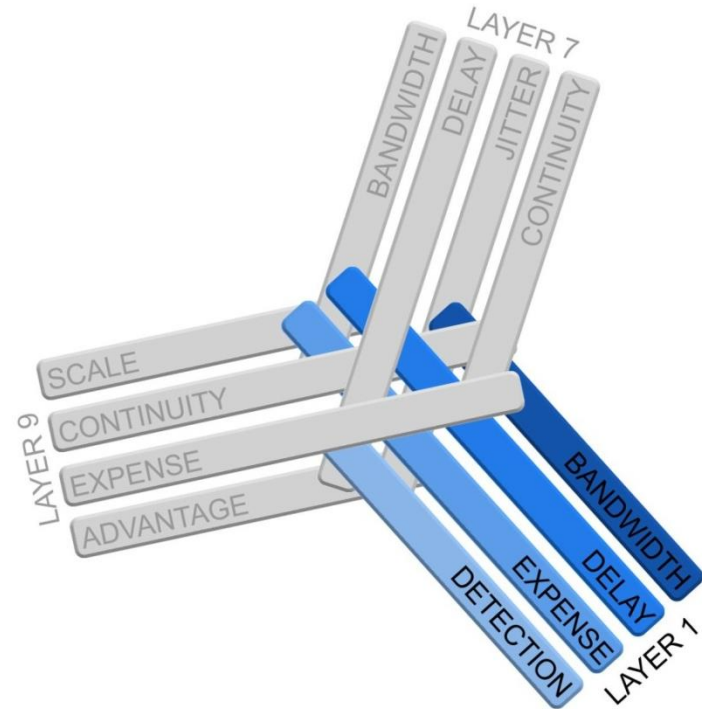
- What are the key applications on the network?
- What are the characteristics of this set of applications?
 - Bandwidth requirements?
 - Tolerance for delay?
 - Tolerance for jitter?
 - Tolerance of network reconvergence?



Link Drivers

Key Questions

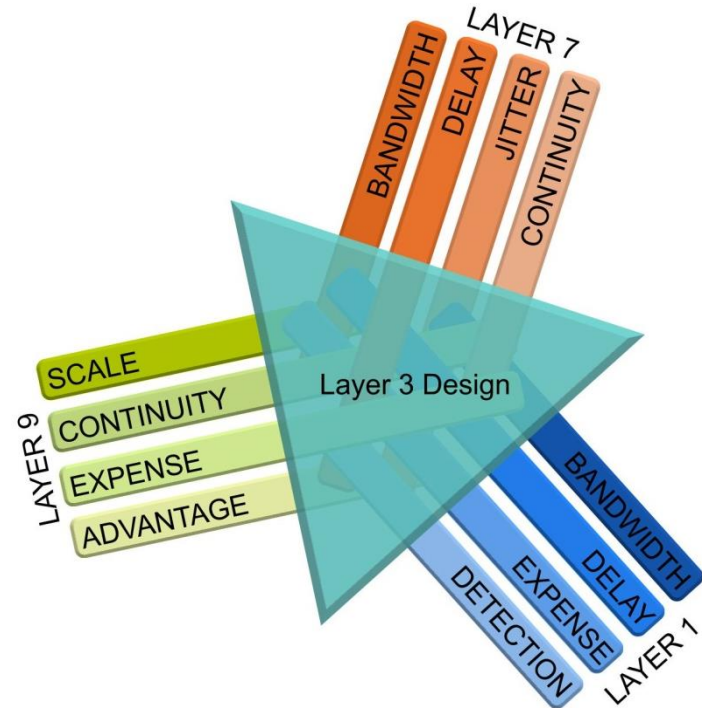
- What's the bandwidth?
- What's the delay profile?
- What's the jitter profile?
- How fast does layer 3 learn about a failure?



Where it all Meets

Layer 3

- All of these drivers meet at layer 3, the control plane
 - This is why the CCDE is a control plane driven test
 - If you can't do layer 3, you can't build a network that will support business requirements and applications on top of available links and hardware
- The control plane provides the tools to bring layer 9, layer 7, and layer 1 together





Technologies Overview

A Word on Technologies...

- Essentially...
 - Any published RFC implemented by more than one vendor
 - Any common industry technology
 - Any common industry practice
 - EIGRP
- This test is broad from a technical viewpoint, but...

A Word on Technologies...

What the CCDE is NOT

- You do not “go forth and configure”
 - This is higher level than the “?”
- This is not about choosing the right equipment in the right place
 - Hardware limitations only come in at a high level
 - Hardware changes occur on a daily basis
- The skills you demonstrate for this certification should be timeless

The One Crucial Question

Why do I care?

OR...

How does this impact design?





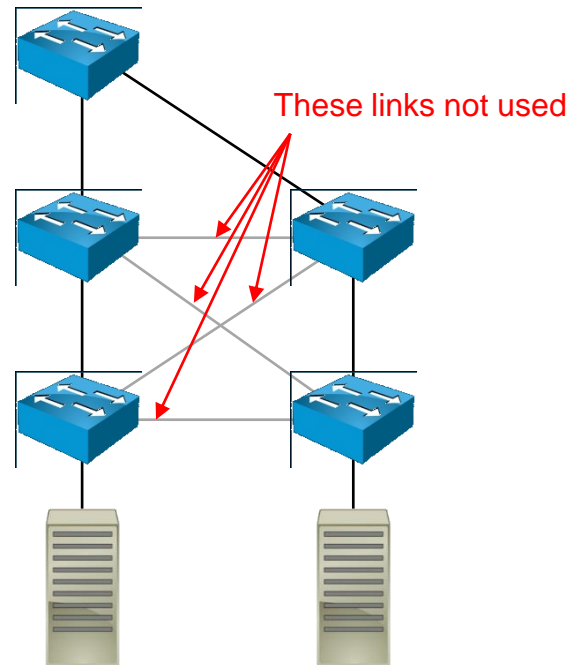
Layer 2 Control Plane

Speeds and Feeds

- What bandwidth do I need where in the network?
- This is only addressed at a very broad level on the CCDE
 - The cost of bandwidth shifts rapidly
 - “In detail” bandwidth requirements tend to be closely tied to application information you would normally need to look up
 - The “in detail” bandwidth requirements of specific applications can shift rapidly
- Speeds and feeds need to be taken into account
 - But it won't ever be a primary concern on the CCDE

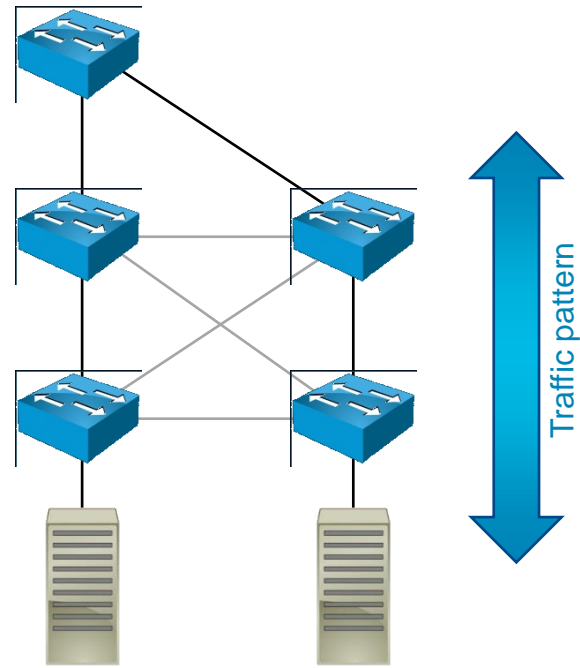
Spanning Tree

- Spanning Tree doesn't use all links
 - This can be a major problem for QoS, network usage, and efficiency
 - Current data centre design is focused on providing “clear bandwidth” server to server
- Spanning tree is a Bellman-Ford distance-vector routing protocol
 - Similar to RIP at the link level, rather than the destination level



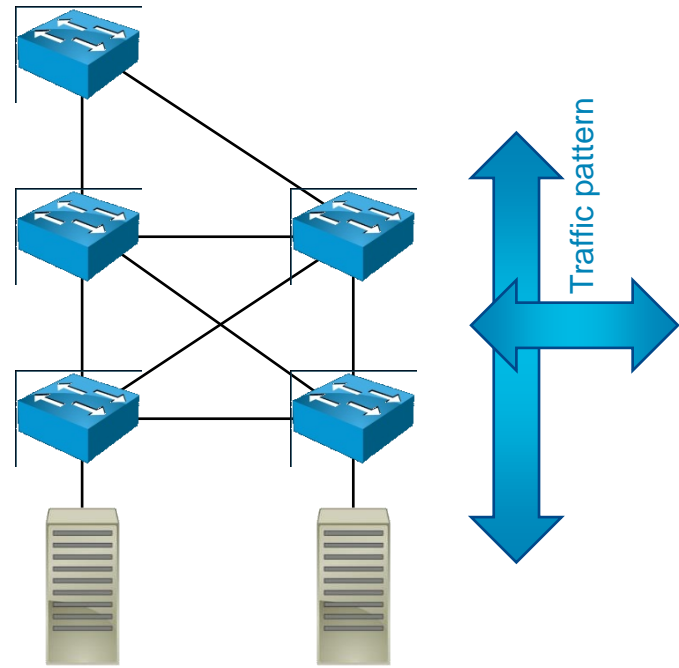
Spanning Tree

- Traffic patterns need to be planned “north south” in a spanning tree network
- The root of the spanning tree needs to be carefully chosen and controlled



Layer 2 Fabric

- Fabrics use link state control planes
 - On a per destination basis (rather than per link)
 - This allows a fabric to take advantage of all links
 - This also allows for better “east west” traffic flows
- Fabrics are generally more difficult to deploy
 - Proprietary (strong hardware restrictions)
 - More complex to manage and troubleshoot



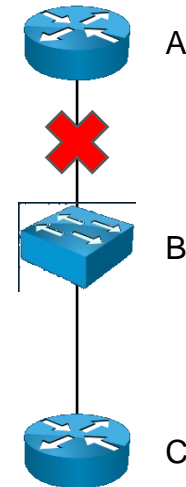
Making the Choice

Spanning Tree or Fabric?

- Assume
 - You are designing a campus network
 - Most traffic will be destined to off-campus destinations
- Spanning tree will be the simplest to deploy in this situation
 - Minimise unusable cross links
 - Anchor the root of the spanning tree

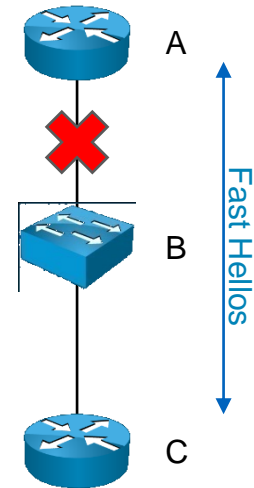
Down Detection

- In general...
 - Event driven notifications are always faster than polling based
 - But you need to take the actual links into consideration
- If the A->B link fails, when does C know about it?
- Only through some layer 3 detection mechanism
 - There is no notification from B that another link on the switch has failed



Down Detection

- Faster layer 3 detection mechanisms are available
 - Routing protocol fast hellos
 - BFD
- When should you use these?
 - If an event driven mechanism is available that will be as faster or faster than polling, use the event driven mechanism
 - Loss of carrier, etc.
 - If no event driven mechanism is available, move to some form of polling



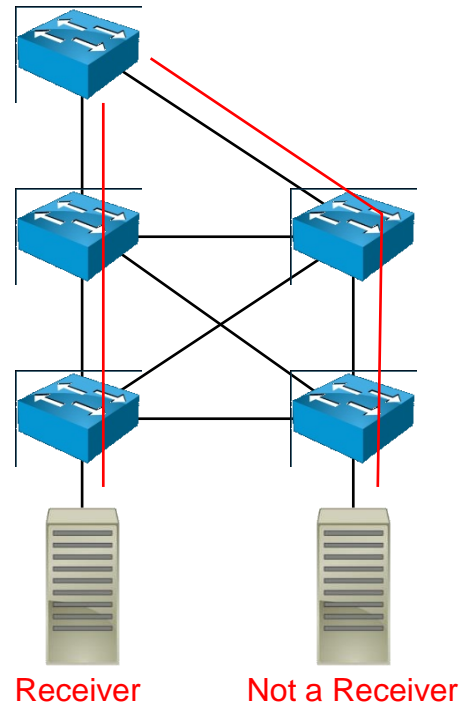
Failure Domains

- Broadcast storms
 - More hosts on a single broadcast domain
 - More broadcasts to find services and reachability
 - Less bandwidth to actually find services and reachability
 - At some point there are more broadcasts than the link can handle
- The only solution for broadcast storms is to break the broadcast domain into smaller pieces
 - Split the broadcast domain into multiple pieces with layer 3

Multicast

IGMP Snooping

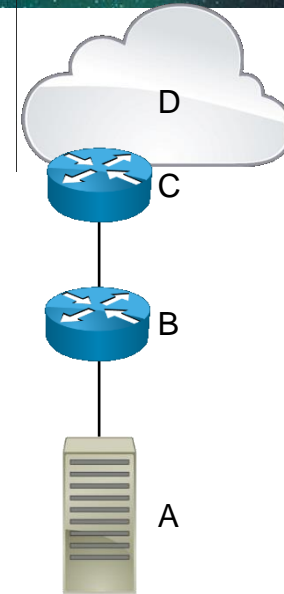
- Layer 2 broadcast domains send all multicast to all destinations
- IGMP snooping allows the switches to see multicast joins
 - Prune off sections of the network that don't need to receive the multicast
 - Saves network resources
- In short, always configure and use IGMP snooping where possible



Multicast

Multicast replication

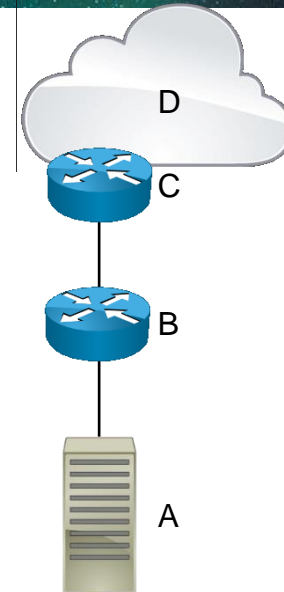
- If A sends a multicast packet
 - A separate packet must be created for each receiver
 - Where will this packet replication take place?
 - At the sender (A)?
 - At the first hop router (B)?
 - At the edge router into the cloud (C)?
 - Within the cloud (D)?



Multicast

Multicast replication

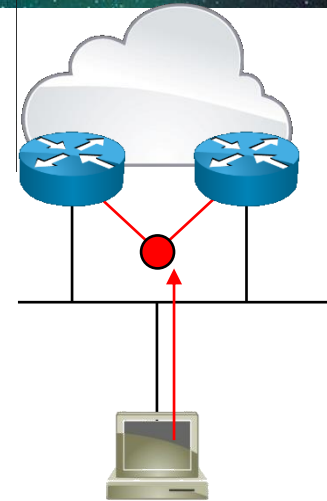
- The answer to this question has a major impact on the performance of multicast in the network
- Key Points
 - Push for layer 2 replication where possible
 - It's more efficient, particularly across broadcast mediums
 - Push replication as close to the receiver as possible



Layer 2 High Availability

First Hop Redundancy

- First hop redundancy protocols allow two routers to act as a single default gateway for a host or server
 - HSRP, VRRP, etc.
- They can also be used as a target for static routes
- The added complexity is minimal
 - But there is still complexity to handle
 - So use judiciously (where it makes sense)





Layer 3 Control Plane

Layer 3 Control Plane

Overview

- Protocol Theory
 - How neighbours are formed and maintained
 - How the best path is calculated
 - How aggregation is configured and deployed
 - How external routing information is handled
 - How protocols interact
- Protocols We Care About
 - OSPF
 - IS-IS
 - EIGRP
 - BGP
 - PIM (Sparse/Bidir)

Layer 3 Control Plane

- Modularity
- Hierarchy
- Resilience
- Example: Choosing the Right Core

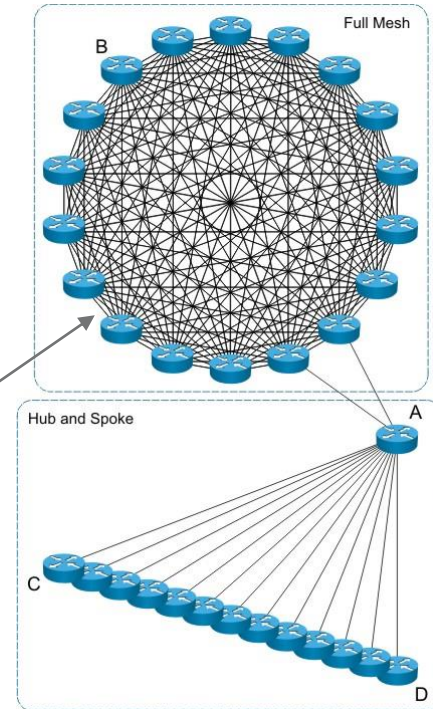
Modularity

- Provides
 - Failure domain separation
 - “Cookie Cutter” Configuration/Design
 - “Choke Points” for:
 - Security
 - Aggregation
 - Information Hiding
 - Etc.
- Primarily achieved through Hierarchy & Virtualisation

Modularity

- Where do you modularise a network?
- The general rule of thumb is divide complexity from complexity
 - Here we have a crazy full mesh and a hub and spoke
 - These two network topologies are complex
 - They require different methods of scaling and operation
 - Separate them!

1970's Spirograph network design



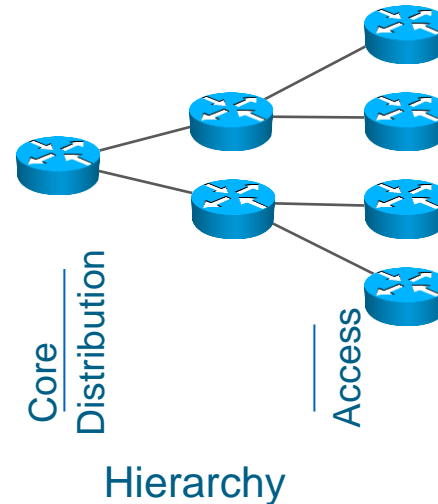
Modularity

Key Questions

- Where is the complexity in this network?
- Where has the network already been broken apart?
- How is the network broken apart?
 - Aggregation of reachability information?
 - Aggregation of topology information?
 - Reorigination of reachability information?
- Are these the best places to break this network apart?

Hierarchy

- Hierarchy is facilitated by summarisation and aggregation
 - Hiding topology information through an OSPF ABR, for instance
 - Hiding reachability information through route aggregation

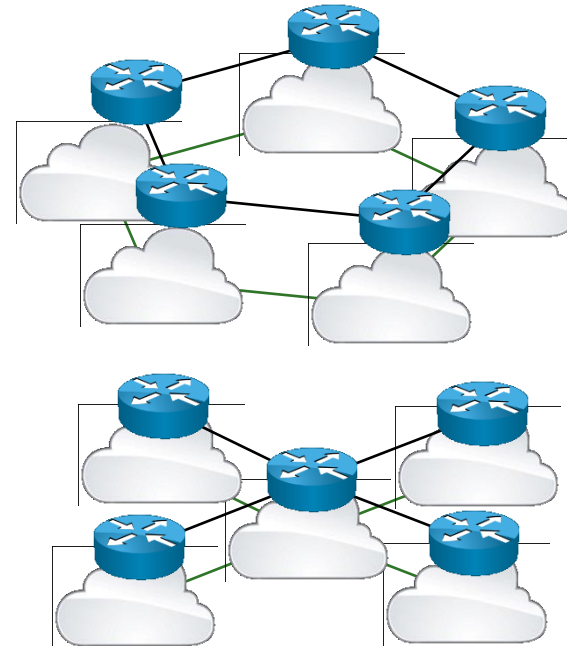


Hierarchy

- Hierarchical design used to be...
 - Three routed layers
 - Core, distribution, access
 - Only one hierarchical structure end-to-end
- Hierarchical design has become any design that...
 - Splits the network up into “places,” or “nodes”
 - Separates these “nodes” by hiding information
 - Organises these “nodes” around a network core
 - IE, roughly “hub and spoke” at a macro level

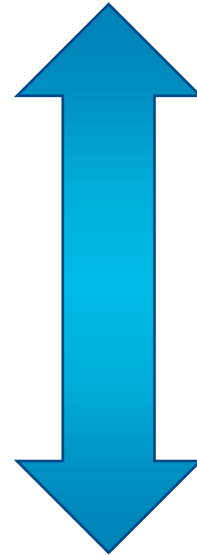
Hierarchy

- Why should hierarchical design follow a “roughly hub and spoke” design?
- Allows traffic forwarding to be segregated and centralised
- Control planes generally converge faster in a hub and spoke than a ring
- It's easier to build and enforce policy this way
- Most traffic “flows” towards a somewhat centralised set of services



Hierarchy Network Functions

- The idea behind hierarchical design is to push these functions into different modules
- Push access, traffic policy, etc., to the edge
- Push forwarding to the core



Forwarding
Core

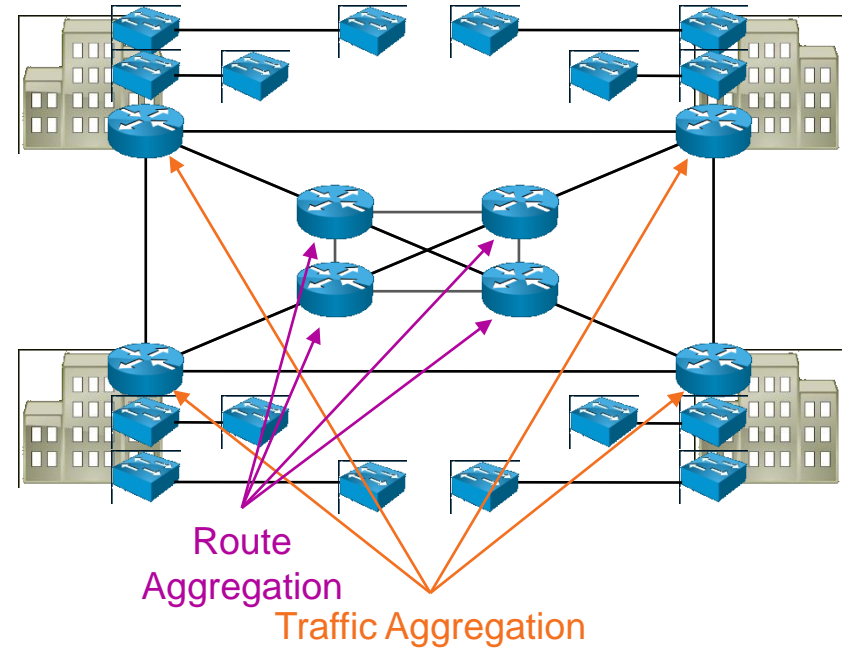
Policy/Route Aggregation
Traffic Aggregation

Edge
Network Access

Hierarchy

Three Layer Example

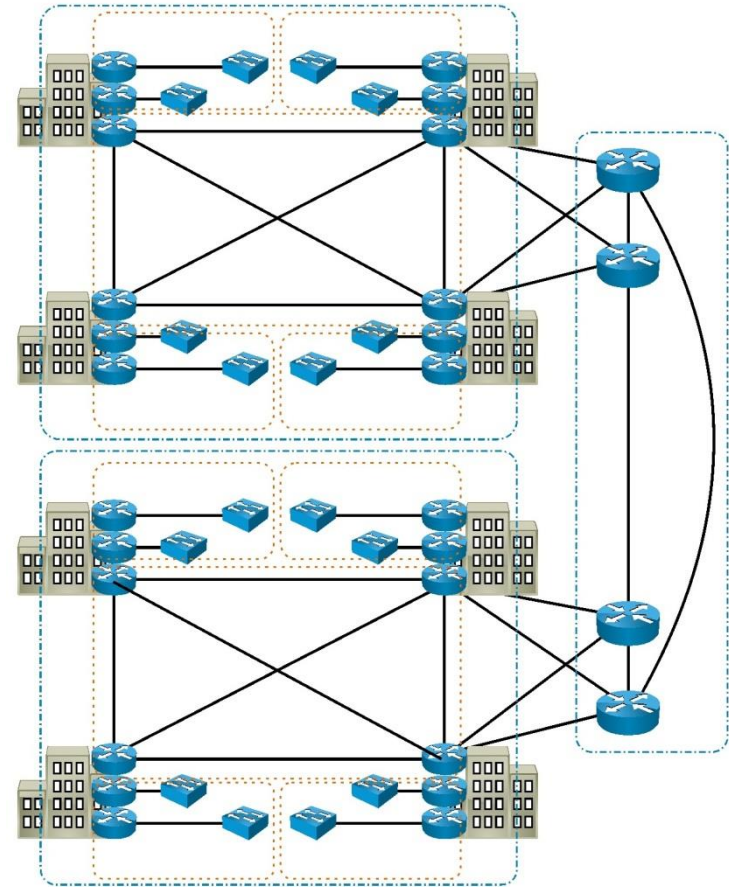
- Traffic aggregation occurs at the access edge routers
- Routing aggregation of some sort at the core edge routers
 - EIGRP route aggregation
 - OSPF ABRs, with or without route aggregation



Hierarchy

Layers Within Layers

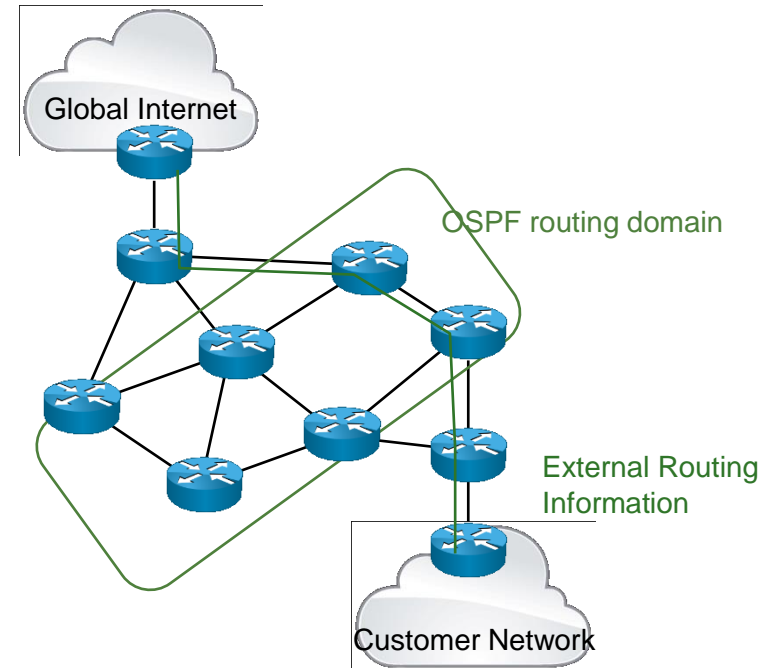
- A series of two layer campuses or data centres may be connected through a core, creating a three layer design
- But, each campus might be administered as a two layer design, and each campus might be treated as a “black box” to the rest of the network



Hierarchy

Layers Within Layers

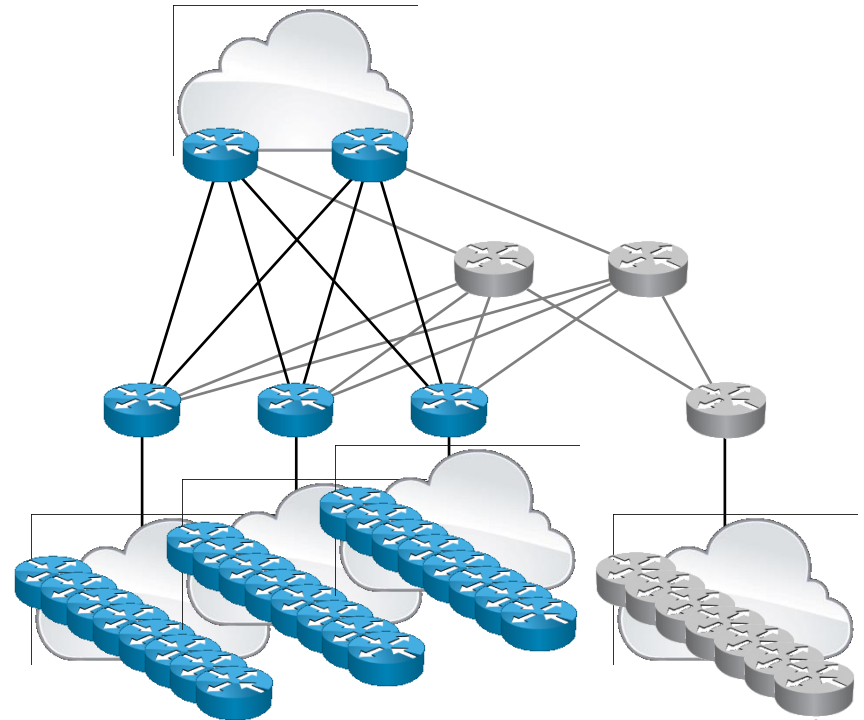
- Service providers actually run layers within layers with their routing structure
- They carry global routes in a separate routing protocol (BGP)
- This separation of control planes is a form of layers within layers
 - But it's more virtualisation than hierarchy



Making the Choice

Adding a Layer

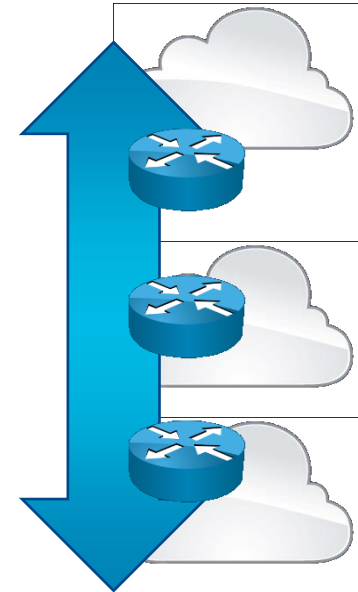
- You need to add a new module to this network
 - Each module is a large scale hub and spoke topology
- Add a new layer, or connect to the existing network?
 - Assume adding to the existing network will entail replacing the second tier routers to support the speeds and feeds



Hierarchy

Key Questions

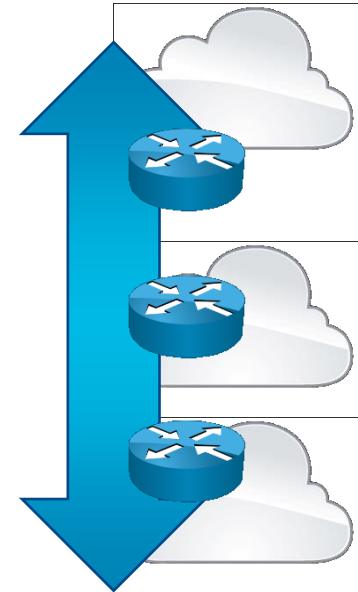
- Evaluating Current Hierarchy
 - Where are the “places” or “modules” in this network?
 - How are they separated from one another?
 - Could I draw a “black box” diagram just showing each “place” with the separation points?
 - Where is each network function placed in this network?



Hierarchy

Key Questions

- Considering Hierarchy within a Design
 - Add a layer? Remove a layer?
 - Can you explain why adding to removing a layer will improve modularity, simplicity, meet a business requirement, or... ??
 - What will adding hierarchy do to traffic flow and stretch?



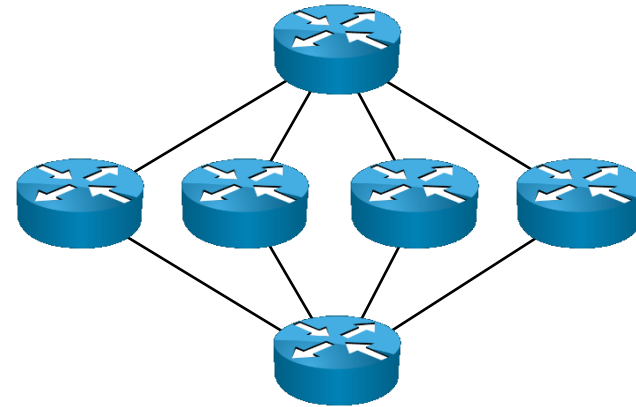
Resilience

- Resilience is the general measure of a network's ability to support business continuity
 - The ability of the network to adapt to changing conditions rapidly with minimal impact on the services running over the network
 - Encompasses “hard” failures, such as network or service outages, expressed through the concept of availability
 - Encompasses “soft” failures, which are much harder to define
 - The network might be “operational,” but....
 - Not usable to solve business problems
 - Not able running optimally
 - Running in a “fragile state,” with little ability to cope with changes beyond the current ones

Resilience

Redundancy

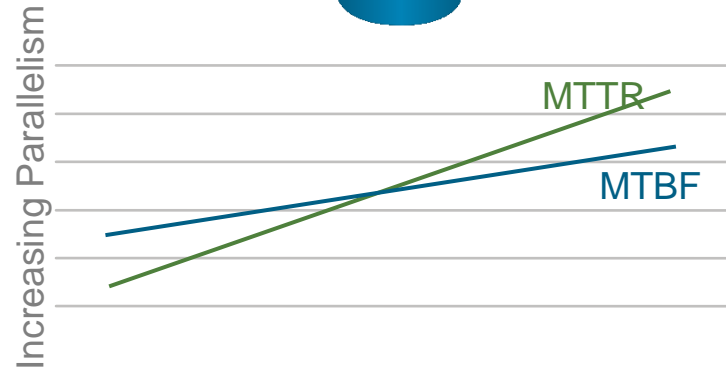
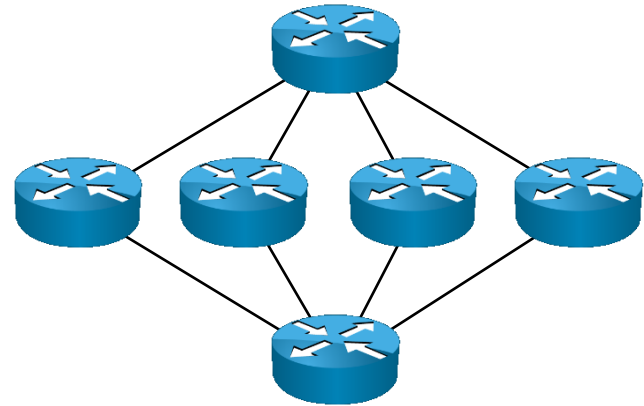
- In principle, redundancy is easy
- Any system with more parallel paths through the system will fail less often
- The problem is a network isn't really a single system
- It's a group of interacting systems
 - Routers
 - Switches
 - Routing Protocols
 - Link Layer Protocols
 - Switching Protocols
 - Etc



Resilience

Redundancy

- In reality, adding paths decreases the MTBF in one layer
 - Increasing parallelism at layer 2 decreases MTBF to nil at some point
- But, it also increases the MTTR at another layer
 - Increasing parallelism increases routing complexity, increasing convergence times
 - The convergence time is directly tied to the MTTR



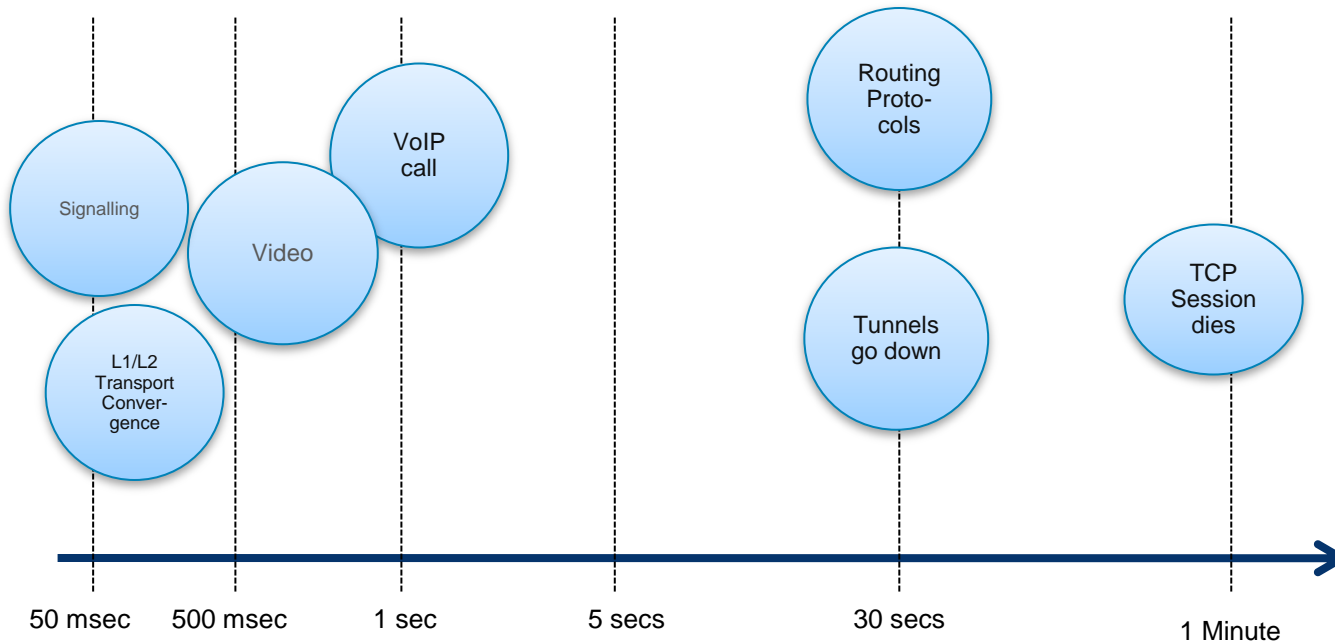
Resilience

Redundancy Key Questions

- Is there redundancy already in the network?
 - If so, it's probably there for a reason
 - Take existing redundancy as a requirement for redundancy in general
 - But use your own judgment in terms of how much redundancy is right
- Don't overdo or underdo
 - Rule of thumb: two paths is enough
 - There are exceptions – CLOS designs, for instance
- What are the requirements?
 - Don't build 500ms convergence if the applications don't demand it
 - Consider the complexity added through additional redundancy

Fast Convergence

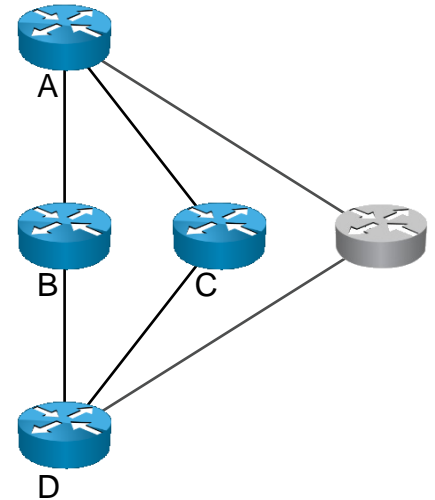
■ Motivation



Making the Choice

Fast Convergence

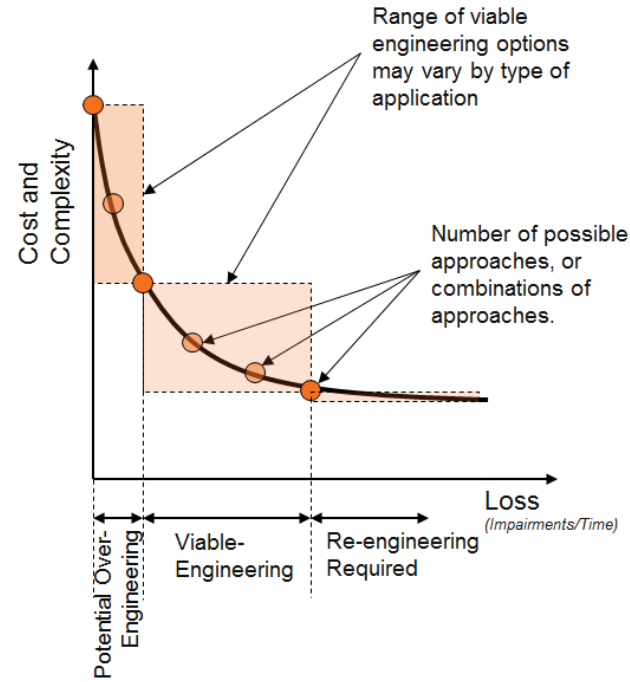
- You need to increase the resiliency across B and C
- Do you...
 - Add another parallel router?
 - Configure fast hellos or BFD?
 - Configure LFAs or EIGRP FS'?
 - Modify the link state flooding timers?



Fast Convergence

Motivation

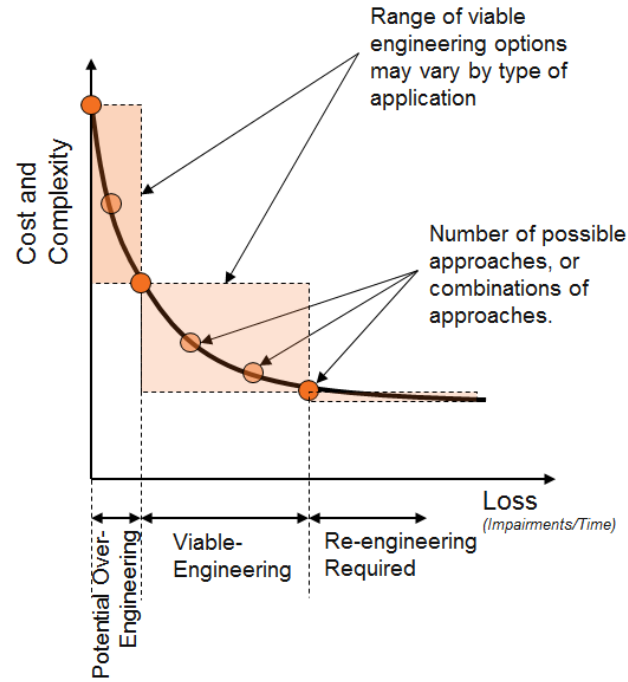
- What are my convergence drivers?
 - What applications require fast convergence?
 - How fast is fast?
- What tools do I have available?
 - Parallel links to provide fast reroute
 - Event driven failure notification



Fast Convergence

Motivation

- What tools do I need to apply to meet my goals?
 - Fast hellos for detection when event driven notifications aren't fast enough or available
 - Fast reroute where alternate paths are available
 - Graceful restart (or nonstop routing) where fast reroute isn't possible



Making the Choice

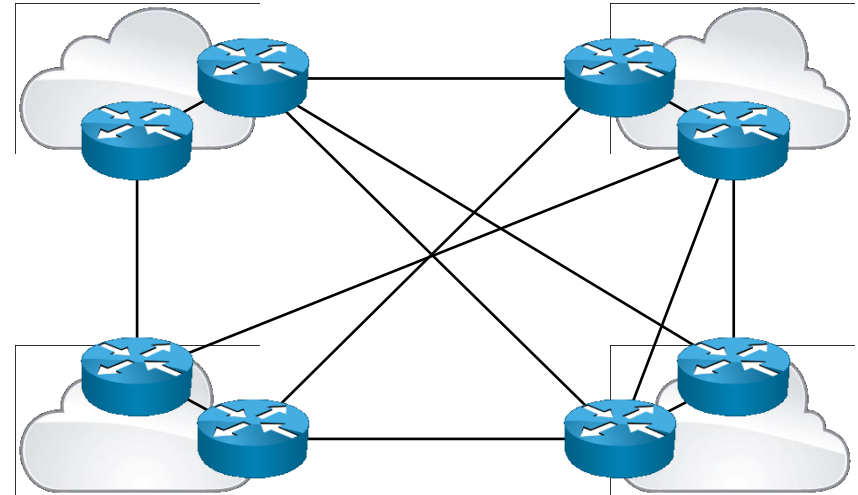
Choosing the Right Core

- During a replace technology work model, you are faced with the problem of replacing the network core
- Which type of core would you choose?
- You're given four options...

Making the Choice

The Single Plane Core Option

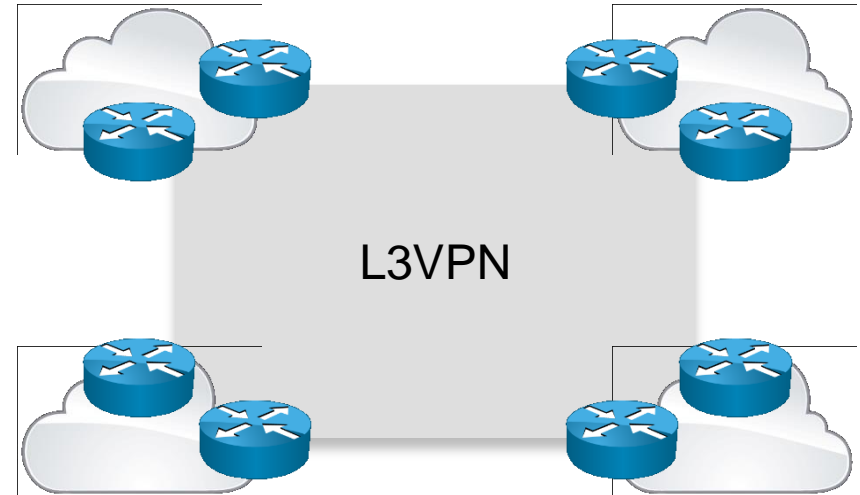
- Design
 - Typically a ring with “cut-throughs” or some form of partial mesh
 - Partial mesh provides resiliency
- Positive Aspects
 - Resiliency through backup links
 - Put links where they are needed to provide optimal routing
 - Support for end-to-end services and virtualisation
- Negative Aspects
 - Resiliency is dependent on a single domain in the control and data planes
 - Can lead to “spaghetti cores,” with little planning beyond traffic level management



Making the Choice

The L3VPN Option

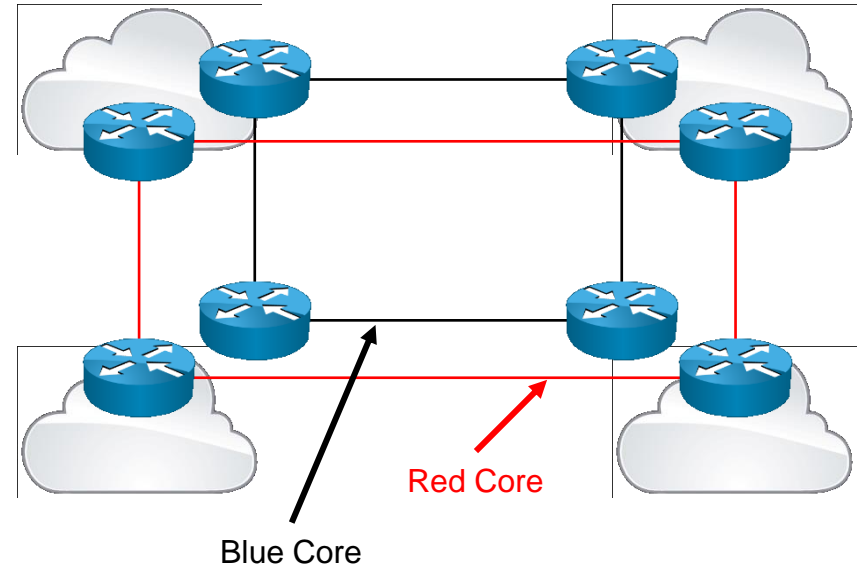
- Design
 - Apparently flat layer 3 domain provided by a service provider
- Positive Aspects
 - SLAs, and hence risk, controlled externally
 - Good low cost option if business goals align with service provider goals
- Negative Aspects
 - Minimal business control over routing infrastructure (traffic engineering, etc.)
 - Ties business plans to service provider performance
 - Very little support for end to end virtualisation, enterprise supported services, etc.



Making the Choice

The Dual Plane Core

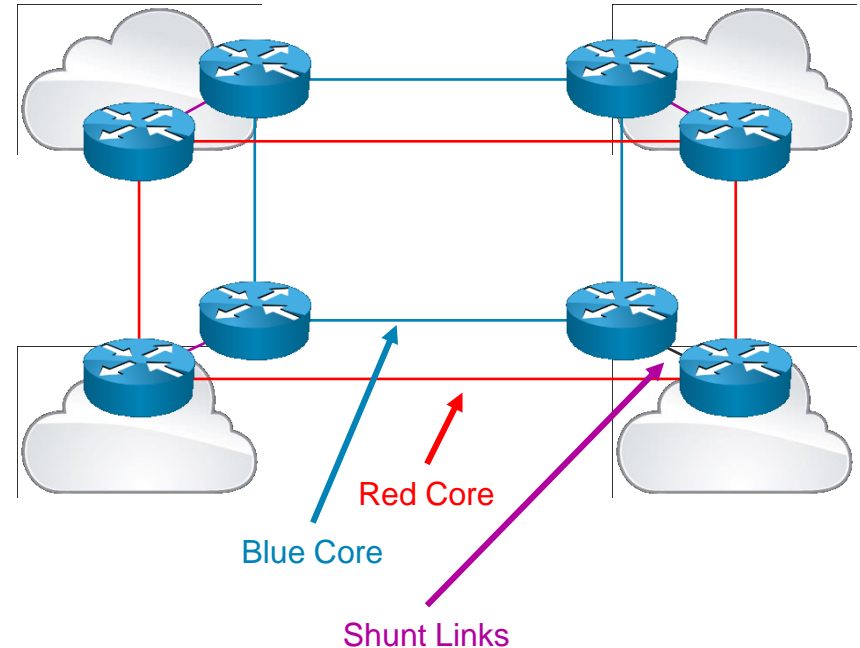
- Design
 - Cores act as independent failure domains
 - Simpler single plane cores placed in parallel
- Positive Aspects
 - Resiliency is by traffic switching to the alternate core
 - Minimal overlap in failure domains
 - Ability to roll out services through a staging process, one core at a time
 - Support for end-to-end services and virtualisation
- Negative Aspects
 - Complex management problems –multiple protocols, etc.
 - Double failures can still wreak havoc on network performance



Making the Choice

The Multiplanar Core

- Design
 - Simpler single plane cores placed in parallel
 - Cores act as semi-independent failure domains
- Positive Aspects
 - Shunt links at “meet points” provide a high cost alternative path
 - Simple to manage, configure, and understand
 - Support for end-to-end services and virtualisation
- Negative Aspects
 - Doesn't provide true multiple control plane failure domains



Making the Choice

Choosing a Core

- Which one would you choose and why?
 - The key is to look for the business drivers
- What are the key business drivers here?
 - Cost
 - Virtualisation within the business (communities of interest)
 - Applications
 - Convergence speed
 - Traffic engineering requirements

Making the Choice

Choosing a Core

- Assume
 - IP Only
 - No traffic engineering requirements
 - High numbers of sites to interconnect
 - Cost is the primary factor
 - Small IT Team
- The best solution would be the L3VPN

Making the Choice

Choosing a Core

- Assume
 - Lots of layer 2 interconnect requirements
 - Strong traffic engineering requirements
 - Communities of interest required (each business unit wants its own virtual network)
- A multiplanar or single plane core with self-deployed MPLS is going to be best

Making the Choice

Choosing a Core

- Assume
 - IP only
 - High availability (6 or 7 9's, or more)
 - The need to roll out new applications quickly without high probabilities of network impact
- A multiplanar or dual plane core is going to be best

Making the Choice

Choosing a Core

- The key is in knowing the pressure points on each core type
 - Single plane cores can be difficult to make clearly resilient
 - L3VPNs hand all routing control and virtualisation over to the provider
 - Multiplanar cores are more expensive to provide optimal linking and routing in
 - Dual plane cores can add a great deal of complexity
 - Especially when virtualisation is involved

Layer 3 Control Plane

Key Questions to Ask

- How does it work?
 - What sorts of signalling are used?
 - Where and how is information aggregated?
- What can it do?
 - What information does this control plane provide?
- What are its limitations?
 - Scaling, convergence speed, etc.
- How does it interact with design?
 - Which topologies work well and which don't?
 - How do I make it meet my convergence requirements?
 - How does it interact with the underlying transport?

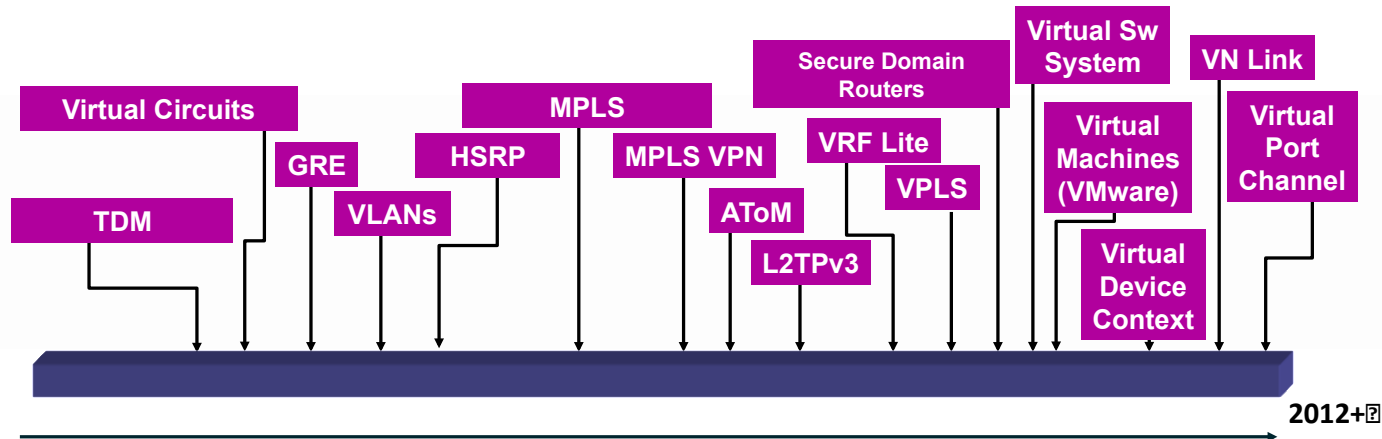


Network Virtualisation

Evolution of “Network” Virtualisation

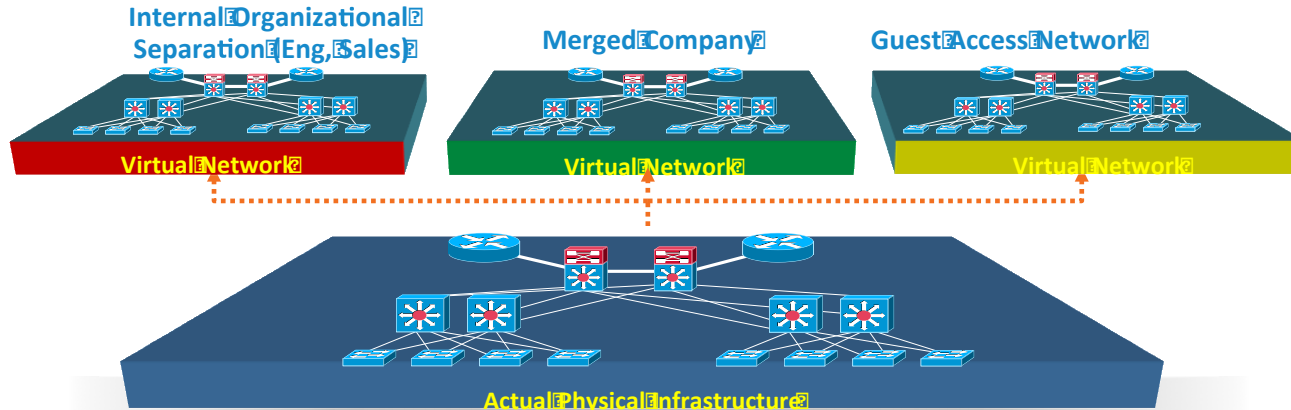
...Means Many Things to Many People

- From TDM, ATM/FR Virtual Circuits in the WAN, to...
- VLANs in the Campus, to...
- Logical/Virtual Routers on routing devices, to...
- Virtual Machines on server clusters in the data centre



What Is Enterprise “Network” Virtualisation?

- Giving one physical network the ability to support multiple virtual networks
- End-user perspective is that of being connected to a dedicated network (security, independent set of policies, routing decisions...)
- Maintains Hierarchy, Virtualises devices, data paths, and services
- Allows for better utilisation of network resources



Why Network Virtualisation?

Key Drivers

- Cost Reduction—allowing a single physical network the ability to support multiple users and virtual networks
- Simpler OAM—reducing the amount of network devices needing to be managed and monitored
- Security—maintaining segmentation of the network for different departments over a single device/Campus/WAN
- High Availability—leverage virtualisation through clustering devices that appear as one (vastly increased uptime)
- Data Centre Applications—require maintained separation, end-to-end (i.e. continuity of virtualisation from server-to-campus-to-WAN) , including Multi-tenant DC's for Cloud Computing
- Common Use Cases
 - Guest Access, Airports, Cloud Computing IaaS, Physical Security Separation, Company Mergers
 - Regulation/Compliance – Health Care (HIPPA), Credit Card (PCI)

Virtualisation

Data Plane Virtualisation

- Conceptually all data plane virtualisation revolves around tunnels
 - Tunnels are a good tool
 - But like any tool, they can be overused
 - When, where, and what's enough?
- Virtualisation can simplify things through indirection
- Too much indirection makes things very complex
 - Where is traffic flowing?
 - Where will it flow if that path fails?
 - How do I measure and plan?

Virtualisation

Common Design Considerations

- Interactions with Applications
 - MTU issues
 - Specific applications (layer 2, layer 3, security issues, etc.)
 - Types of tunnels vs. transport (IP vs. MPLS)
- Interaction with Routing
 - Steering traffic into tunnels
 - Overlay impact on routing complexity
 - Fast convergence
- Scaling
 - The impact of various topologies on scaling
 - Management scaling
 - Routing scaling on tunnel topologies
- Inter-provider
 - Technologies
 - Problems and challenges

Making the Choice

Virtualisation Technology

- Assume
 - There are 50 sites to interconnect
 - The core is already MPLS enabled
 - The networking staff is technically skilled
 - Layer 2 connectivity is a requirement (say for DCI)
- VPLS is going to be the best option

Virtualisation

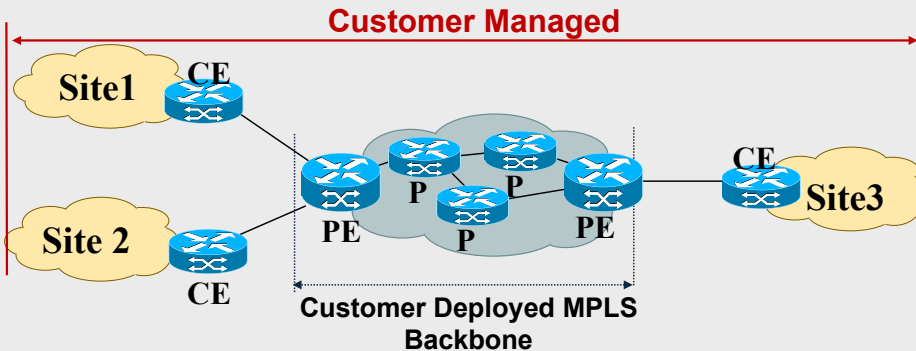
Key Questions

- What type of traffic does this virtualisation technology support?
 - Layer 2? Layer 3?
- How does this interact with the control plane?
 - Does it provide its own control plane at L3 (L3VPN)?
 - Does it provide an L2 control plane?
- How do I get traffic into the virtual topology?
- What about QoS?
- Management and Monitoring?

Self Deployed MPLS vs. SP L3 Managed

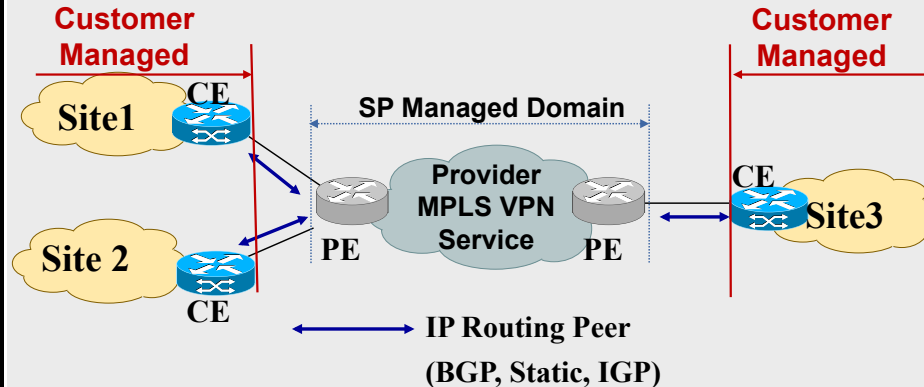
Network Deployment Options over the WAN

Self Deployed MPLS



- Customer manages and owns:
 - IP routing, provisioning
 - Transport for PE-P, P-P, PE-CE
 - SLA's, to "end" customer
 - QoS, Traffic Engineering
- Allows customer full control E2E

SP Managed IP/VPN Service

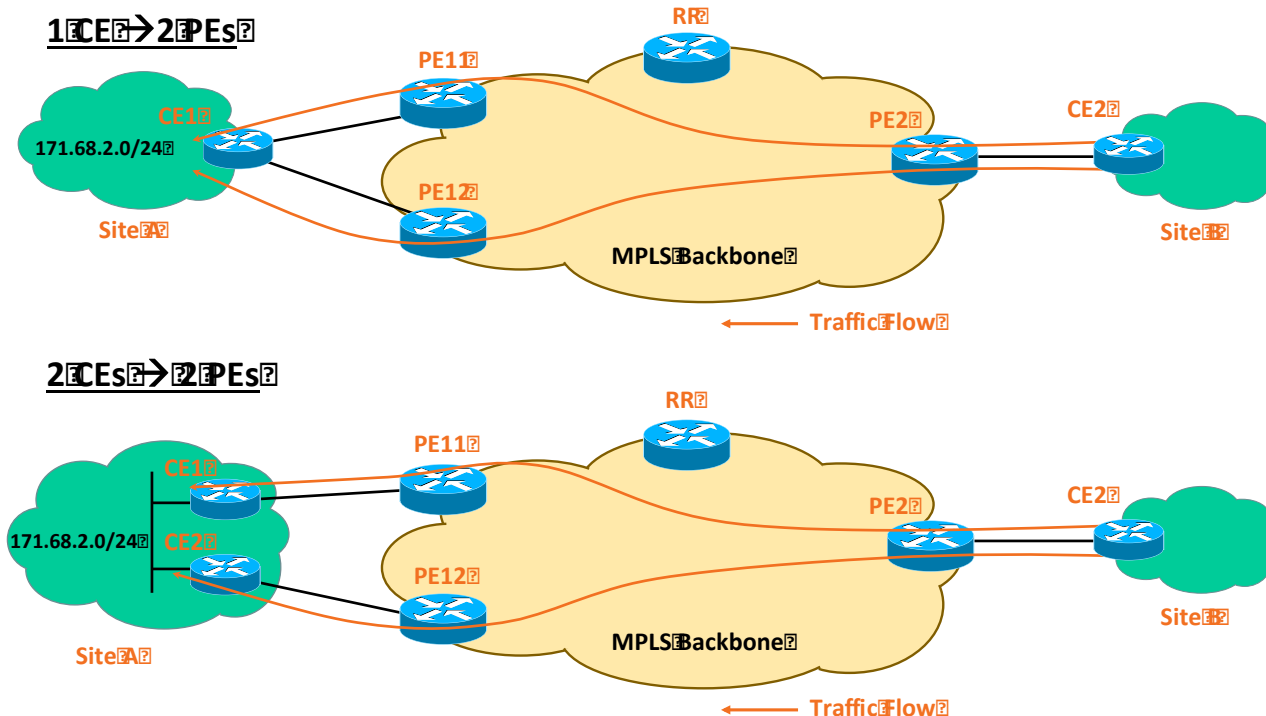


*** No Labels Are Exchanged with the SP**

- CE Routers owned by customer
- PE Routers owned by SP
- Customer "peers" to "PE" via IP
- Exchanges routing with SP via routing protocol (or static route)
- Customer relies on SP to advertise routes to reach other customer CEs

L3 VPN Services

Load-sharing of VPN Traffic



L3 VPN

Best Practices

- Use RR to scale BGP; deploy RRs in pair for the redundancy
 - Keep RRs out of the forwarding paths (scalability, resource optimisation)
- Choose AS format for RT and RD i.e., ASN: X
 - Reserve first few 100s of X for the internal purposes such as filtering
- Consider unique RD per VRF per PE,
 - Helpful for many scenarios such as multi-homing, hub&spoke etc.
 - Helpful to avoid add-path, shadow RR etc.
- Don't use customer names (V458:GodFatherNYC32ndSt) as the VRF names; nightmare for the NOC.
 - Consider v101, v102, v201, v202, etc. and Use VRF description for naming
- Utilise SP's public address space for PE-CE IP addressing
 - Helps to avoid overlapping; Use /31 subnetting on PE-CE interfaces

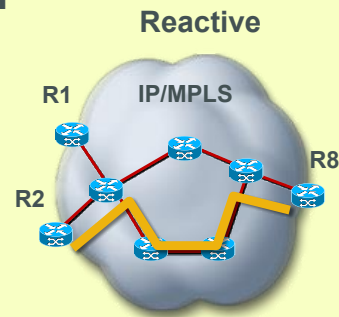
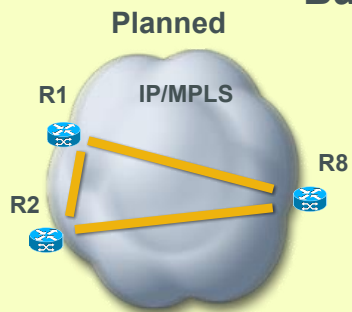
L3 VPN

Best Practices (cont'd)

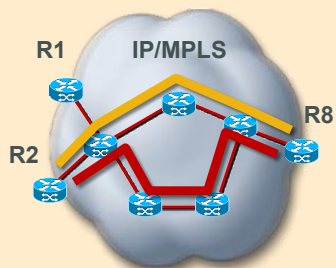
- Limit number of prefixes per-VRF and/or per-neighbour on PE
 - Max-prefix within VRF configuration; Suppress the inactive routes
 - Max-prefix per neighbour (PE-CE) within OSPF/RIP/BGP VRF

MPLS TE Deployment Models

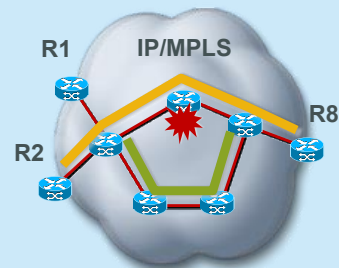
Bandwidth Optimization



Point-to-Point SLA



Protection



Traffic Engineering

Key Questions

- What is the impact of aggregation?
 - Consider stretch
- What devices and flows need consistent/symmetric paths
 - For stateful inspection and other purposes
- What is my goal in directing traffic?
 - Meeting Quality of Service objectives?
 - Providing failover?
 - Management bandwidth & link usage?
- What techniques can I use to solve these problems?
 - Allowing more specific routing information, MPLS/TE, etc.

Traffic Engineering

Key Questions

- TE – especially MPLS/TE – adds a lot of complexity to a network
- Think before deploying
 - Can I solve this problem another way?
 - What is the level of complexity using this solution verses some other solution?
 - What will the impact mean time between mistakes (MTBM) and other resilience characteristics



QoS – Quality of Service

Quality of Service

Interpret QoS Metrics

- DiffServ and IntServ Architectures
- Tools for implementing PHBs
- Class Starvation
- Interaction with Transport
- Interaction with Applications

Quality of Service

Defined

- To the end user
 - User's perception that their applications are performing properly
 - Voice – No drop calls, no static
 - Video – High quality, smooth video
 - Data – Rapid response time
- To The Network Manager
 - Maximise network bandwidth utilisation while meeting performance expectations
 - Control Delay – The finite amount of time it takes a packet to reach the receiving endpoint
 - Jitter – The difference in the end-to-end delay between packets.
 - Packet Loss – relative measure of the number of packets that were not received compared to the total number of packets transmitted.

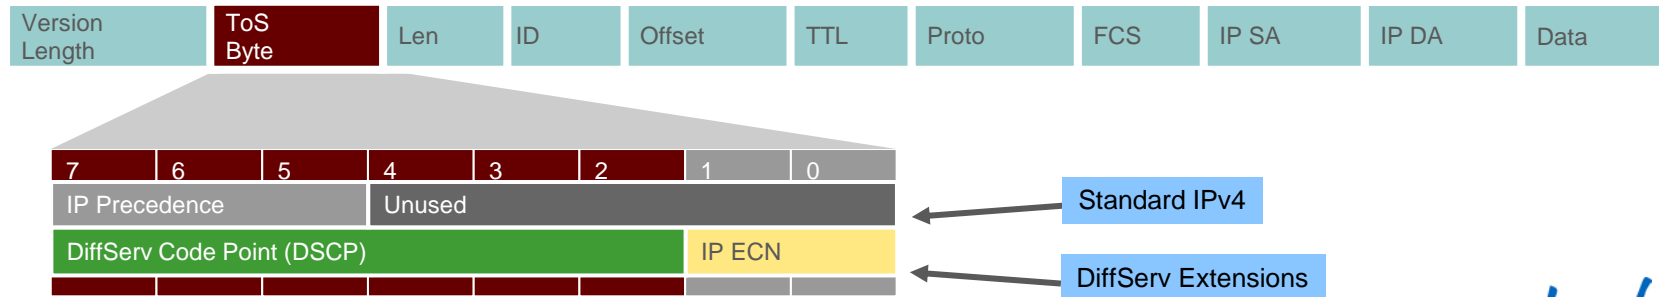
Quality of Service

Objectives

- Clearly define the organisational objectives
 - Protect voice? Video? Data?
DoS/worm mitigation?
- Assign as few applications as possible to be treated as “mission-critical”
- Seek executive endorsement of the QoS objectives prior to design and deployment
- Determine how many classes of traffic are required to meet the organisational objectives
 - More classes = More granular service-guarantees

Quality of Service IPv4 Classification

- IPv4
 - Three most significant bits of ToS byte are called IP Precedence (IPP)—other bits unused
- DiffServ
 - Six most significant bits of ToS byte are called DiffServ Code Point (DSCP)—remaining two bits used for flow control
- DSCP is backward-compatible with IP precedence



QoS with GRE, MPLS over GRE - Example

ToS Reflection

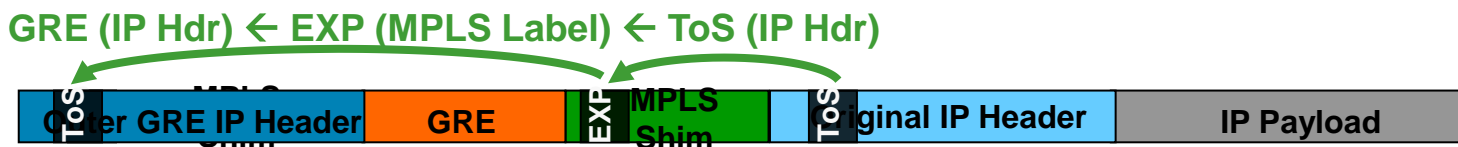
GRE Header



GRE Header with ToS Reflection



MPLS over GRE Header with ToS Reflection



- Router will copy original ToS marking to outer GRE header
- For MPLS over GRE, the EXP marking is copied to the outer header of the GRE tunnel
- This allows the IPv4 “transport” to perform QoS on the multi-encapsulated packet

QoS Deployment Models in a Virtualised Environment

■ **Aggregate Model**

- A common QoS strategy is used for all VRFs
 - i.e. same marking for voice, video, critical data, best effort
- Allows identical QoS strategy to be used with/without virtualisation

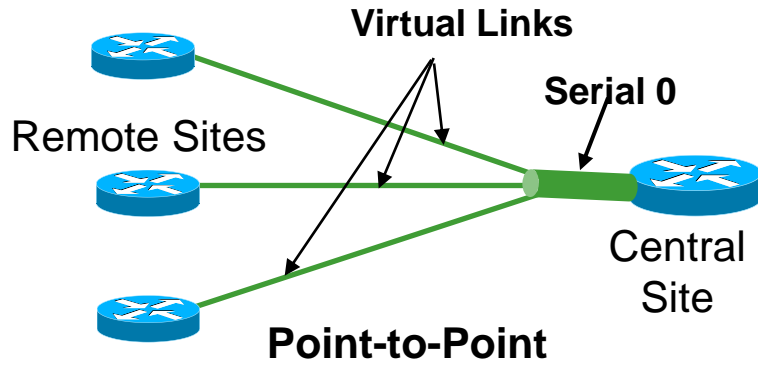
■ **Prioritised VRF Model**

- Traffic in some VRFs are prioritised over other VRFs (i.e. Production over Guest VRF)

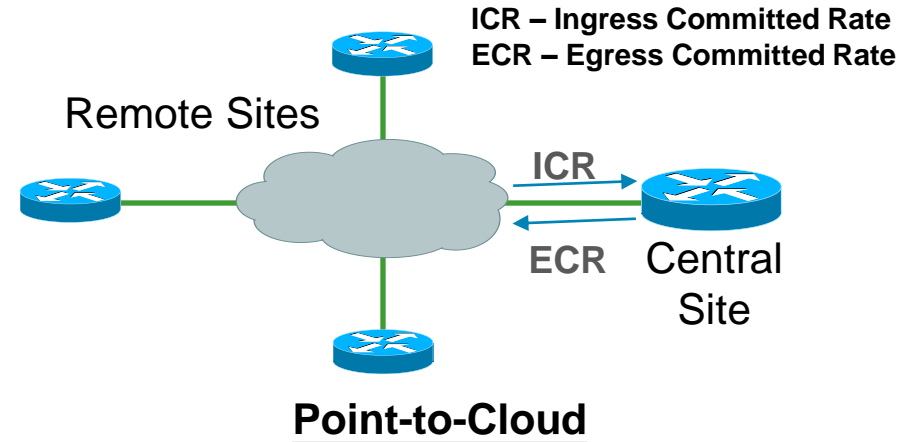
Aggregate vs. Prioritized Model

- Following the “**Aggregate Model**” Allows the Identical QoS Strategy to Be Used With/Without Network Virtualisation

QoS Models



- Connection oriented service
- Logical/physical interfaces/connections
- Logical - GRE, FR/ATM/Ethernet VC
- Virtualisation options connection oriented:
 - VRF-Lite over P2P GRE
 - MPLS VPN over P2P GRE
 - * VRF-Lite over DMVPN
 - * 2547 over DMVPN



- No point-to-point (site-to-site) guarantees
- Any site can transmit up to ICR into the cloud
- Any site can receive up to ECR from the cloud
- SLA offers guarantees for conforming traffic
- Virtualisation options leveraging point-to-cloud:

- MPLS VPNs over mGRE

QoS for Virtualisation – Summary

- **Aggregate** QoS model is the simplest and most straight forward approach
(Recommended)
- Simplification using the **Aggregate** model recommends:
 - Traffic class marking identical to non virtualisation scheme
 - Traffic class marking identical between VRF's
 - Leverage H-QoS on virtualised interfaces (GRE, .1Q)
 - Router dynamically copies ToS→EXP→ToS (GRE)
- Prioritised VRF model can be used to prefer traffic originating in one VRF over another (e.g. guest access, mission critical apps)
- **Summary:** Consider implementing the same QoS approach that is used for non-virtualised, when deploying QoS in virtualised enterprise network designs

Quality of Service

Key Questions

- What does this QoS mechanism do?
- What's the impact on delay, jitter, etc?
- Does it impact some applications more than others?
- How's it normally used?

Quality of Service

Key Questions

- What do each of the applications on this network require?
- What will happen if I place application A and application B in the same queue?
- Where is the best place to mark traffic?
- What does each queuing mechanisms do?
 - Does it drop traffic?
 - Does it delay traffic?

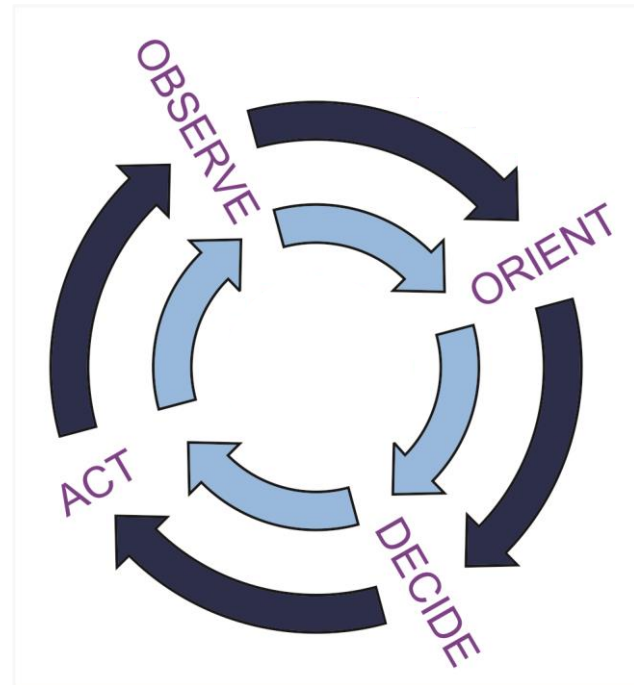


Network Management

The OODA Loop

Management and Security Background

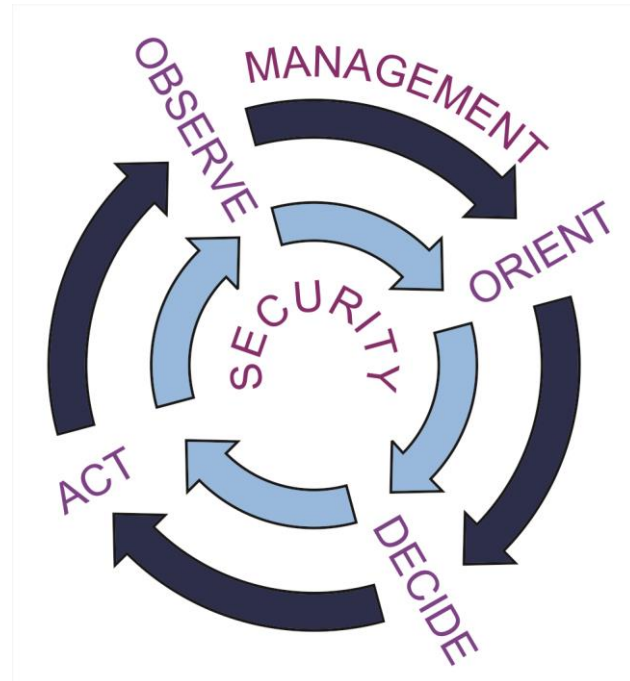
- The OODA loop is a way of looking at the process of reacting to...
 - Events
 - Threats
- It's a useful model for breaking down the management and security processes
 - The OODA loop isn't explicitly covered on the exam
 - But the principles here will help you handle the security scenarios



The OODA Loop

Management and Security Background

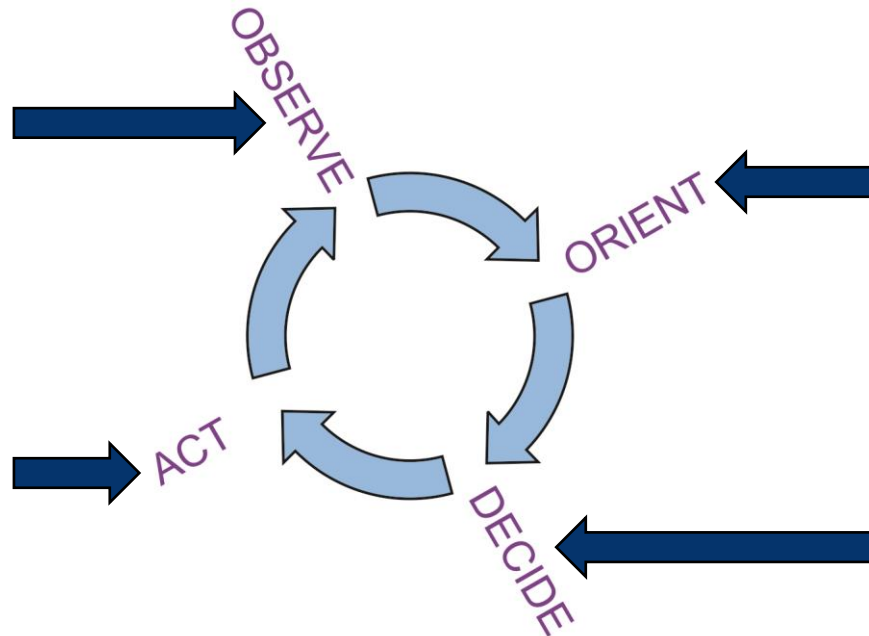
- Management
 - A “slower” loop
 - Reacts to organic threats
 - Changes in the business, technology, etc.
 - Business drivers
- Security
 - A “faster” loop
 - Reacts to inorganic threats
 - Attacks designed to deny service, obtain access, discover information, etc.



Management


Network Documentation
Baseline Performance
Baseline Utilisation
Change Analysis
Business Ecosystem
Business Processes
Technology Ecosystem

Design Modification
Replace Technology
Inject Technology
Add Services
Policy Modification



Best Fit Analysis
Root Cause Analysis
Business Trends
Best Practices
Case Studies

Shape to Models
Technology Trends
Business Trends
Best Practices
Case Studies
Change Management
Worst Case Analysis

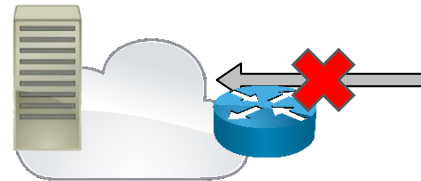


Security

Security

Aspects

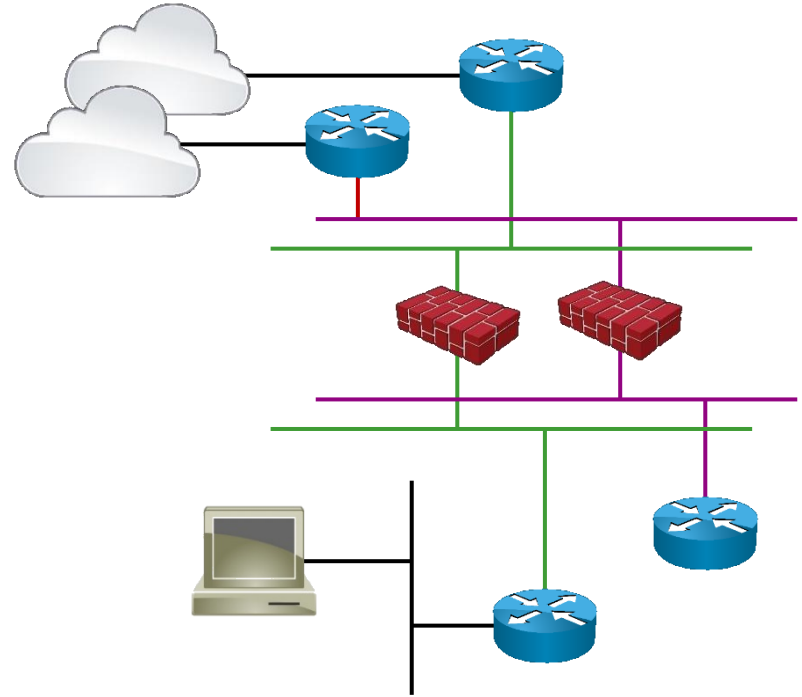
- Protecting Devices
 - Routers and switches are vulnerable to the same attacks as hosts
- Protecting the System
 - DDoS and other attacks attempt to compromise the infrastructure
- Protecting the Data
 - Protect service and application information carried across the network



Security

Crunchy on the Outside...

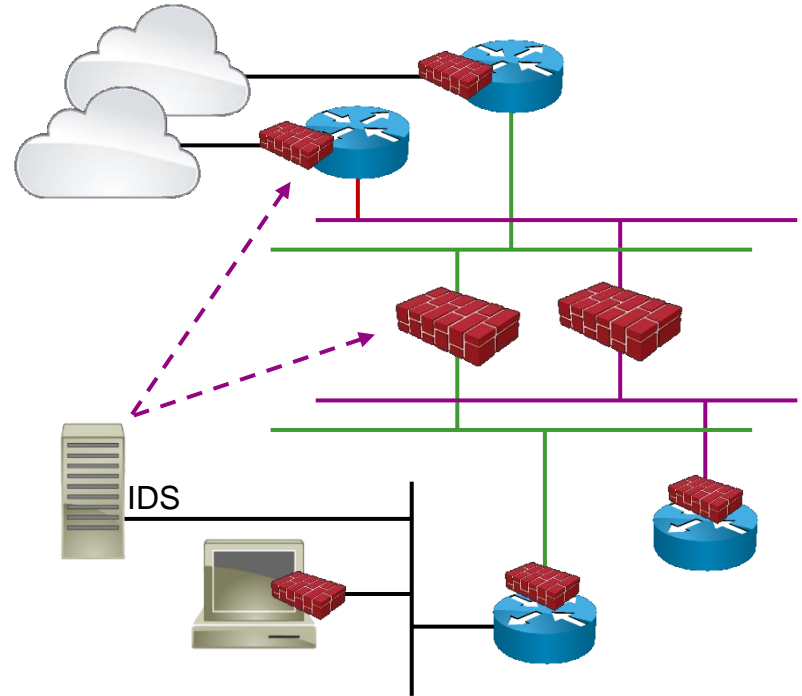
- A solid DMZ was once the best you could do in security design for your network...
- Crunchy on the outside, chewy in the middle



Security

Crunchy Through and Through

- Every device has security
- Automatic (fast) feedback loops
- Crunchy through and through



Security

Practical Points to Protect the System

- Observe
 - What? Where? How?
 - What can I put in place before hand to help me observe?
- Orient
 - Baselines, current trends
 - Network management feeds your ability to orient
- Decide
 - What? Where? How?
- Act
 - What?
 - Policies and procedures taken from the network management side
 - Current trends
 - Where?
 - Network modularisation provides this answer
 - How?
 - Network management should facilitate rapid policy dispersion

Security

Protecting the Data

- Key questions
 - Is the information confidential?
 - Is there an inside threat?
- Answering these questions will provide the answers...
 - What sort of protection do I need to provide?
 - IPsec?
 - Making certain traffic of certain classes doesn't pass across specific links?
- Virtualisation isn't a solid answer to protect data
 - There's no encryption here, just another tag
 - Virtual circuits across public networks aren't *really* more secure than just sending the data natively across the public network

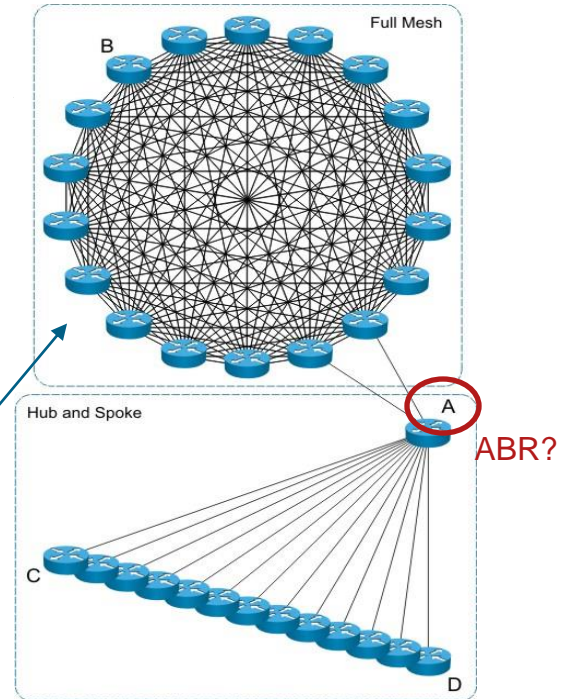


CCDE Written Examples

Routing Examples

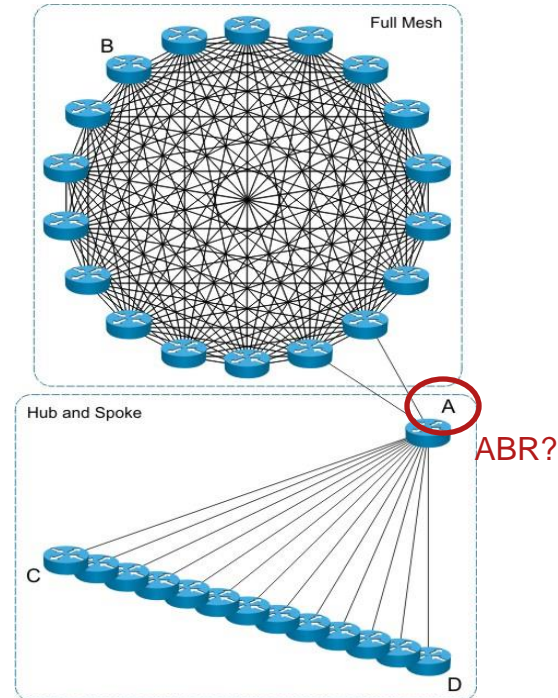
- Assume you cannot aggregate in this network
- What justification would you give for configuring an ABR?

1970's Spirograph network design



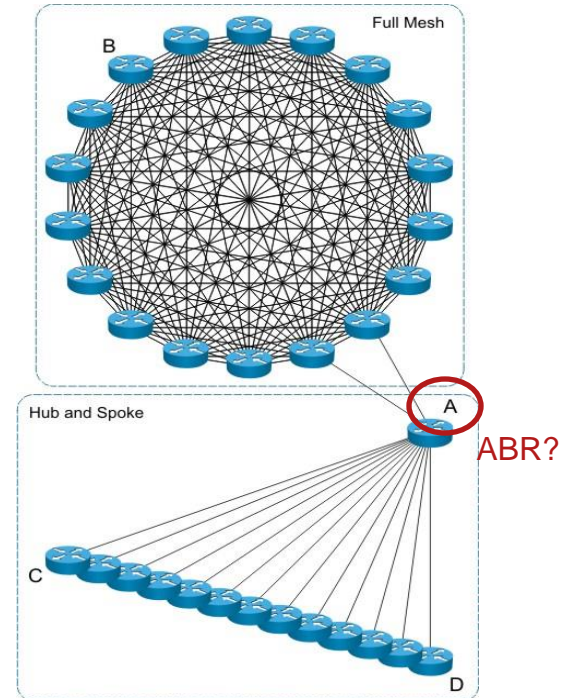
Routing Examples

- To reduce Router B's routing table size
- To reduce the complexity of the full mesh
- To reduce the impact of Router B failing at Router C
- To reduce the total time required to run SPF at Router A



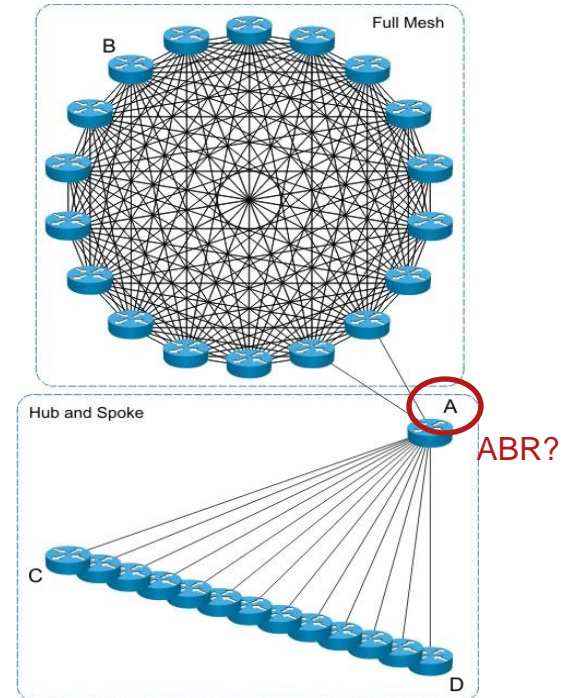
Routing Examples

- To reduce Router B's routing table size
 - An ABR doesn't aggregate reachability unless you configured aggregation
 - The routing table is made up of routes, which is reachability information
 - An ABR at A wouldn't impact the size of the routing table at B



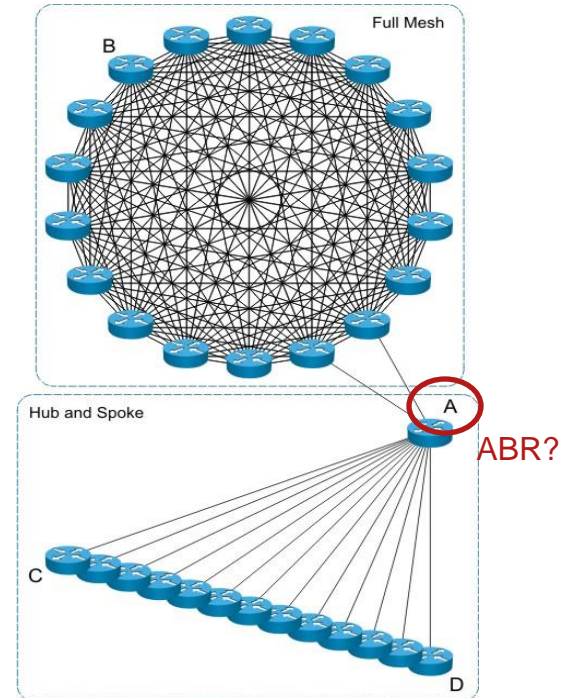
Routing Examples

- To reduce the complexity of the full mesh
 - Router A doesn't belong to the full mesh, so there's no apparent relationships
- To reduce the total time required to run SPF at Router A
 - A's single SPF run is replaced by two SPF runs
 - Is this faster, or slower?
It's probably a wash



Routing Examples

- To reduce the impact of Router B failing at Router C
 - In a single flooding domain, B failing would cause a full SPF run at C
 - If A is an ABR, C will only receive a type 3 change (if any change at all), which only causes a partial SPF

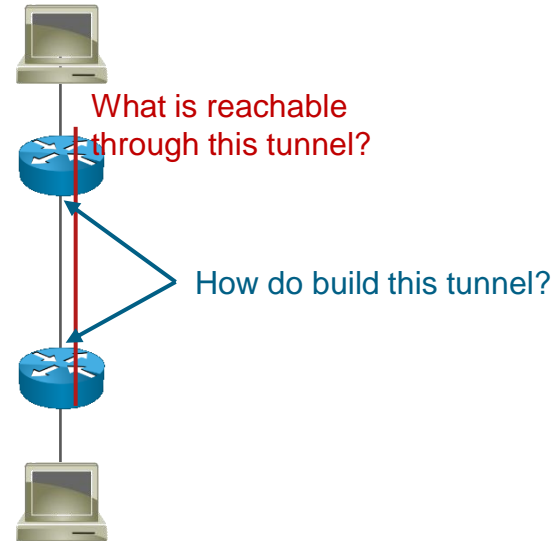


Tunnelling Examples

- When using any form of IPSec over GRE tunnelling (for instance, DMVPNs) over a public or private network, how many routing instances will you need to provide full reachability?
 - One
 - Two
 - Three
 - Four

Tunnelling Examples

- We know we need a control plane to find out what is reachable through the tunnel
 - This is one routing instance
- We also need to find the tunnel endpoints
 - This is a second routing instance
 - You might not own this routing instance



Tunnelling Examples

- We could have more routing instances
 - But we don't need them
 - We need two
- Two is the correct answer
- Why do we care?
 - Complexity is spread across the two control planes
 - Moving complexity through tunnelling really doesn't decrease overall complexity
 - You need to remember to plan for both control planes when designing tunneled solutions

Quality of Service Examples

- Which would you deploy to increase the throughput of multiple TCP traffic flows on a single link?
 - Head of queue dropping
 - Traffic Policing
 - Weighted RED
 - Traffic Shaping

Quality of Service Examples

- What is the basic problem with multiple TCP flows on a single link?
 - Synchronisation of the flow windows and restart
- All three of these do nothing about synchronisation (or encourage it)
 - Head of queue dropping
 - Traffic Policing
 - Traffic Shaping

Quality of Service Examples

- Weighted RED
 - Weighted RED prevents multiple TCP flows from synchronising
 - This allows underlying links to be used at a higher rate of overall utilisation
- Why do we care?
 - This is an interaction between layer 3 and transport behaviour required by specific applications

Network Management Examples

- Which two of the following are true of Syslog?
 - Syslog always provides a wider variety of information than SNMP Traps
 - Syslog is more reliable than SNMP traps because it is transported over TCP
 - Syslog provides severity information by default, SNMP traps do not
 - Syslog information can be easily read, SNMP traps must be decoded

Network Management Examples

- Syslog always provides a wider variety of information than SNMP Traps
 - While this is generally true, it's not always true
- Syslog is more reliable than SNMP traps because it is transported over TCP
 - SNMP can use either UDP or TCP as a transport
- Syslog provides severity information by default, SNMP traps do not
 - While SNMP traps may have severity information coded in them, the severity is not part of the base protocol
- Syslog information can be easily read, SNMP traps must be decoded
 - You can only read SNMP traps if you're part of the Borg, or you spent your early years working on Wellfleet routers...

Network Management Examples

- SNMP vs. Syslog
 - Why do we care?
 - As a network designer, you need to know which tools can do what
 - Specifying the correct network management tool for the job is an important part of network design



CCDE Practical Examples

Background

Business Information

- A single network infrastructure to lower cost

Longhorn Infrastructure Services and Provisioning

Longhorn (LISP) provides customer support services for a wide range of clients throughout the world. Their business model is to provide a semi-customizable database environment for large companies who would like to offload telephone and web support to an outside agent.

Their business is centered in the San Francisco area, with two primary locations, with a data center and small corporate office building located on one campus, and two buildings containing a call center on another campus. The two campuses are connected via four 10GB Metro Ethernet circuits.

They provide these services at a low cost by aggregating product usage, call pattern, and other information onto a single network infrastructure. They sell this aggregated information off as a “byproduct” of their normal operations.

Background

SLA - Impacts on Network Design

- Privacy protection for customers
- No client-to-client leakage
 - Routing separation
 - Traffic separation
- Resilience and fast convergence

LISP Service Levels

LISP offers their customers a high degree of service, including the following elements:

- Total privacy for all customer and client data. No data from one client's customer is ever transferred to another client.
- Total privacy of customer and client bases. No client can ever discover what other companies are using LISP services, unless the client themselves provides permission for LISP to release this information. The information protected to prevent the disclosure of clients includes, but is not limited to routing and DNS tables.

From time to time, LISP may ask a client to provide a reference.

- The LISP network provides five 9's of reliability, so clients can access customer databases, statistics on call volumes per product, and other data, at any time. Traffic from a client to the databases containing their information is never mixed with other client data traffic.
- No more than one percent of calls made into the LISP call center will be dropped for technical reasons.
- The network bandwidth resources are divided into two pools and two appropriate DSCP markings are used. Bandwidth is carved out for ALL VoIP traffic based upon amount of all VoIP calls and that traffic is given priority treatment. All the other traffic uses the rest of the available bandwidth. It is customers responsibility to mark all the traffic with appropriate DSCP marking and LISP honors those markings.
- Each call center agent will be fully trained, and have a complete understanding of the products being offered.
- Each client's web service is hosted in a highly reliant virtualized server environment. The web interfaces offered to clients are fully customizable, so each client can make the LISP web server appear to be a part of their web infrastructure. DNS redirection and virtualization services can be used to make the URL appear to be part of the client's web site, rather than a separate service, if desired.

Background

Internal Standards

- Traffic must pass through a firewall between the client and the data centre
- Client routes carried in BGP, not in local IGP
- Traffic carried outside the LISP network will be carried in SSL or IPsec

Longhorn Infrastructure Services and Provisioning

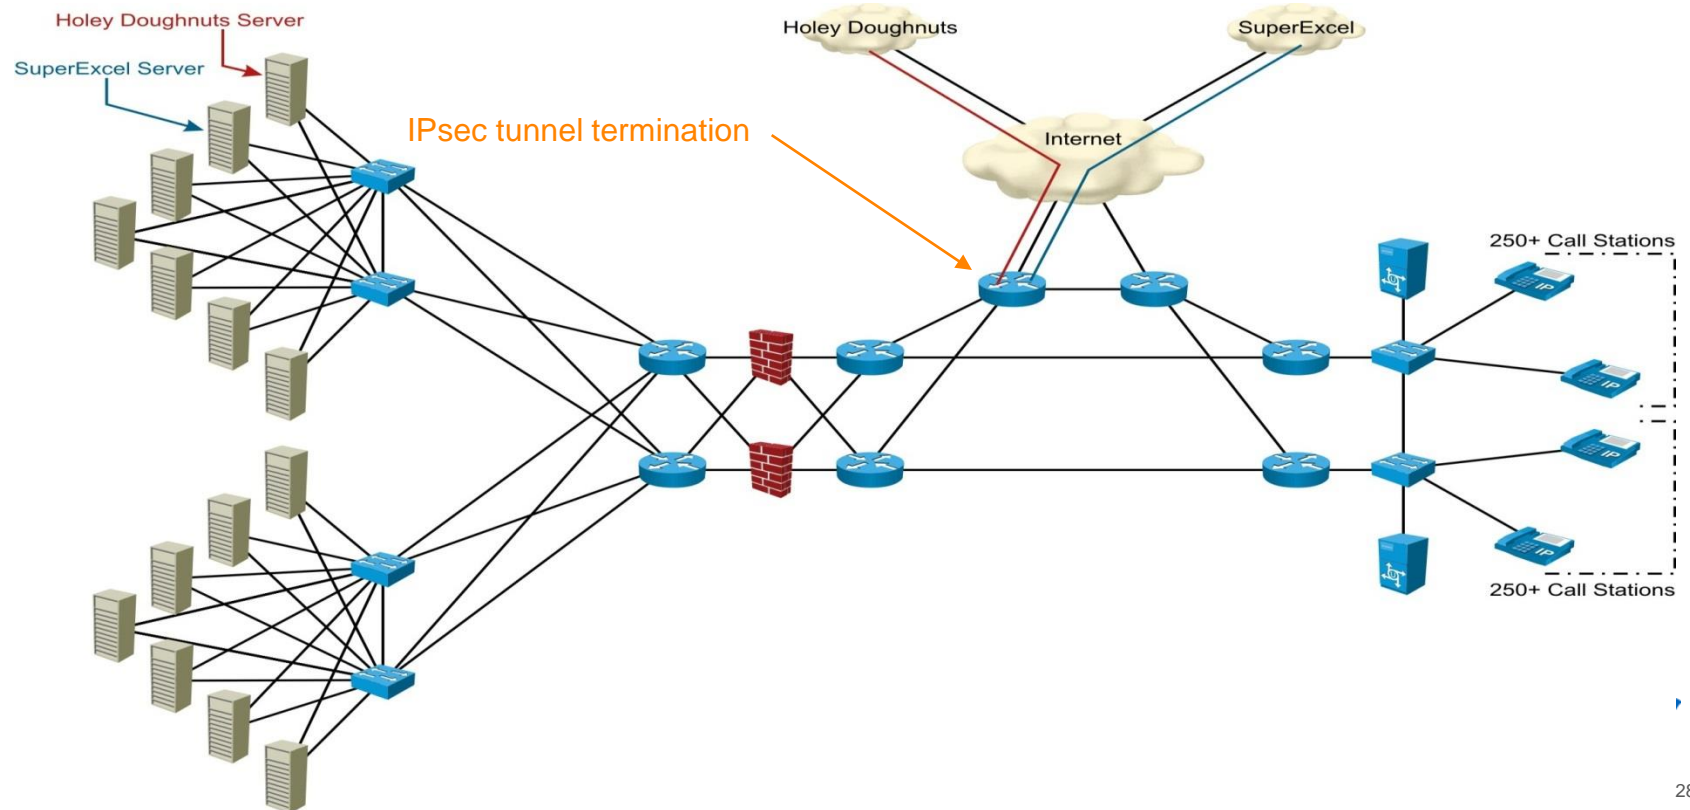
Internal Network Information

LISP has several security policies which are implemented to prevent the LISP network from being breached or attacked.

- All traffic passing from a client connection to a server must pass through a stateful layer 3 filtering device. Traffic must be checked for known viruses, and conform to other commonly used security policies.
- Client routes cannot be mixed with internal LISP IGP routes. Client routes must be carried in an exterior gateway protocol, at this time BGP, to prevent changes in the client network from impacting internal network operations.
- Traffic carried over a public network will be transported in a secured tunnel to ensure private client information is not compromised in any way. LISP will support both IPsec and SSL tunnels, as the customer requests, but will not allow open flow of traffic between the LISP edge and the customer network.

Background

Network Diagram



The Problem

- Holey Doughnuts has done a security audit
- Their auditors discovered they can reach servers not belonging to Holey Doughnuts!
 - They specifically mention being able to reach the servers used by SuperExcel
 - They note there are routes they don't recognise on their CE

Item 1

Root

- Is Holey Doughnut's ability to see routes from SuperExcel, and reach SuperExcel's servers, a problem that needs to be addressed by LISP?
 - Yes
 - No



Item 1

Root Explained

LISP Service Levels

LISP offers their customers a high degree of service, including the following elements:

- Total privacy for all customer and client data. No data from one client's customer is ever transferred to another client.
- Total privacy of customer and client bases. No client can ever discover what other companies are using LISP services, unless the client themselves provides permission for LISP to release this information. The information protected to prevent the disclosure of clients includes, but is not limited to routing and DNS tables.

From time to time, LISP may ask a client to provide a reference.

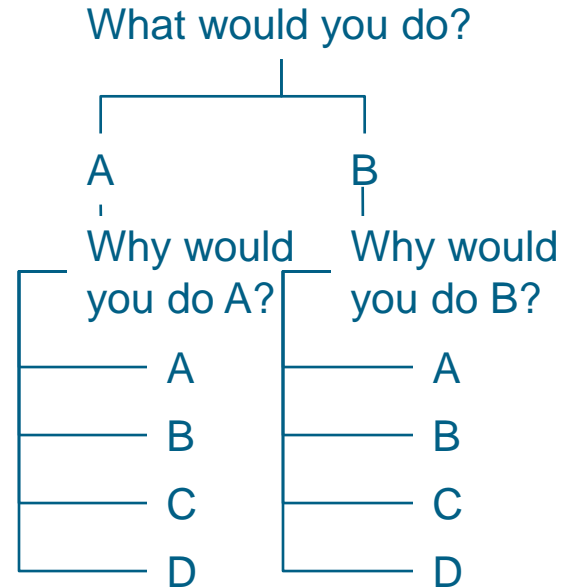
- The LISP network provides five 9's of reliability, so clients can access customer databases, statistics on call volumes per product, and other data, at any time. Traffic from a client to the databases containing their information is never mixed with other client data traffic.
- No more than one percent of calls made into the LISP call center will be dropped for technical reasons.

The answer is yes

Branching Questions

Explained

- A large number of the questions are “branching”
- Two parts: What would you do, and why?
- If you go down the wrong branch, it’s possible to get all wrong answers on the next question
- We’ll go down the “no” branch for this question to illustrate



Item 1

“No” Branch





- This Is Not a Problem LISP Needs to Solve
- Why don't you think this is a problem that needs to be addressed by LISP?
 - Holy Doughnut's engineers cannot get to SuperExcel's data, so this is not a breach of the SLA
 - Holy Doughnuts should address this by reconfiguring their route and packet filters
 - The LISP Service Level Agreement with Holy Doughnuts does not cover this situation
 - The LISP engineering team has bigger things to worry about than this problem

Item 1

“No” Branch Explained

- This Is Not a Problem LISP Needs to Solve

- Why don't you think this is a problem that needs to be addressed by LISP?

-  – Holy Doughnut's engineers cannot get to SuperExcel's data, so this is not a breach of the SLA
-  – Holy Doughnuts should address this by reconfiguring their route and packet filters
-  – The LISP Service Level Agreement with Holy Doughnuts does not cover this situation
-  – The LISP engineering team has bigger things to worry about than this problem

Item 1

“Yes” Branch

- This Is a Problem LISP Needs to Solve
- Why do you think this problem needs to be addressed by LISP?
 - LISP should not be leaking the SuperExcel routes to their other clients
 - LISP should not be allowing telnet traffic from Holey Donuts to SuperExcel’s servers
 - LISP is not fulfilling their service level agreement correctly

Item 1

“Yes” Branch

- This Is a Problem LISP Needs to Solve
- Why do you think this problem needs to be addressed by LISP?
 - ✗ – LISP should not be leaking the SuperExcel routes to their other clients
 - LISP should not be allowing telnet traffic from Holey Donuts to SuperExcel’s servers
 - LISP is not fulfilling their service level agreement correctly

Item 1

“Yes” Branch Explained

LISP Service Levels

LISP offers their customers a high degree of service, including the following elements:

- Total privacy for all customer and client data. No data from one client’s customer is ever transferred to another client.
- Total privacy of customer and client bases. No client can ever discover what other companies are using LISP services, unless the client themselves provides permission for LISP to release this information. The information protected to prevent the disclosure of clients includes, but is not limited to routing and DNS tables.

From time to time, LISP may ask a client to provide a reference.

- The LISP network provides five 9’s of reliability, so clients can access customer databases, statistics on call volumes per product, and other data, at any time. Traffic from a client to the databases containing their information is never mixed with other client data traffic.
- No more than one percent of calls made into the LISP call center will be dropped for technical reasons.

We need traffic separation, not just route filters

Item 1

“Yes” Branch

- This Is a Problem LISP Needs to Solve
- Why do you think this problem needs to be addressed by LISP?
 - ✗ – LISP should not be leaking the SuperExcel routes to their other clients
 - ✗ – LISP should not be allowing telnet traffic from Holey Donuts to SuperExcel’s servers
 - LISP is not fulfilling their service level agreement correctly

Item 1

“Yes” Branch Explained

LISP Service Levels

LISP offers their customers a high degree of service, including the following elements:

- Total privacy for all customer and client data. No data from one client’s customer is ever transferred to another client.
- Total privacy of customer and client bases. No client can ever discover what other companies are using LISP services, unless the client themselves provides permission for LISP to release this information. The information protected to prevent the disclosure of clients includes, but is not limited to routing and DNS tables.




From time to time, LISP may ask a client to provide a reference.

- The LISP network provides five 9’s of reliability, so clients can access customer databases, statistics on call volumes per product, and other data, at any time. **Traffic from a client to the databases containing their information is never mixed with other client data traffic.**
- No more than one percent of calls made into the LISP call center will be dropped for technical reasons.

We need traffic separation, not just packet filters

Item 1

“Yes” Branch

- This Is a Problem LISP Needs to Solve
- Why do you think this problem needs to be addressed by LISP?
 -  - LISP should not be leaking the SuperExcel routes to their other clients
 -  - LISP should not be allowing telnet traffic from Holey Donuts to SuperExcel's servers
 -  - LISP is not fulfilling their service level agreement correctly

Item 2

- Which solution would resolve this problem in a way that supports both the SLA and the LISP business model?
 - Using Policy Based Routing to separate customer traffic
 - Creating a physical per customer to separate customer traffic
 - Provisioning a separate virtual topology per customer
 - Configuring a separate IGP instance per customer to separate customer routes



Item 2

- Which solution would resolve this problem in a way that supports both the SLA and the LISP business model?
 - ✗ – Using Policy Based Routing to separate customer traffic
 - Creating a physical per customer to separate customer traffic
 - Provisioning a separate virtual topology per customer
 - Configuring a separate IGP instance per customer to separate customer routes

We need traffic separation, not just packet filters

Item 2

- Which solution would resolve this problem in a way that supports both the SLA and the LISP business model?

-  – Using Policy Based Routing to separate customer traffic
-  – Creating a physical per customer to separate customer traffic
- Provisioning a separate virtual topology per customer
- Configuring a separate IGP instance per customer to separate customer routes

Item 2

- A single network infrastructure to lower cost

Longhorn Infrastructure Services and Provisioning




Longhorn (LISP) provides customer support services for a wide range of clients throughout the world. Their business model is to provide a semi-customizable database environment for large companies who would like to offload telephone and web support to an outside agent.

Their business is centered in the San Francisco area, with two primary locations, with a data center and small corporate office building located on one campus, and two buildings containing a call center on another campus. The two campuses are connected via four 10GB Metro Ethernet circuits.

They provide these services at a low cost by aggregating product usage, call pattern, and other information onto a single network infrastructure. They sell this aggregated information off as a "byproduct" of their normal operations.

Item 2





- Which solution would resolve this problem in a way that supports both the SLA and the LISP business model?

-  – Using Policy Based Routing to separate customer traffic
-  – Creating a physical per customer to separate customer traffic
 - Provisioning a separate virtual topology per customer
-  – Configuring a separate IGP instance per customer to separate customer routes

We need traffic separation as well a routing separation

Item 2

- Which solution would resolve this problem in a way that supports both the SLA and the LISP business model?

-  – Using Policy Based Routing to separate customer traffic
-  – Creating a physical per customer to separate customer traffic
-  – Provisioning a separate virtual topology per customer
-  – Configuring a separate IGP instance per customer to separate customer routes

We need traffic separation as well a routing separation

Item 3


- What additional piece of information would you need to determine which virtualisation service you would use to separate customer traffic in this network?
 - The tunnelling technologies the firewalls support
 - The tunnelling technologies the call managers support
 - Tunnelling technology used to connect to customer CEs

Item 3

Explained

- From time to time during the test, you will be asked what further information you need to complete the task
- If you already have a specific piece of information, don't ask for it!
 - This might not match “real life” as well as you'd like
 - But the question is trying to evaluate whether or not you know the information you need to solve the problem
- This is one of the analytical skills the test is designed to measure

Item 3

- What additional piece of information would you need to determine which virtualisation service you would use to separate customer traffic in this network?
 - The tunnelling technologies the firewalls support
 - The tunnelling technologies the call managers support
 -  – Tunnelling technology used to connect to customer CEs

Item 3

External Tunnels

- Traffic carried outside the LISP network will be carried in SSL or IPsec

Longhorn Infrastructure Services and Provisioning

Internal Network Information

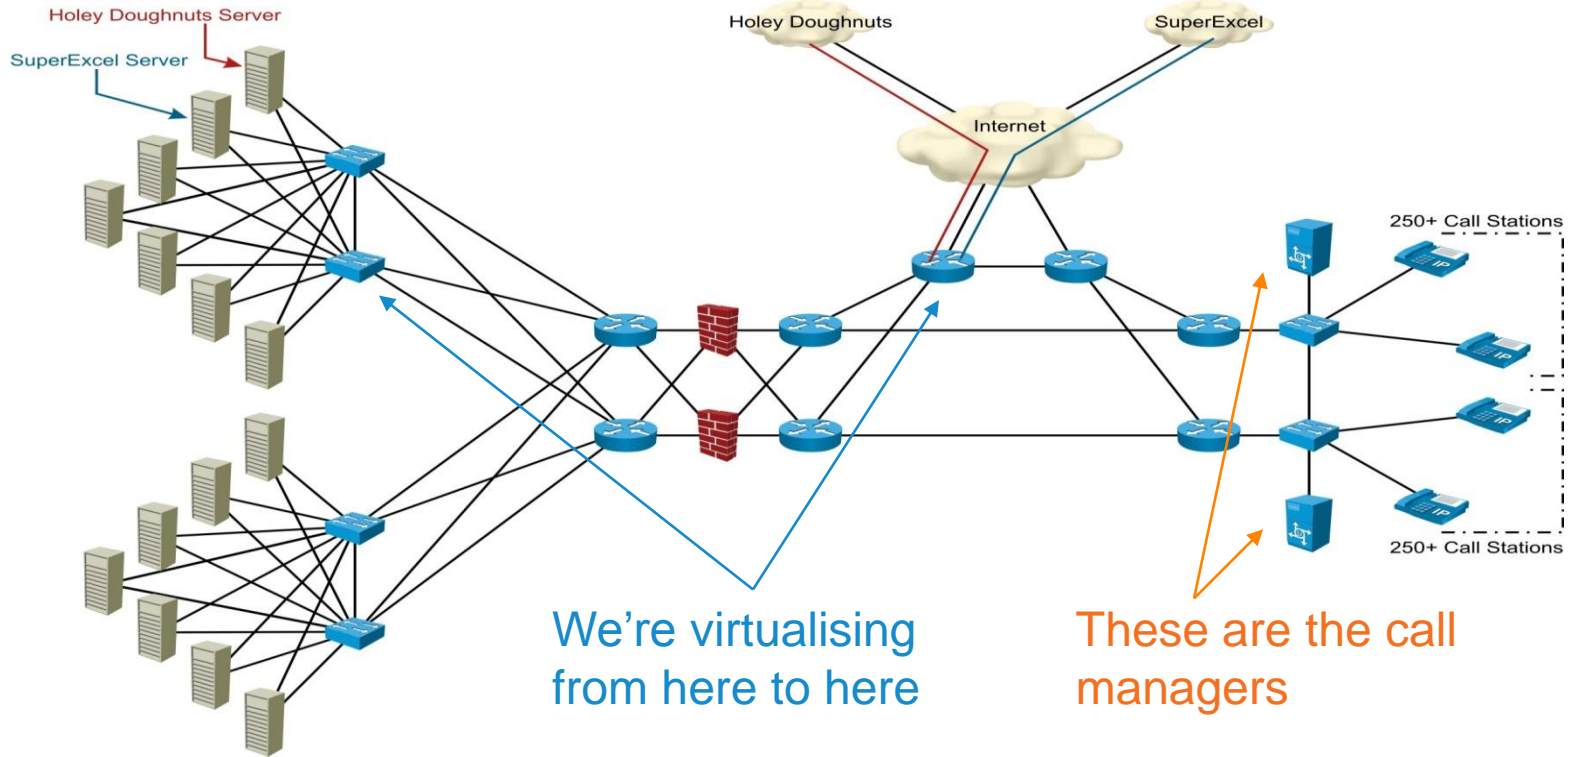
LISP has several security policies which are implemented to prevent the LISP network from being breached or attacked.

- All traffic passing from a client connection to a server must pass through a stateful layer 3 filtering device. Traffic must be checked for known viruses, and conform to other commonly used security policies.
- Client routes cannot be mixed with internal LISP IGP routes. Client routes must be carried in an exterior gateway protocol, at this time BGP, to prevent changes in the client network from impacting internal network operations.
- Traffic carried over a public network will be transported in a secured tunnel to ensure private client information is not compromised in any way. LISP will support both IPsec and SSL tunnels, as the customer requests, but will not allow open flow of traffic between the LISP edge and the customer network.

Item 3

- What additional piece of information would you need to determine which virtualisation service you would use to separate customer traffic in this network?
 - The tunnelling technologies the firewalls support
 - ~~X~~ – The tunnelling technologies the call managers support
 - ~~X~~ – Tunnelling technology used to connect to customer CEs

Item 3

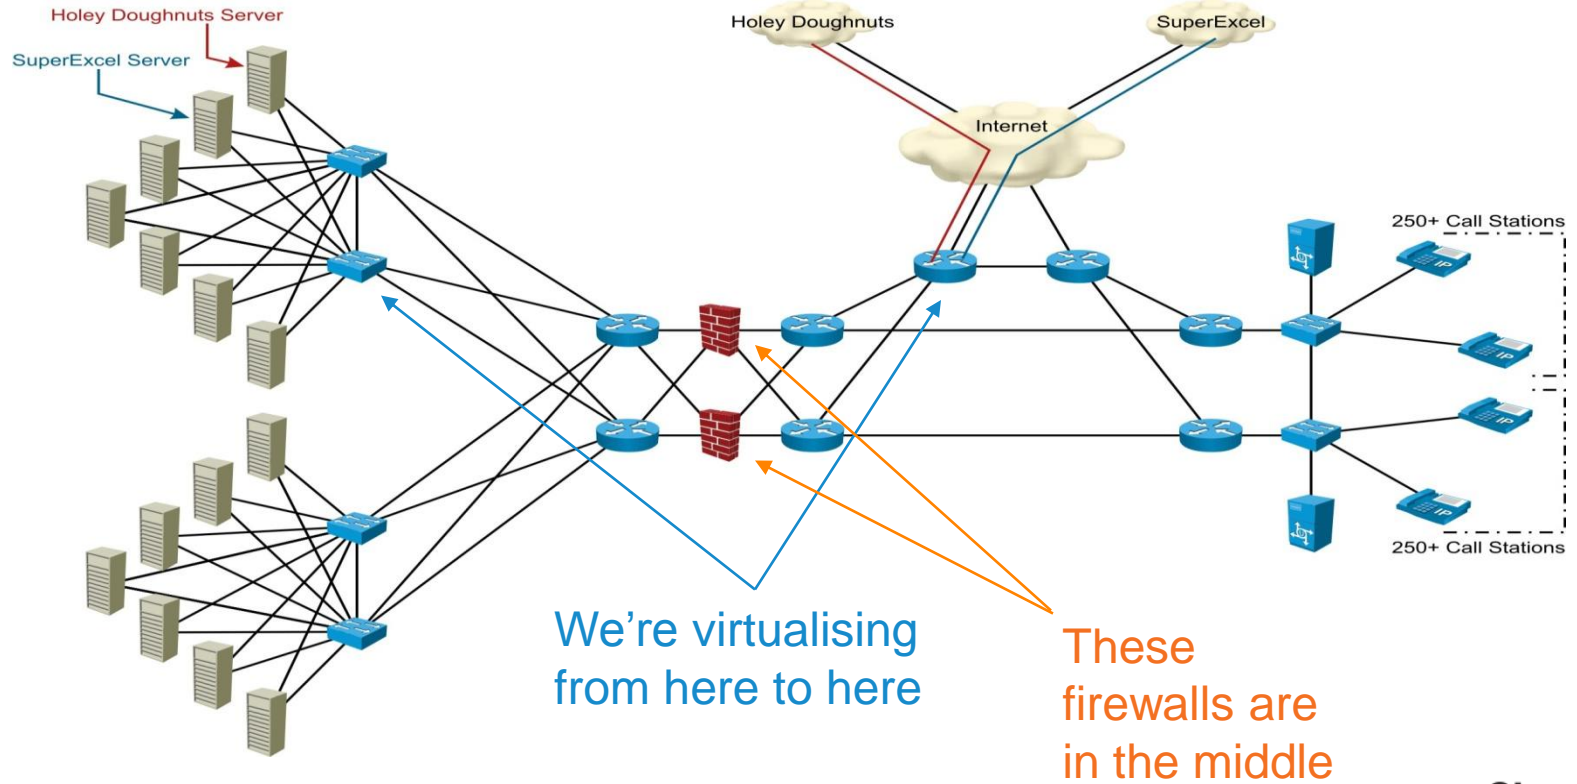


Item 3

- What additional piece of information would you need to determine which virtualisation service you would use to separate customer traffic in this network?

- ✓ – The tunnelling technologies the firewalls support
- ✗ – The tunnelling technologies the call managers support
- ✗ – Tunnelling technology used to connect to customer CEs

Item 3



Additional Information

- At various points throughout the test, new documents will be supplied that provide new information
- In this case, we need to know what tunnelling technologies these firewalls will support

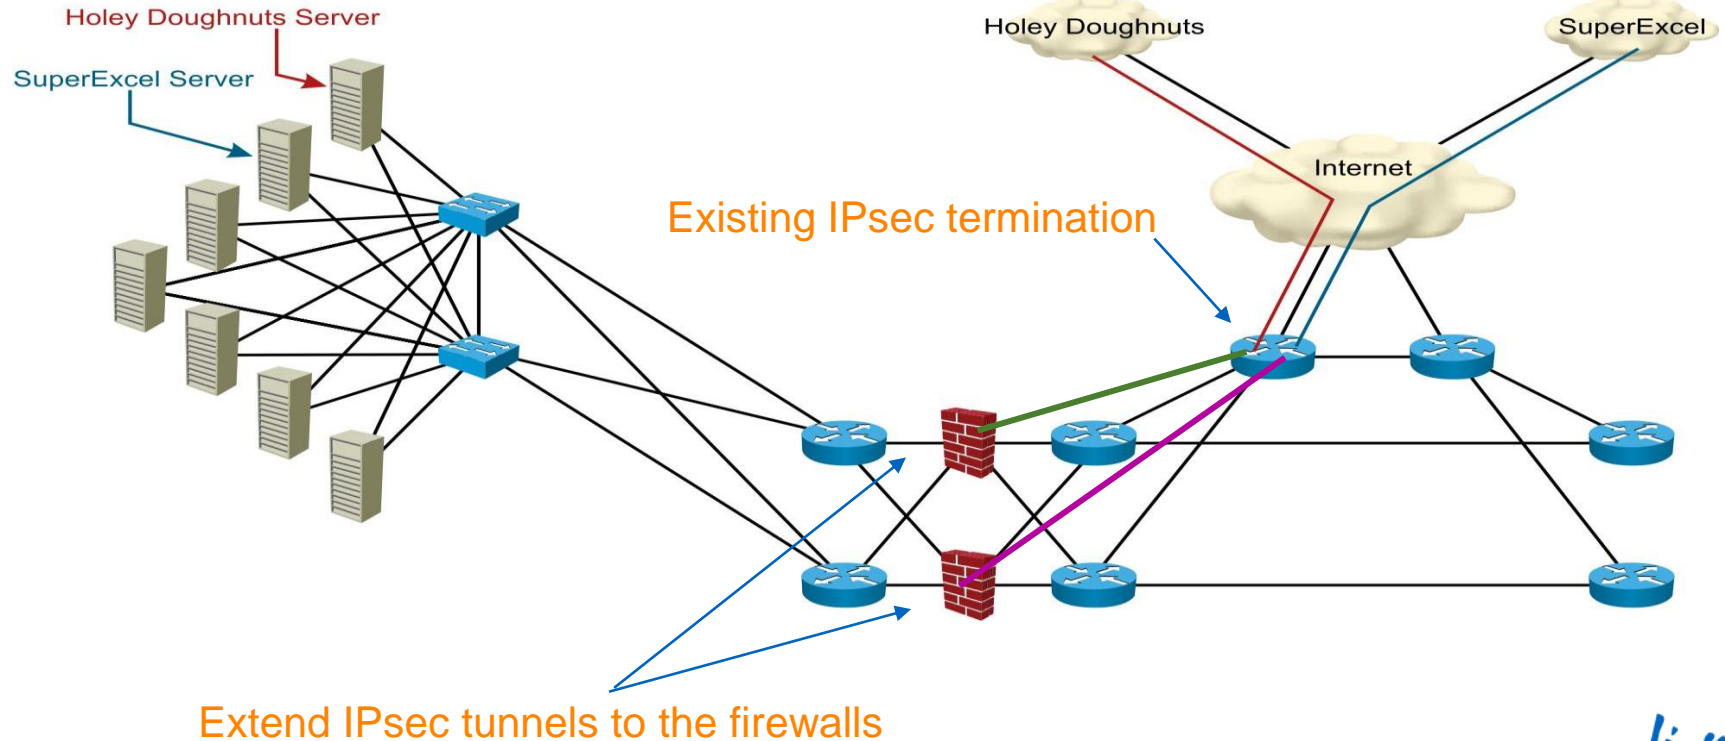
From: Bruce
Sent: Monday 12:12 PM
To: Amy
Cc: [lsip-eng](#)
Subject: Possible Solution

I took a cruise through the documentation for the firewall, and discovered they can switch traffic through multiple virtual instances within a single hardware platform based on an MPLS or .1q tag, and they can terminate IPsec tunnels, but they can't switch traffic out of any sort of IP based tunnel through multiple virtual instances in the device.

The servers can terminate IPsec tunnels, but doing so would be a major burden on their processors. The firewalls do have a hardware assist blade built just to offload IPsec processing, and dump the tunnels into a .1q VLAN on their back side.

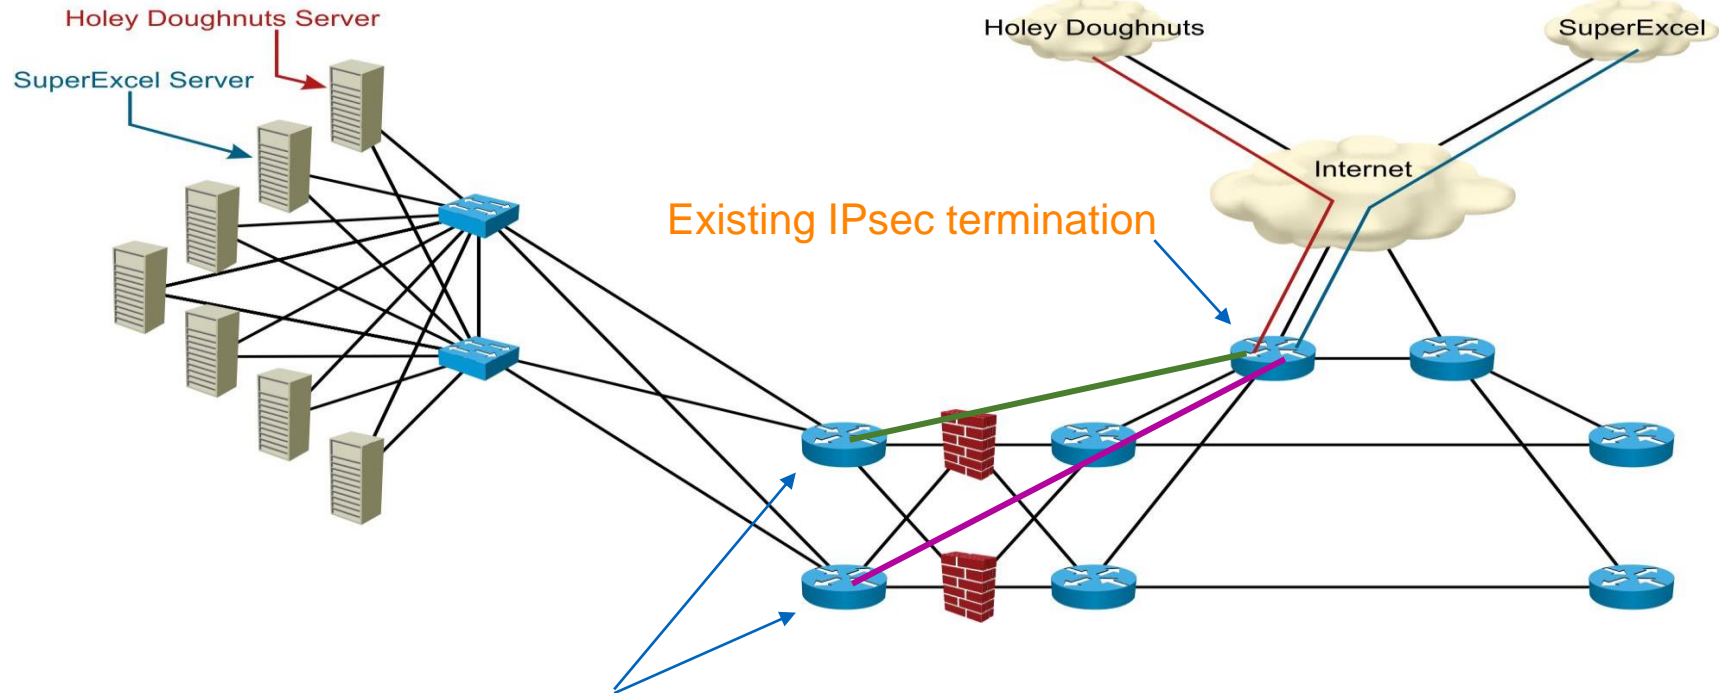
Item 4

Proposal 1



Item 4

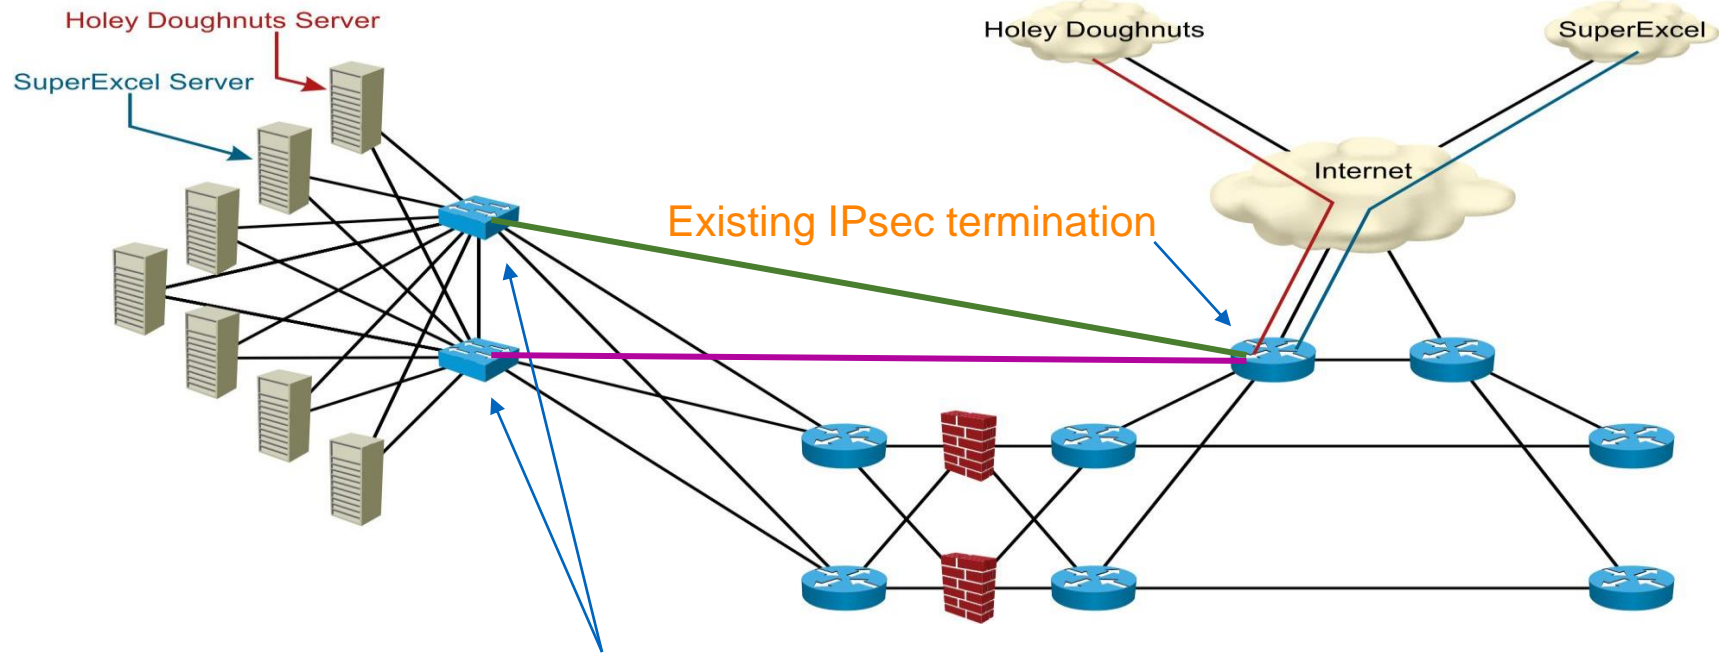
Proposal 2



Create an L3VPN between the IPsec termination and the default gateway routers

Item 4

Proposal 3



Create an L2TPv3 tunnel between the IPsec termination and the access switch





Item 4

- Note which requirement each proposal would fulfill below

	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment / management simplicity			

Item 4

- Note which requirement each proposal would fulfill below

	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment / management simplicity			
Supported by existing equipment			

It's best to treat this type of question as a set of multiple choice questions

Item 4

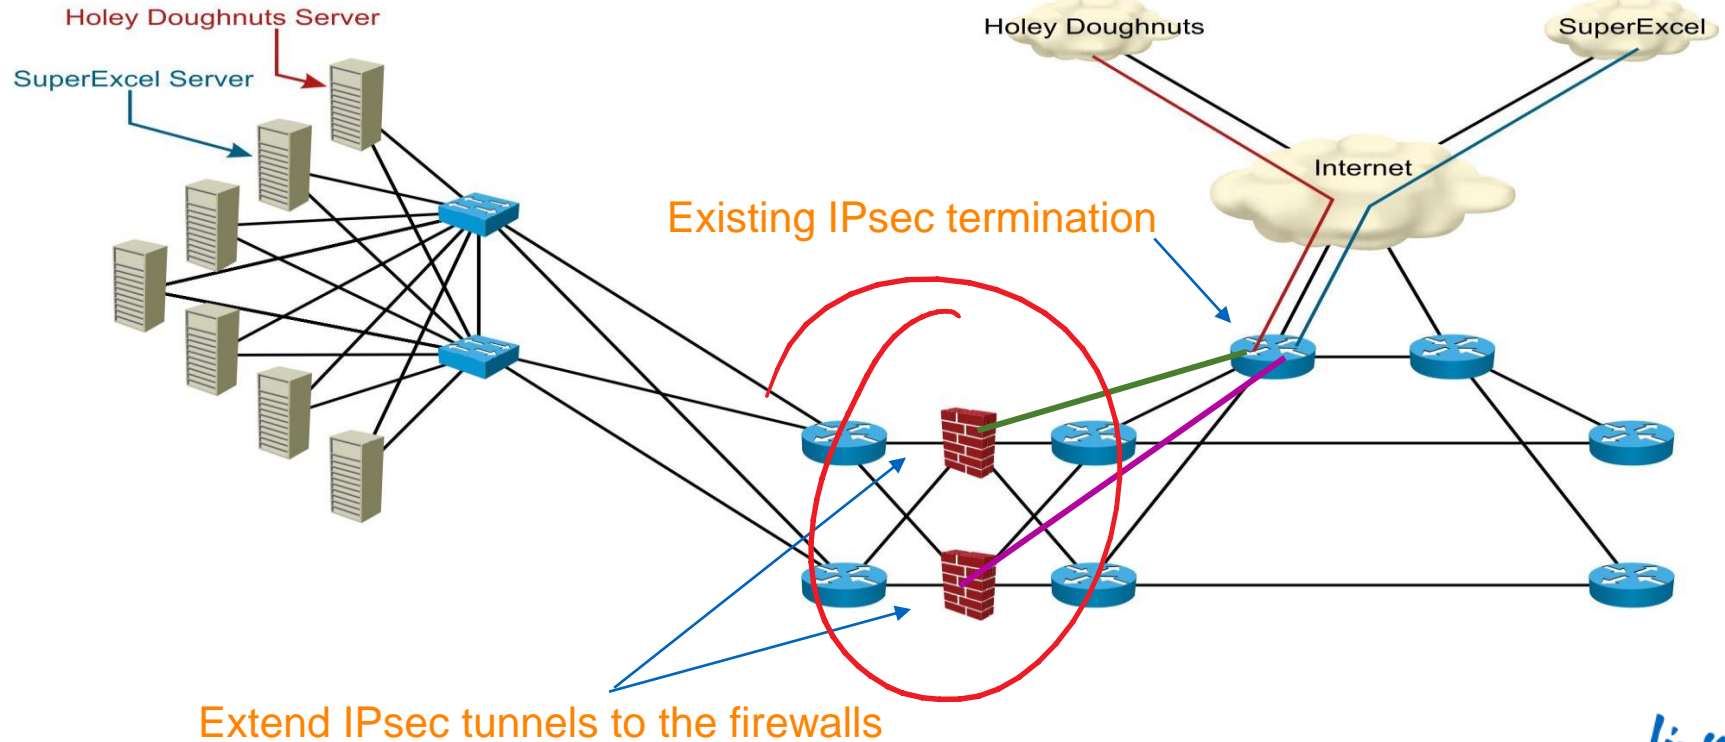
- Note which requirement each proposal would fulfill below

	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment / management simplicity			
Supported by existing equipment			

Would proposal 1 separate the traffic as required by the SLA?

Item 4


Proposal 1



Item 4

Proposal 1

- Note which requirement each proposal would fulfill below



	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment/ management simplicity			
Supported by existing equipment			

Since this could be argued either way, this would be graded correct with either answer

Item 4

Proposal 1

- Note which requirement each proposal would fulfill below




	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment/ management simplicity			
Supported by existing equipment			

IPsec doesn't have any connection to routing, but you could carry BGP over the IPsec tunnels to keep the routing tables separated

Item 4

Proposal 1

- Note which requirement each proposal would fulfill below





	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment/ management simplicity			
Supported by existing equipment			

This doesn't increase the number of tunnels being managed, so it's generally neutral on the complexity front

Item 4

Proposal 1

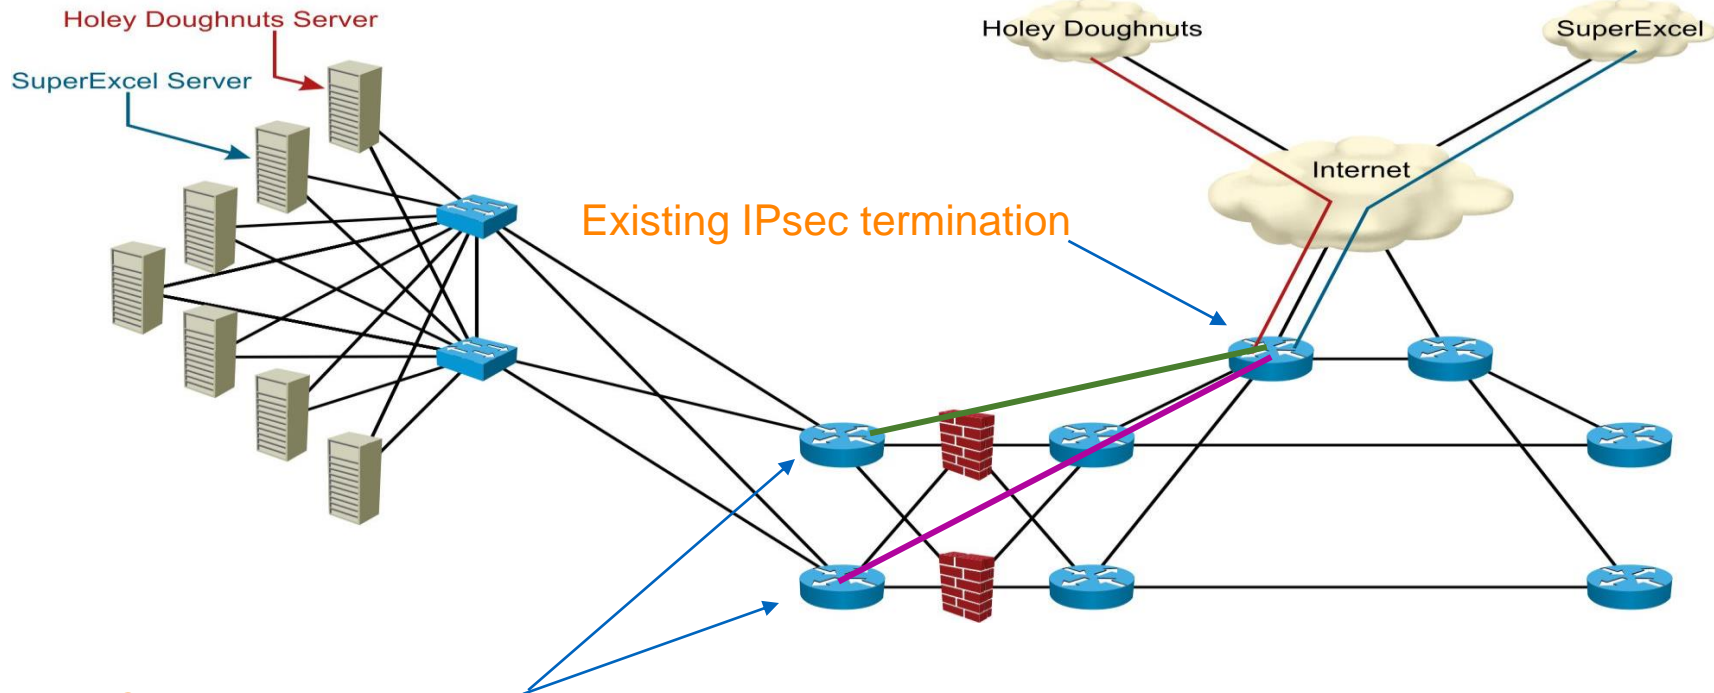
- Note which requirement each proposal would fulfill below

	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment/ management simplicity			
Supported by existing equipment			

The firewalls do support terminating IPsec tunnels

Item 4

Proposal 2











Create an L3VPN between the IPsec termination and the default gateway routers

Item 4

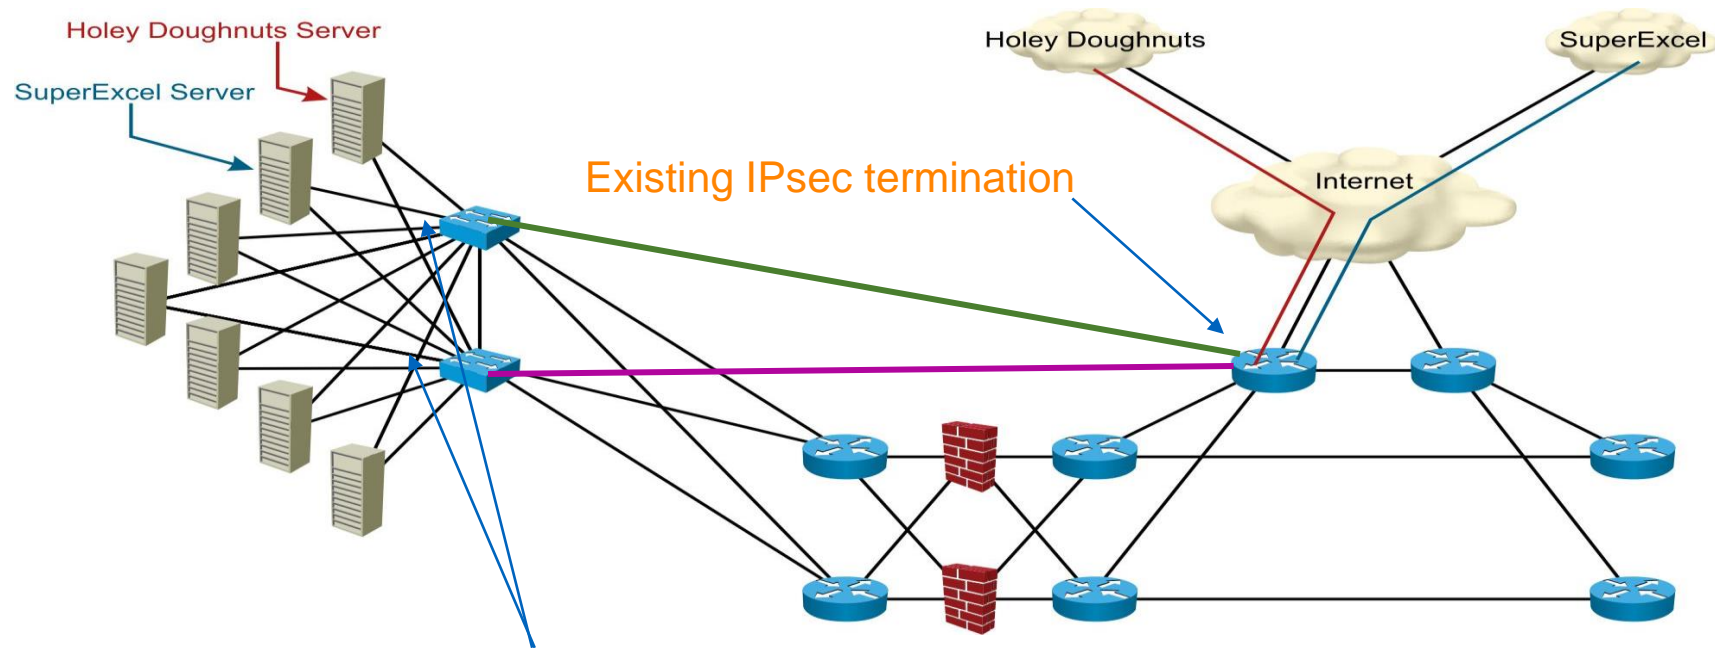
Proposal 2

- Note which requirement each proposal would fulfill below

	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment/ management simplicity			
Supported by existing equipment			

Item 4

Proposal 3















Create an L2TPv3 tunnel between the IPsec termination and the access switch

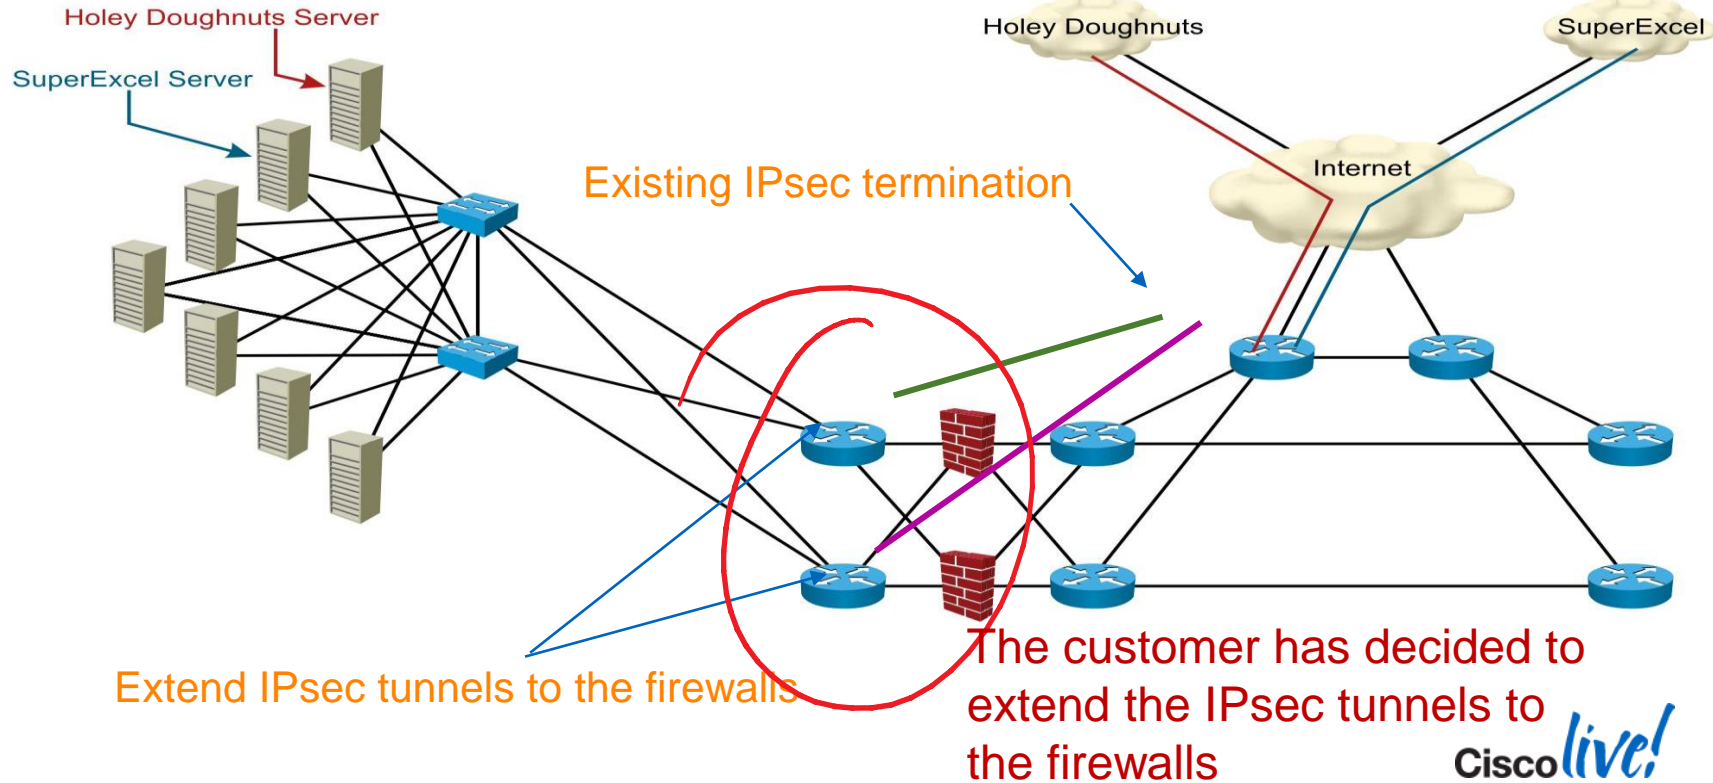
Item 4

Proposal 3

- Note which requirement each proposal would fulfill below

	Proposal 1 Extend IPsec	Proposal 2 L3VPN to default	Proposal 3 L2TPv3 to access
Traffic separation			
Route separation			
Deployment/ management simplicity			
Supported by existing equipment			

Information



Customer Decisions

- The customer in each situation might actually choose a way forward that you don't agree with
- You have to design about the problem you have in front of you, not the problem you wish you had
- You don't always make the right design decision, either. 😊

Item 5

Root

- What routing problem remains to be solved to meet the LISP SLA?
 - Separating the traffic from the firewall to the servers
 - Providing routing information to the virtual servers
 - Terminating the BGP sessions carried in the IPsec tunnels

This is “tricky” branching question; the root looks like just another multiple choice question...

Item 5

Root

- What routing problem remains to be solved to meet the LISP SLA?



- ✗ – Separating the traffic from the firewall to the servers
- Providing routing information to the virtual servers
- Terminating the BGP sessions carried in the IPsec tunnels

Traffic Separation is not a routing problem...

Item 5

Root

- What routing problem remains to be solved to meet the LISP SLA?




-  – Separating the traffic from the firewall to the servers
-  – Providing routing information to the virtual servers
- Terminating the BGP sessions carried in the IPsec tunnels

This makes a huge assumption that the servers, themselves, are running routing protocols

Item 5

Root

- What routing problem remains to be solved to meet the LISP SLA?



-  – Separating the traffic from the firewall to the servers
-  – Providing routing information to the virtual servers
-  – Terminating the BGP sessions carried in the IPsec tunnels

Where are we going to terminate those BGP sessions so we can share routing information with the customer?

Item 5

BGP Branch

- What piece of information is critical in determining where to terminate the BGP sessions flowing through the IPsec tunnel?

-  – The number of routes each customer is sending to LISP
-  – Whether or not the firewalls support BGP peering

Information

From: Bruce
Sent: Monday 12:12 PM
To: Amy
Cc: [lsip-eng](#)
Subject: Possible Solution

The firewalls definitely cannot terminate BGP peering sessions –they just aren't designed for routing functionality except at a very minimal level. The two routers between the firewalls and the access switches are capable of running multiple "virtual routers" on a single same platform, and tying .1q tunnels at layer 2 into each virtual machine.

HTH

Bruce

Bruce answers our question, but is there anything else here we care about?

Information

From: Bruce
Sent: Monday 12:12 PM
To: Amy
Cc: [lsip-eng](#)
Subject: Possible Solution

The firewalls definitely cannot terminate BGP peering sessions –they just aren't designed for routing functionality except at a very minimal level. The two routers between the firewalls and the access switches are capable of running multiple "virtual routers" on a single same platform, and tying .1q tunnels at layer 2 into each virtual machine.

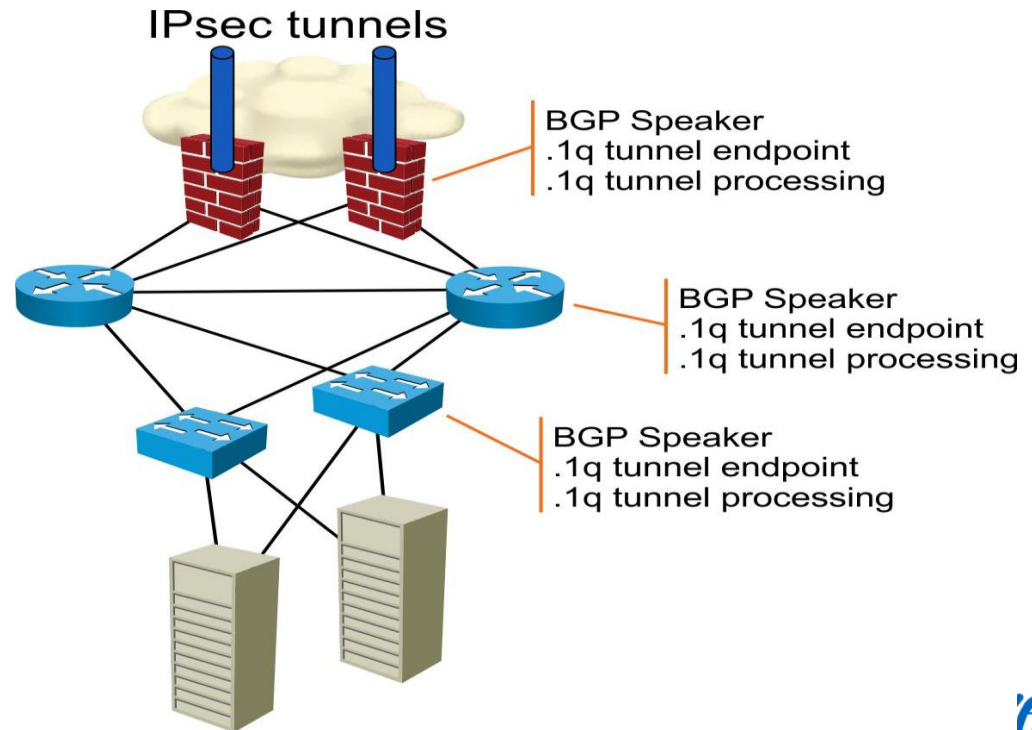
HTH

Bruce

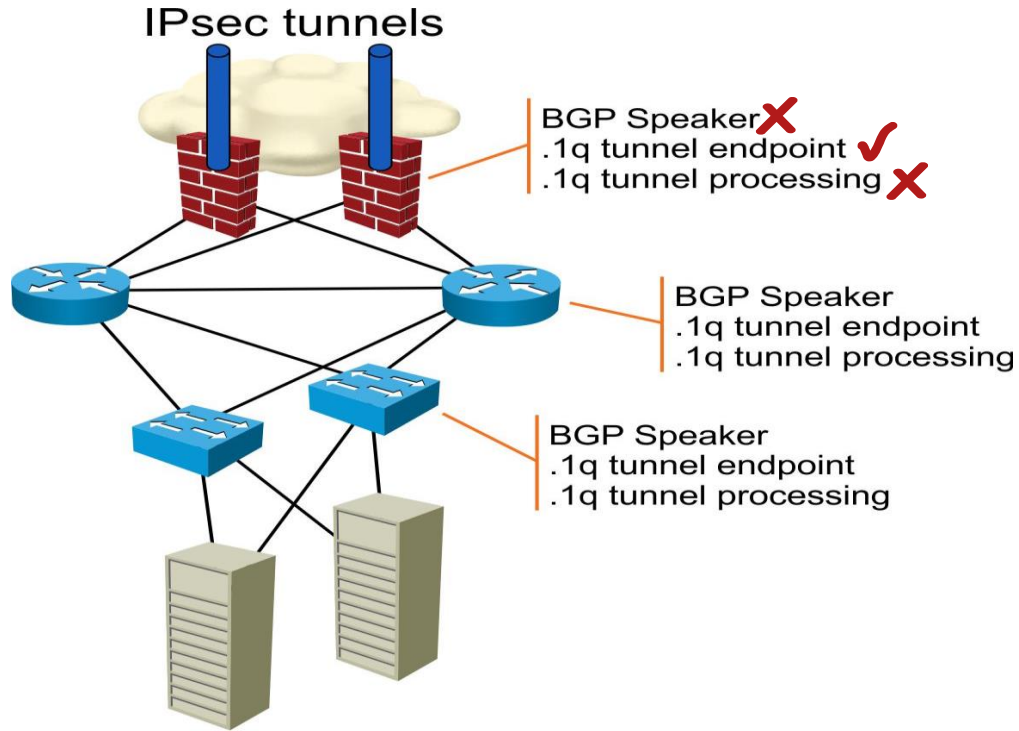
Sometimes information is buried in an answer that becomes very important later on...

Item 6

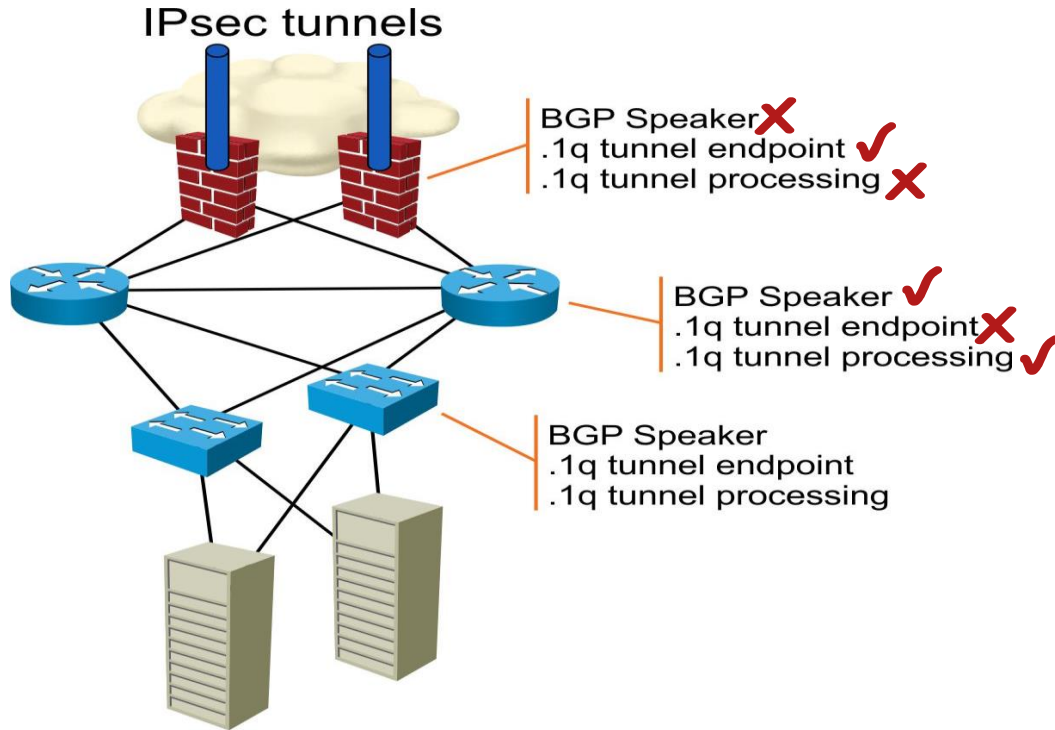
- Mark each item you would configure in order to complete the design between the firewalls and the servers



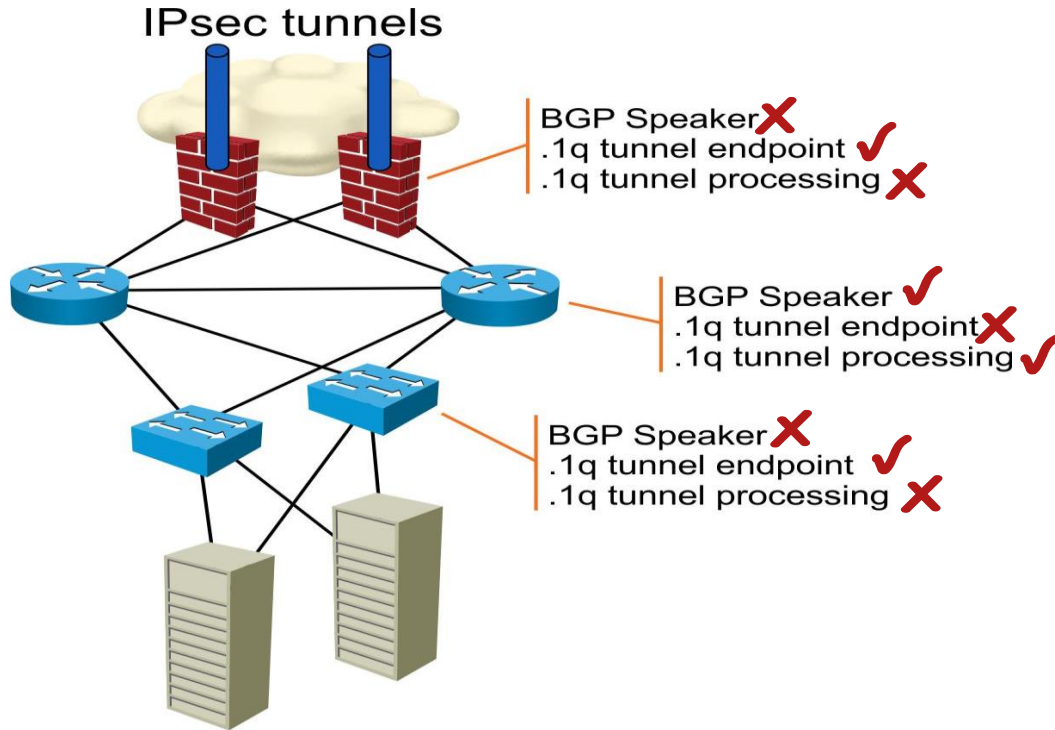
Item 6



Item 6



Item 6



Second Problem Information

From: Amy
Sent: Thursday 11:21 AM
To: lisp-eng
Cc:
Subject: New Video Application

The folks in the call center are rolling out a new application, and wanted to know what we would need to do to the network to support this new application. It's a video application they're going to use for training new folks who come into the call center, to bring them up to speed on specific products we are providing support for.

What do y'all think?

Amy

Item 7



- What application information do you need to determine the impact on the network of this new video application?
 - What the average length of video session will be
 - Whether the video is interactive or streaming
 - What CODEC is being used to encode the video onto the network
 - What desktop application will be used for this video stream

Item 7

- What application information do you need to determine the impact on the network of this new video application?
 - ✗ – What the average length of video session will be
 - Whether the video is interactive or streaming
 - What CODEC is being used to encode the video onto the network
 - What desktop application will be used for this video stream




Item 7

- What application information do you need to determine the impact on the network of this new video application?

-  – What the average length of video session will be
- Whether the video is interactive or streaming
- What CODEC is being used to encode the video onto the network
-  – What desktop application will be used for this video stream





Item 7

- What application information do you need to determine the impact on the network of this new video application?

-  – What the average length of video session will be
- Whether the video is interactive or streaming
-  – What CODEC is being used to encode the video onto the network
-  – What desktop application will be used for this video stream

Item 7

- What application information do you need to determine the impact on the network of this new video application?

-  – What the average length of video session will be
-  – Whether the video is interactive or streaming
-  – What CODEC is being used to encode the video onto the network
-  – What desktop application will be used for this video stream

Information

From: Bruce
Sent: Thursday 2:12 PM
To: lisp-eng
Cc:
Subject: RE: New Video Application

I called over to the vendor they're buying the application from, and also talked to the business managers over there who are championing this project. The application is a streaming on demand system used to provide training information to new agents, or agents who are picking up support for new products.

The two servers are sized to feed five or six streams at a time. If a new stream is attempted beyond what the server will support, the application returns an error. To get around this limitation, they're planning on implementing a scheduling system, so they group folks watching the same video as much as possible. They are expecting, based on their pilot runs, for ten to twenty people to be watching a given video at any time during the day.

HTH

Bruce

Item 8

- What further information would you need to determine what network modifications will be needed to support this new application (choose two)?
 - The utilisation of the DC to call centre links
 - The average delay between the DC and the call centre
 - The quality of service configurations on DC to call centre links
 - The average jitter level between the DC and the call centre

Item 8

- What further information would you need to determine what network modifications will be needed to support this new application (choose two)?

- The utilisation of the DC to call centre links
- The average delay between the DC and the call centre
- The quality of service configurations on DC to call centre links
- The average jitter level between the DC and the call centre

We can knock out both of these because this is not a streaming application

Item 8

- What further information would you need to determine what network modifications will be needed to support this new application (choose two)?

- ✓ – The utilisation of the DC to call centre links
- ✗ – The average delay between the DC and the call centre
 - The quality of service configurations on DC to call centre links
- ✗ – The average jitter level between the DC and the call centre

Item 8

■ What further information would you need to determine what network modifications will be needed to support this new application (choose two)?

- ✓ – The utilisation of the DC to call centre links
- ✗ – The average delay between the DC and the call centre
- ✓ – The quality of service configurations on DC to call centre links
- ✗ – The average jitter level between the DC and the call centre

Item 9

Root

- Would you recommend deploying multicast to support this new application?
 - Yes
 - No

Item 9

“No” Branch

- Why don't you believe LISP should deploy multicast to support this application?

X – The video is on demand, which is best handled point to point

X – The servers are easily able to handle the load without the complexity of deploying multicast

X – The servers and watchers are close together, so multicast isn't required

Item 9

“Yes” Branch

- What multicast routing protocol would you recommend for this network?

X – Multicast OSPF (RFC 1584)

– PIM Sparse (RFC 2362)

X – PIM Dense (RFC 3973)

– PIM BiDir (RFC 5105)

We can knock out both of these easily

Item 9

“Yes” Branch

- What multicast routing protocol would you recommend for this network?

X – Multicast OSPF (RFC 1584)

– PIM Sparse (RFC 2362)

X – PIM Dense (RFC 3973)

X – PIM BiDir (RFC 5105)

We don't have requirements for BiDir

Item 9

“Yes” Branch

- What multicast routing protocol would you recommend for this network?

 – Multicast OSPF (RFC 1584)

 – PIM Sparse (RFC 2362)

 – PIM Dense (RFC 3973)

 – PIM BiDir (RFC 5105)



Conclusion / Review

Review

- The CCDE Practical is scenario-driven: it uses a hypothetical story as a basis to thinking through a complex design test that is credible, real-life, and cover multiple steps
- You begin with a network and a challenge
- Using this information, you...
 - Decide what information you'll need to finish the challenge
 - Determine what the critical technical and business constraints are
 - Determine which technology fits where
 - Determine how to deploy specific technologies to meet the challenge
 - Respond to customer ideas and queries

Review

- Branching Questions
 - They don't take a specific form
 - You won't even know when you're taking the test if you've encountered one
 - Just do the best you can on each and every question, and don't worry about the underlying structure
- Chart and Drag 'n Drop Questions
 - Don't freak out
 - Just treat each row/column combination as if it were an individual yes/no question
- Diagram Questions
 - Just break the problem down into a set of individual, more manageable problems

Review

- If you don't know the answer...
 - There's probably some piece of information missing someplace
 - Think through all the documentation you have, and see if you can figure out what's missing
- Ask the “key questions” about each and every question you encounter

Final Thoughts

- Get hands-on experience with the Walk-in Labs located near the Registration Desk
- Come see demos of many key solutions and products in the main Cisco booth in the World of Solutions
- Visit www.ciscoLive365.com after the event for updated PDFs, on-demand session videos, networking, and more!
- Follow Cisco Live! using social media:
 - Facebook: <https://facebook.com/CiscoLiveMEL>
 - Twitter: <https://twitter.com/#clmel>
 - Linked In Group: <http://www.linkedin.com/groups/Cisco-Live-Australia-4744933?gid=4744933>
 - Flickr: <http://flickr.com/photos/ciscolivemel>



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™



Reference Material



CCIE® and CCDE® Program Overview

CCIEs and CCDEs Worldwide

- CCIE: Most highly respected IT certification for more than 15 years
- CCDE: Launched in 2008
- Industry standard for validating expert skills and experience
- Demonstrate strong commitment and investment to networking career, life-long learning, and dedication to remaining an active CCIE / CCDE



Overview: CCIE and CCDE Tracks

Routing & Switching

- Expert level knowledge of networking across LAN and WAN interfaces and variety of routers and switches
- Solve complex connectivity problems to increase bandwidth, improve response times, maximise performance, and support global operations

Security

- Expert level knowledge of security and VPN solutions
- Demonstrate in-depth understanding of Layer 2 and 3 network infrastructure; Solid understanding of Windows, Unix, Linux and HTTP, SMTP, FTP and DNS

Voice

- Expert level knowledge of Cisco Voice over IP (VoIP) products and solutions
- Capable of building and configuring complex end-to-end telephony network, troubleshooting and resolving VoIP-related problems

Design

- Expert level knowledge of network design principles for the Layer 2 and 3 network infrastructure
- Capable of assessing and translating network business requirements into technical designs

Data Centre

- Expert level knowledge of Data Centre Technologies, including DC infrastructure, storage, compute and virtualisation
- Capable of building, configuring, and troubleshooting an end-to-end virtualised Data Centre using Cisco DC technologies

Service Provider

- Expert level knowledge of IP fundamentals and technologies
Expertise in building an extensible service provider network
- Expert level knowledge to troubleshoot and maintain complex service provider networks

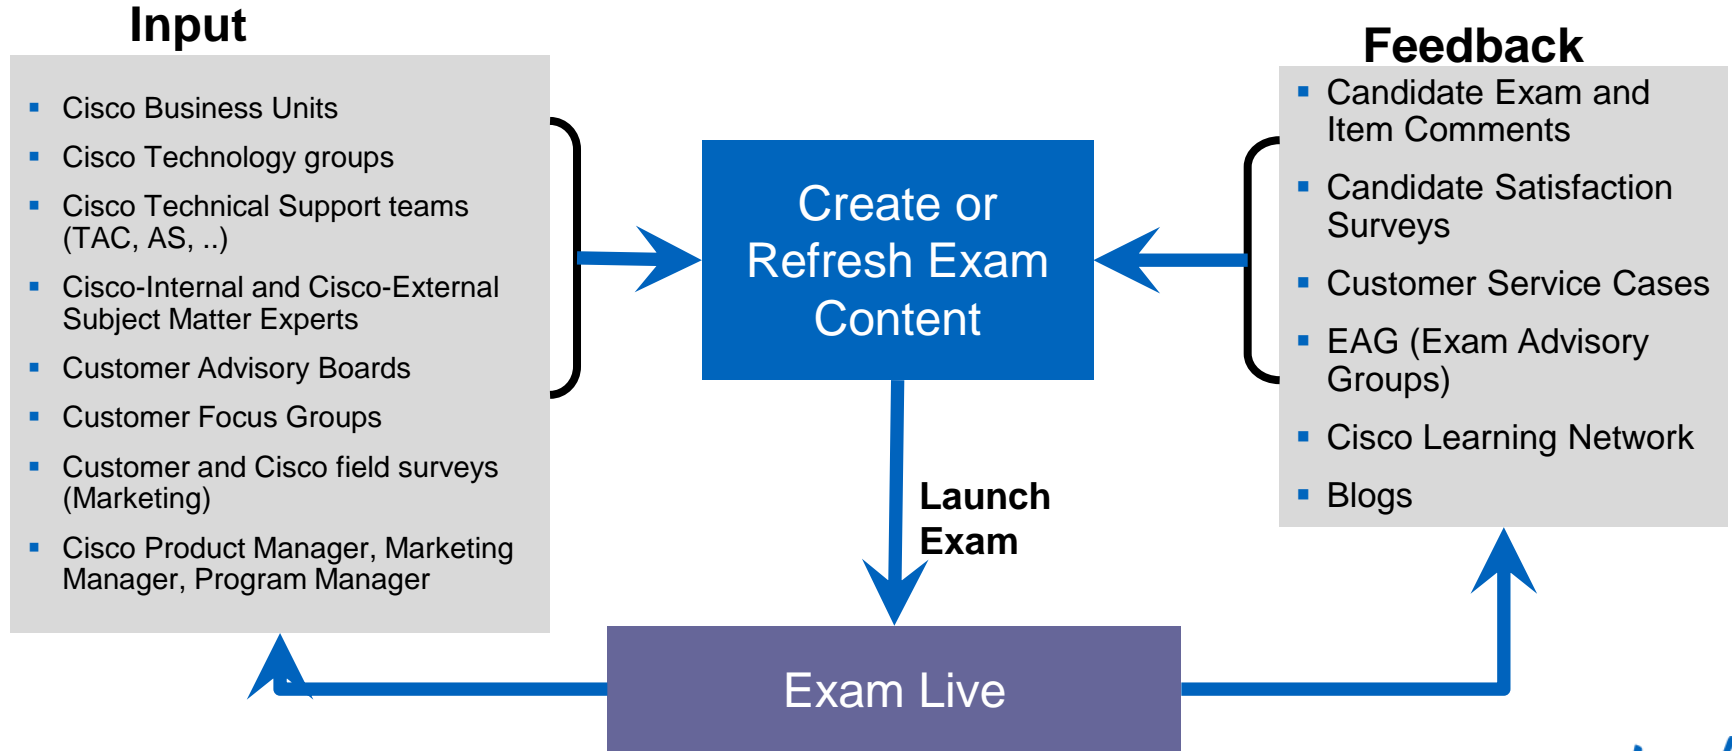
SP Operations

- Expert level knowledge of SP IP NGN technologies
- Capable of troubleshooting SP networks, managing SP processes (incident, fault, change, configuration, and performance), and knowledge of NMS technology

Wireless

- Expert level knowledge of WLAN technologies
- Provides next step for individuals interested in a career in managing or working with Cisco wireless technologies

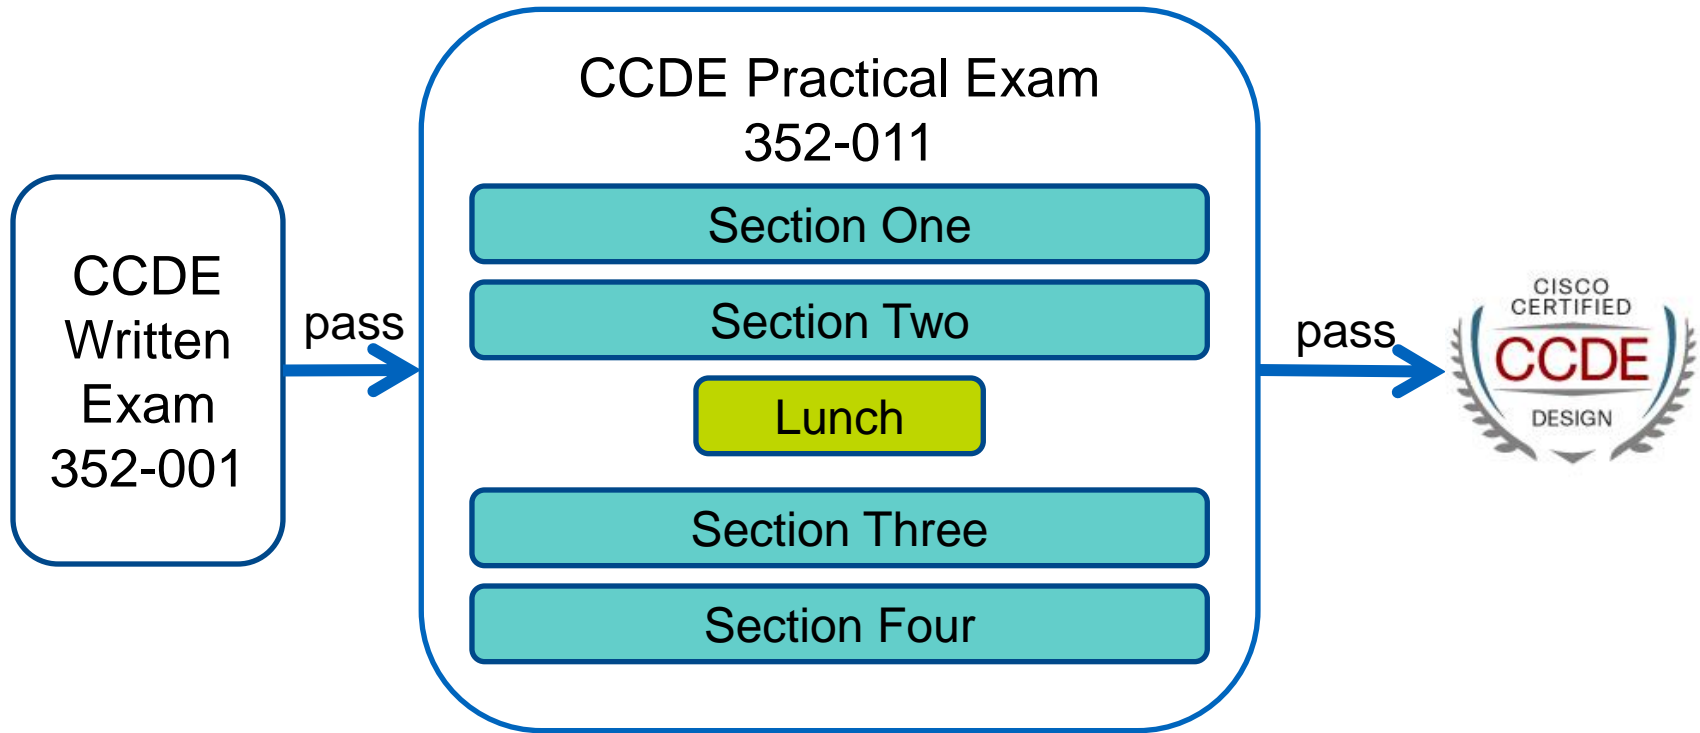
Proactive and Holistic Candidate Feedback Process





Taking the CCDE

Certification Process



Certification Process

- No pre-requisites
- 2-hour written exam
 - Design Theory
 - Protocol Operation
 - Technology Operation
- 8-hour scenario-based practical exam
 - Analyse design requirements
 - Develop network designs
 - Implement network design
 - Validate and optimise network design



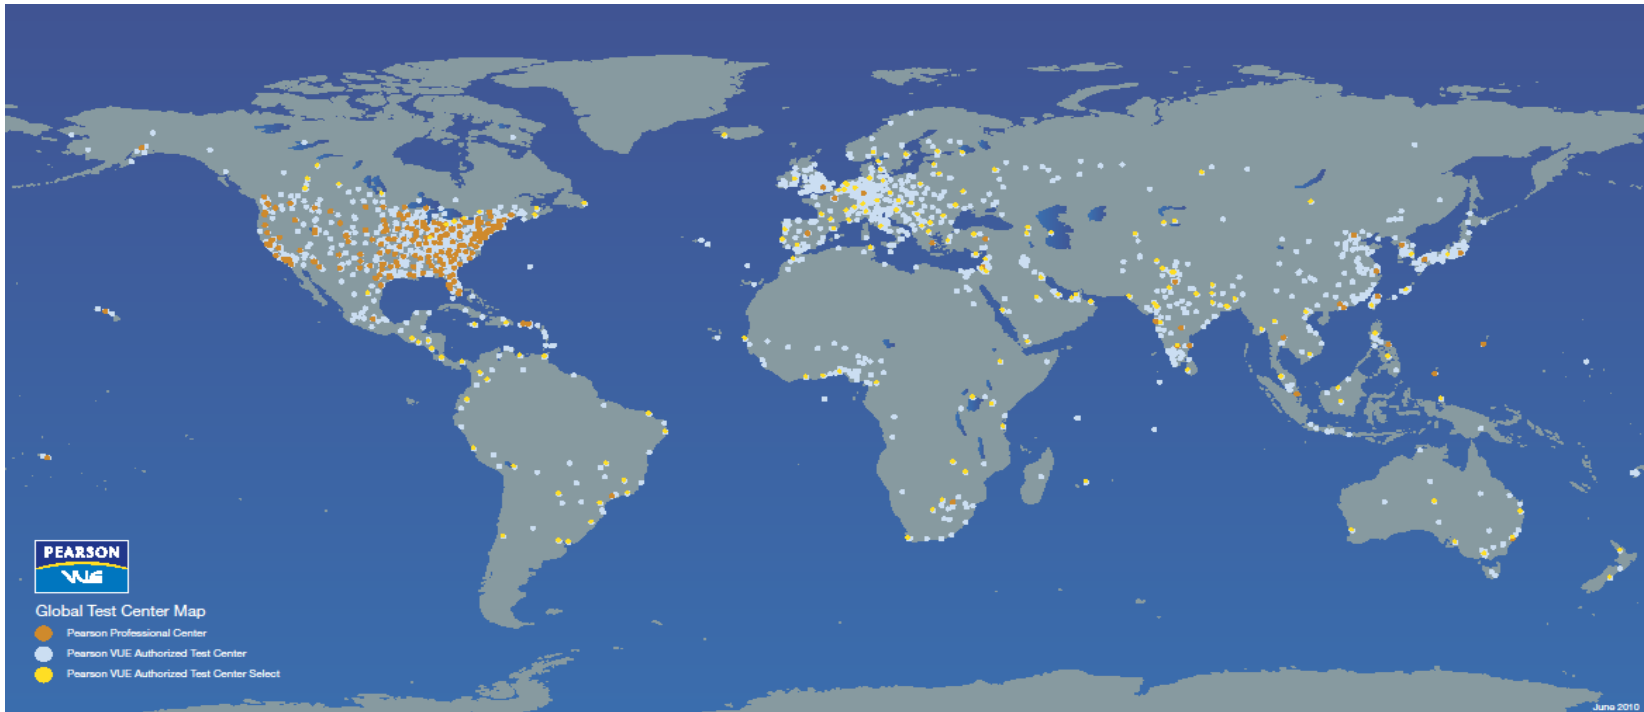
CCDE Written

352-001 Exam Information

- **Location:** Pearson VUE Test Centres
- **Cost:** \$350 USD
- **Duration:** 2 hours
- **Format:** 85-110 items
 - Multiple Choice
 - Drag & Drop
 - Scored & non-scored items
 - NO “skip question”, NO “go back”
- **Pre-requisite:** None, though following is recommended:
 - minimum of seven (7) years job experience in network engineering
 - thorough understanding of networking infrastructure principles
 - in-depth understanding of the topics in the exam topics on CLN

CCDE Written

Locations (Pearson VUE Testing Centres)



CCDE Written

Registration

- Appointments made in advance (phone/web) on VUE testing centres or on day of test (at certain testing centres)
- Personal information required to register for exam:
 - Legal Name
 - Cisco Certification ID or Test ID number (if you have taken Cisco exams before, you will use the same ID as before)
 - Company name
 - Valid email address
 - Payment (payment required at time of registration)

CCDE Written

Policies

- No limit to number of attempts that can be made
- If fail:
 - Candidate must wait 5 calendar days between exam attempts before retesting (beginning the day after the failed attempt)
- If pass:
 - Candidate must wait minimum of 180 days (6 months) before retesting

CCDE Written

Scoring

- Auto-scored, based on statistical analysis
 - Must achieve minimum score to pass
 - Pass/fail results available immediately
- Candidates receive score report, which indicates broad areas where additional study may be useful; score report mapped to exam topics

CCDE Practical

352-011 Exam Information

- **Location:** Pearson VUE Professional Test Centres; event based (only offered on specific dates/locations)
- **Cost:** \$1500 USD
- **Duration:** 9 hours (8 hour exam plus 1 hour lunch break)
- **Format:**
 - 4 scenario-based sections
 - Approximately 28-38 items per scenario
 - Scored & non-scored items
 - NO “skip question”, NO “go back”
 - 24 inch monitors
 - NEW: Auto-scored
- **Pre-requisite:** CCDE Written

CCDE Practical

Locations (Pearson Professional Centres – approximately 300 worldwide)



CCDE Practical

2014 Upcoming Test Dates

Exam Date	Date Registration Closes	Pearson Professional Centre Location
Thursday, May 15, 2014	Wednesday, May 14, 2014	Worldwide locations Schedule your exam at location nearest to you
Thursday, August 21, 2014	Wednesday, August 20, 2014	Worldwide locations. Schedule your exam at location nearest to you
Thursday, November 20, 2014	Wednesday, November 19, 2014	Worldwide locations Schedule your exam at location nearest you

CCDE Practical

Registration

- Appointments made in advance (phone/web) to the VUE PPCs
- Personal information required to register for exam:
 - Legal Name
 - Cisco Certification ID or Test ID number (if you have taken a Cisco exam before, you will use the same ID as before)
 - Company name
 - Valid email address
 - Payment (payment required at time of registration)
 - Confirmation of success in passing CCDE Written exam and acknowledgement that exam results are still valid

CCDE Practical

Policies

- To take exam, candidate must:
 - Pass CCDE Written
 - CCDE Written exam attempt must be valid
- How to Keep CCDE Written Valid
 - Initial attempt of CCDE Practical must be within 18 months of passing CCDE written
 - Candidates who do not pass CCDE Practical must re-attempt CCDE Practical within 12 months of last scored attempt
 - Candidates who do not pass CCDE Practical within three years of passing CCDE Written must retake the CCDE Written before being allowed to re-attempt CCDE Practical

CCDE Practical

Retake Policies

- No limit to number of attempts that can be made
- If fail:
 - Candidate must wait until the next testing window to retest
- If pass:
 - Candidate must wait minimum of 180 days (6 months) before retesting

CCDE Practical

Recertification

- To maintain active CCDE status, CCDEs are required to pass any expert-level certification exam every 24 months
 - Recertification deadlines based on original certification date
 - If CCDE recertification requirements are not completed on or before expiration date, CCDE certification is suspended
 - Suspended candidates have one year to recertify by passing CCDE Written. If candidate does not recertify in time, all requirements must be completed again

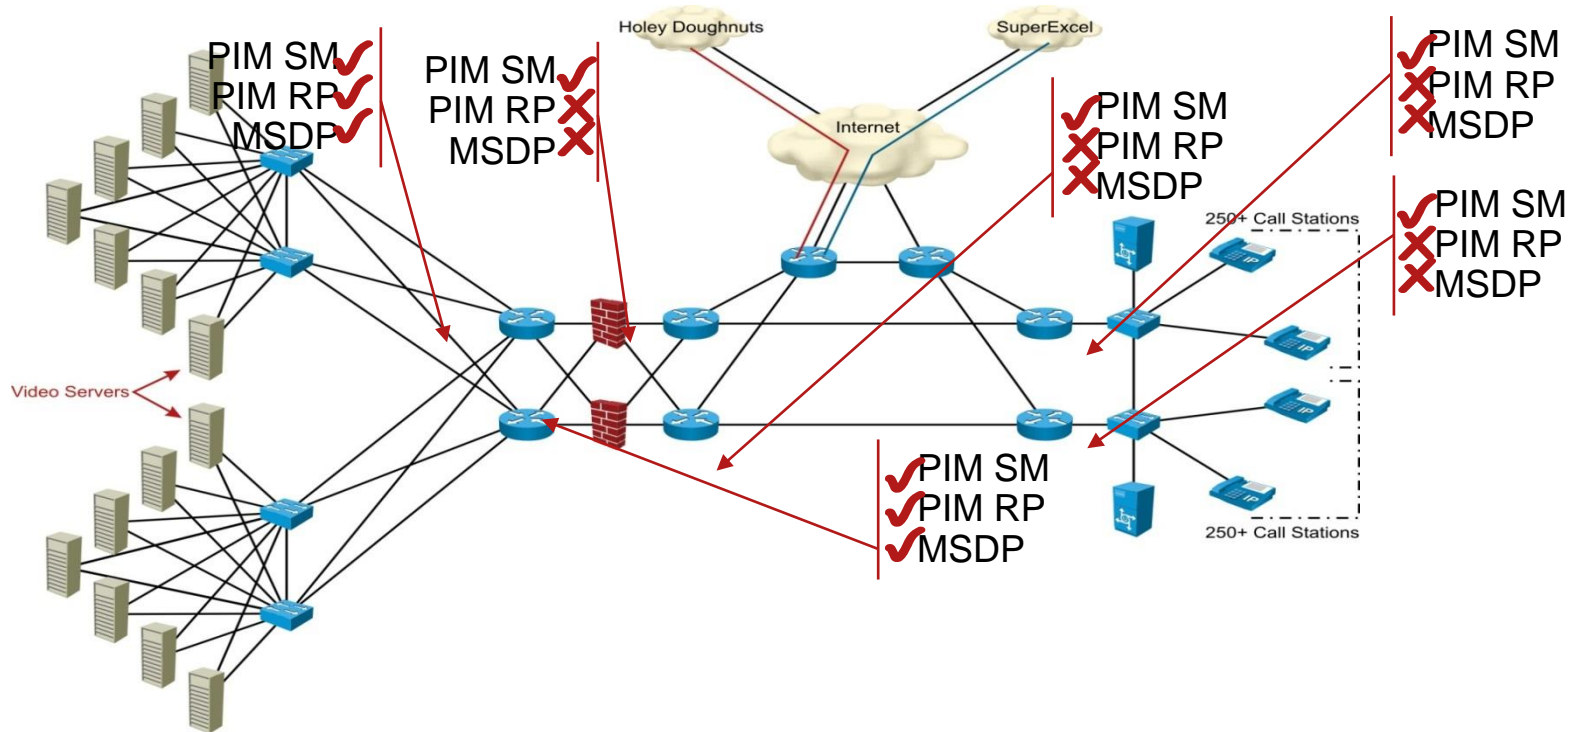
CCDE Practical

Scoring

- Exam is now auto-scored!
 - Employs statistical analysis to assure all exams are in the same length and difficulty level
 - Multiple valid solutions are taken into consideration for scoring
 - Some items are worth multiple points, with partial scoring
 - Must achieve minimum score to pass
- Candidates receive score report, which indicates broad areas where additional study may be useful; score report mapped to exam topics

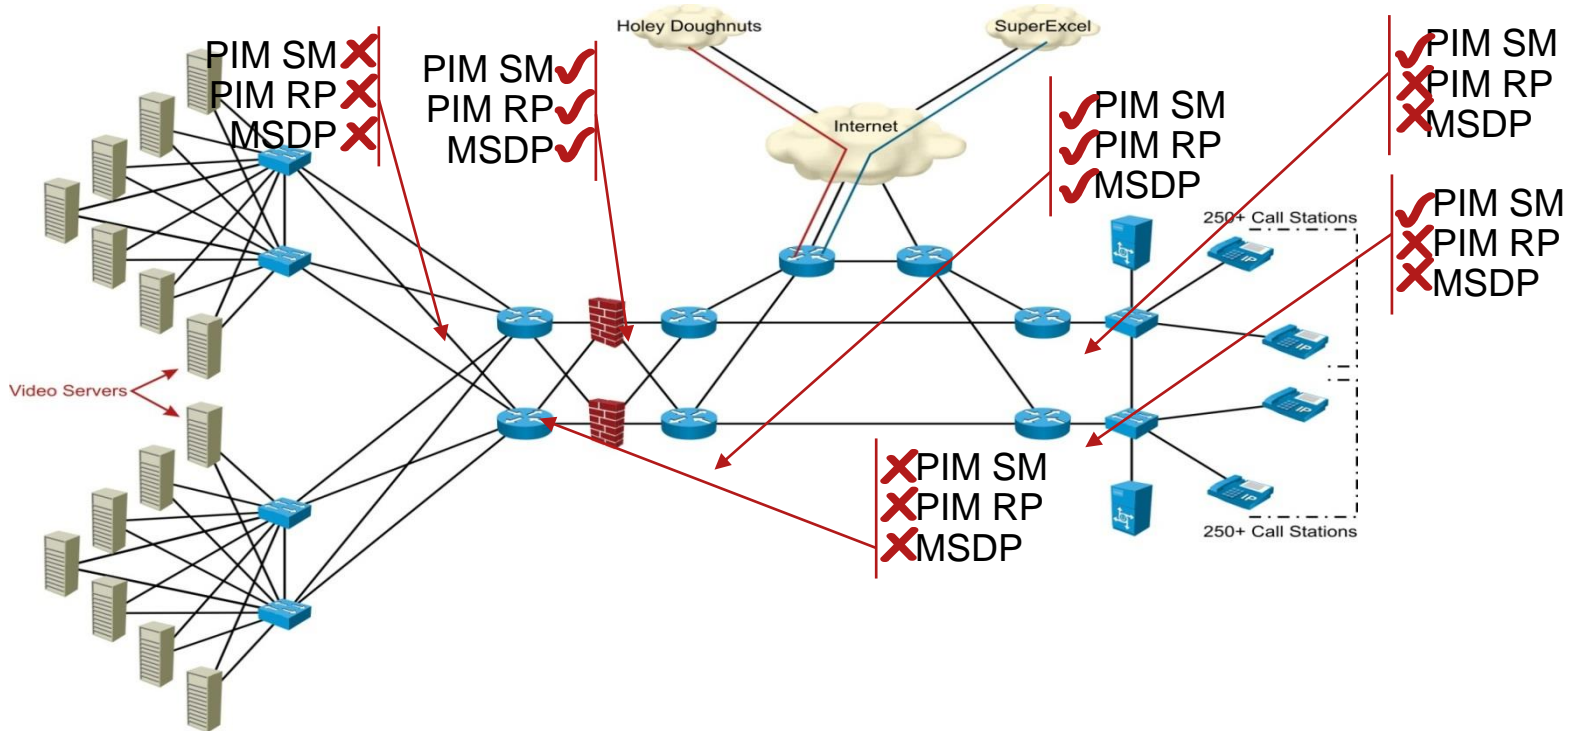
Partial Scoring Example

Correct



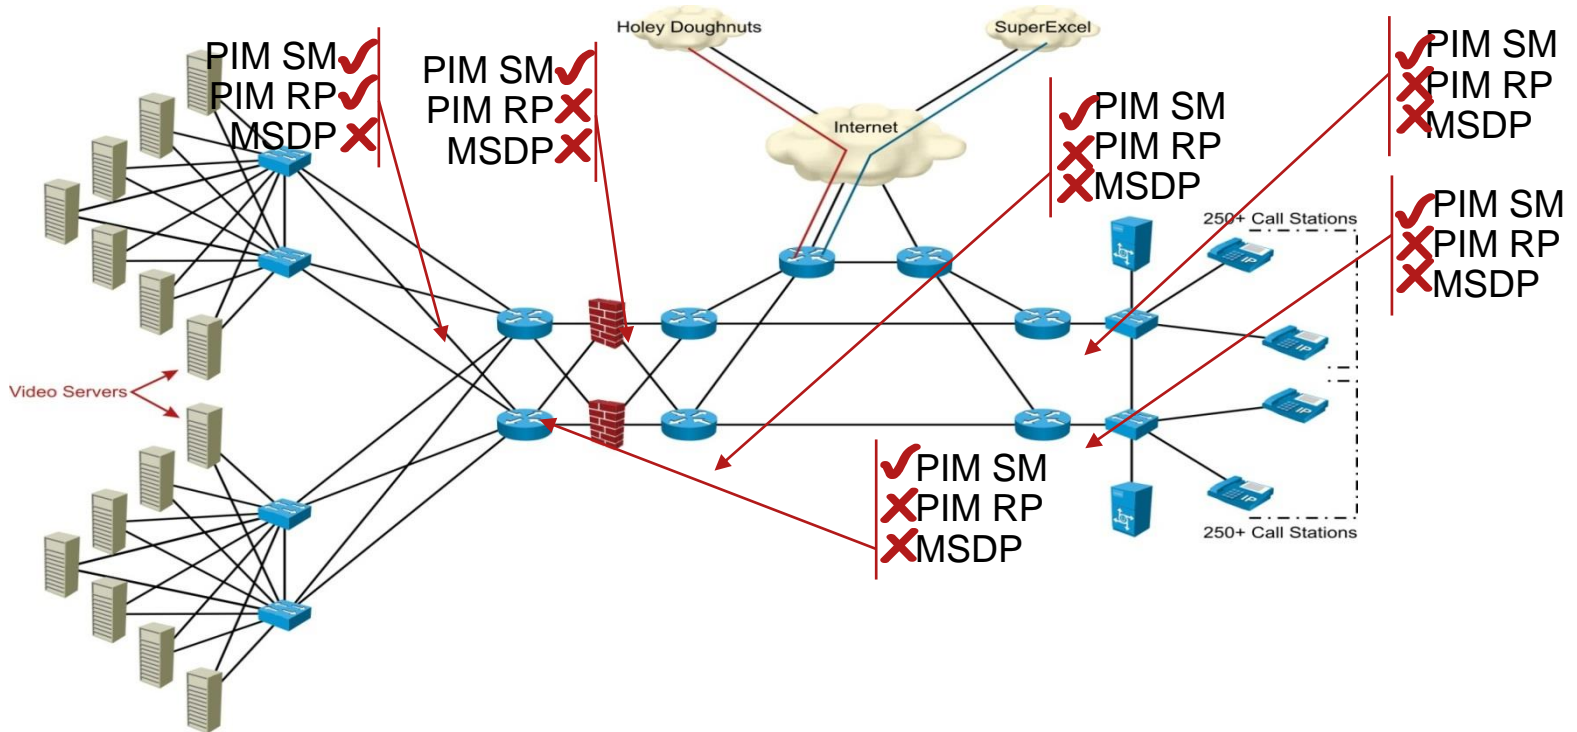
Partial Scoring Example

Partially Correct



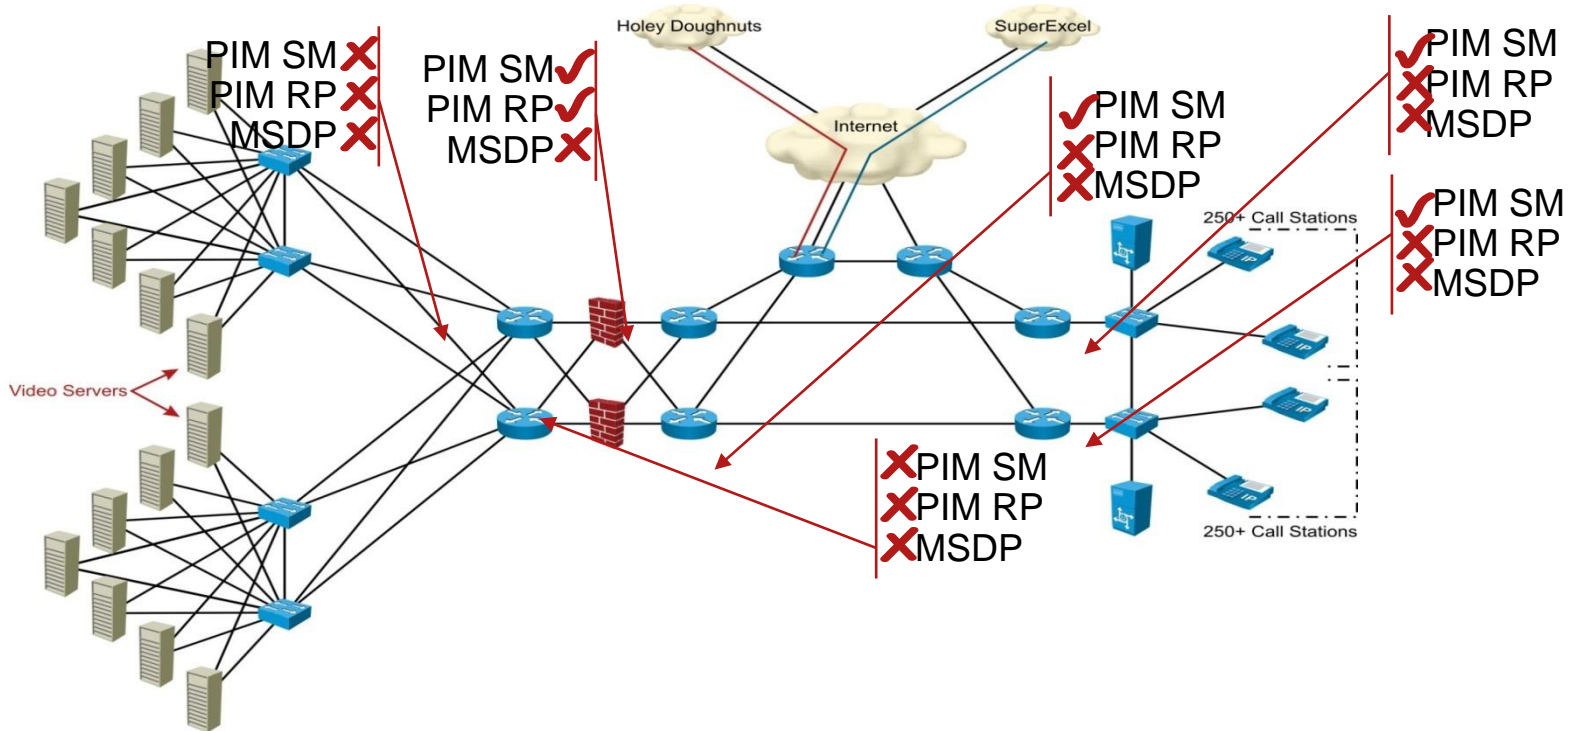
Partial Scoring Example

Partially Correct



Partial Scoring Example

Partially Correct



Partial Scoring Example

Correct

Note which requirements each technology listed below would fulfill

	LISP Traffic Security Requirements	LISP Routing Security Requirements	SLA Requirements
L3VPN (MPLS)	✓	✓	✓
L2TPv3	✗	✗	✓
VPLS (MPLS)	✓	✗	✓
GRE Tunnels	✗	✗	✓

Partial Scoring Example

Partially Correct

Note which requirements each technology listed below would fulfill

	LISP Traffic Security Requirements	LISP Routing Security Requirements	SLA Requirements
L3VPN (MPLS)	✓	✓	✓
L2TPv3	✗	✗	✓
VPLS (MPLS)	✓	✗	✓
GRE Tunnels	✗	✗	✓



Preparation Resources

Preparation Resources

- Reference exam topics and other tools and resources posted on CLN to assess your experience and knowledge and determine if you are ready to take the exam
- Utilise other available resources for preparation, like the Cisco Press CCDE Quick Reference Guide
- Start by preparing for and taking the CCDE Written exam; Do not try to prepare for CCDE Written and CCDE Practical at same time
- Use CCDA and/or CCDP as a starting point (to confirm your body of knowledge required for CCDE), but do not limit your study to CCDA/CCDP material

Preparation Resources

- Cisco Learning Network (CLN)
 - CCDE homepage: <https://learningnetwork.cisco.com/community/certifications/ccde>
 - CCDE Written: https://learningnetwork.cisco.com/community/certifications/ccde/written_exam
 - CCDE Practical: https://learningnetwork.cisco.com/community/certifications/ccde/practical_exam
 - CCDE Reading List: <https://learningnetwork.cisco.com/docs/DOC-2462>
- CCDE Study Group: <https://learningnetwork.cisco.com/groups/ccde-study-group>
- Cisco Product Pages (former Cisco Documentation CD) <http://www.cisco.com/web/psa/products/index.html>
- CCDE Customer support and FAQ: <http://www.cisco.com/go/certsupport>
- Cisco Press: <http://www.ciscopress.com>

Crossing the CCDE Border Webinar

Webinar Objectives:

- Provide audience guidance & mentorship in achieving their professional goals as CCDEs
- Share the information needed for the audience to:
 - Cross CCDE border and evolve their careers
 - Prepare for the CCDE certification

Outline:

- Starting up the CCDE Journey
- Characteristics of Successful CCDE Candidates
CCDE Preparation Tools, Tips & Resources
- CCDE Leaders on stage
- Open Q&A Session

Recording: <https://learningnetwork.cisco.com/docs/DOC-16003>

Panel



Ahmed Abro



Bruce Pinsky



Aaron Aday



Elaine Lopes

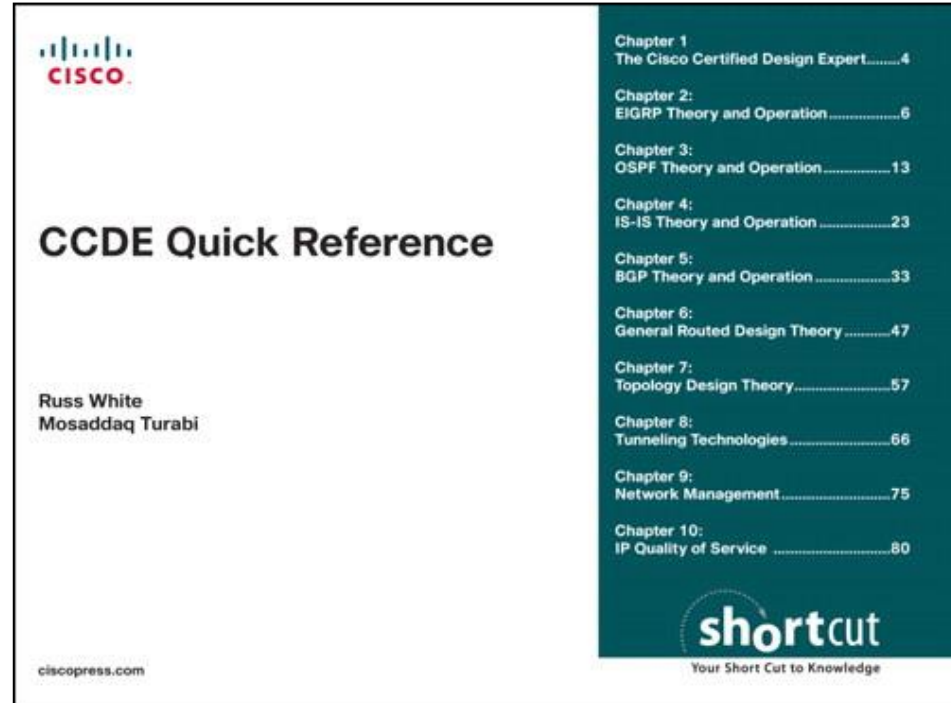


Errol Hayward



Cisco *live!*

Preparation Resources – Written Exam

- How do I study for this?
- Start here
 - This is essentially the written exam topics “explained”
 - The technologies covered on both the written and the practical are the same
 - The written exam topics is focused on the technology, the practical on analysis



Preparation Reading – Practical Exam





Optimal Routing Design

Techniques for optimizing large-scale IP routing operation and managing network growth

Russ White, CCIE® No. 2653
Don Slice, CCIE No. 1929
Alvaro Retana, CCIE No. 1609

ciscopress.com





MPLS and VPN Architectures

CCIP™ Edition

Prepare for CCIP certification as you learn to design and deploy MPLS-based VPNs

Ivan Pepelnjak, CCIE® No. 1354
Jim Guichard, CCIE No. 2069

ciscopress.com




Comparing, Designing, and Deploying VPNs

A practical guide for comparing, designing, and deploying IPsec, MPLS Layer 3, L2TPv3, L2TPv2, AToM, and SSL virtual private networks

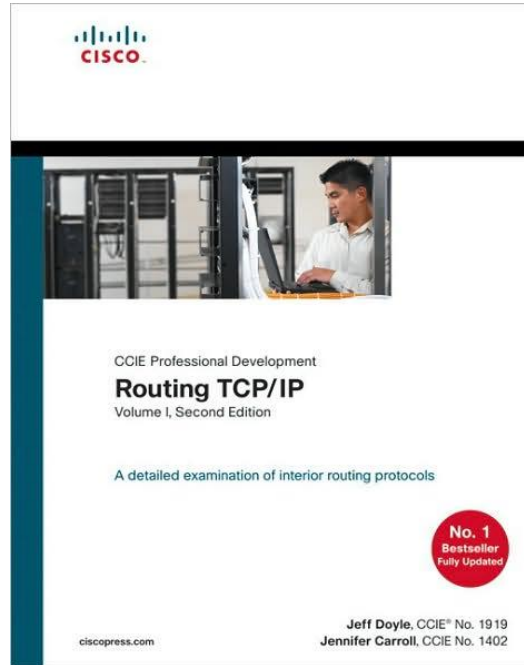
Mark Lewis, CCIE® No. 6280

ciscopress.com

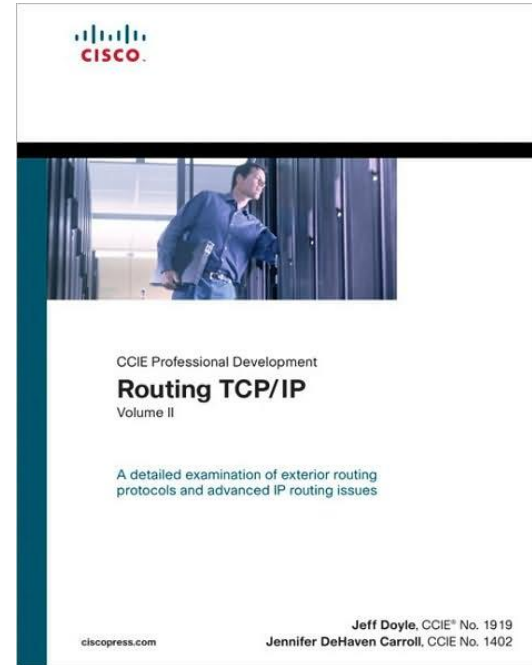
Preparation Reading



BRKCRT-8002



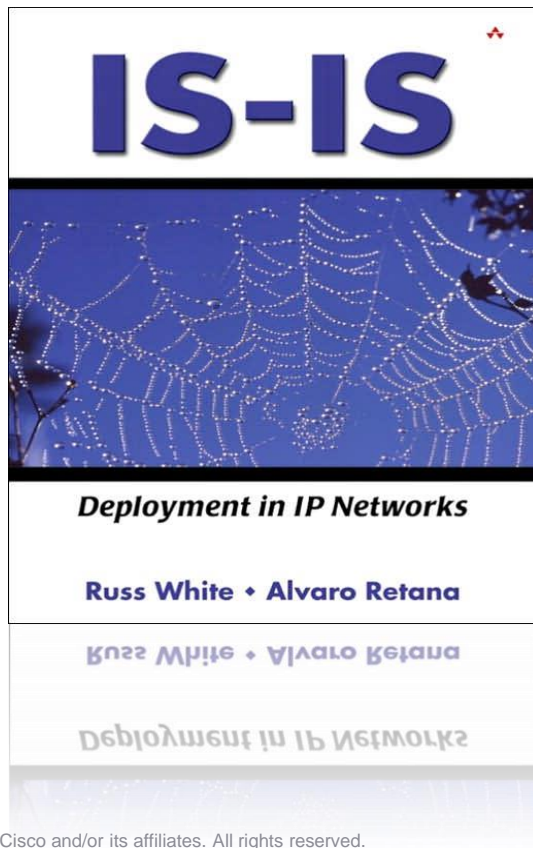
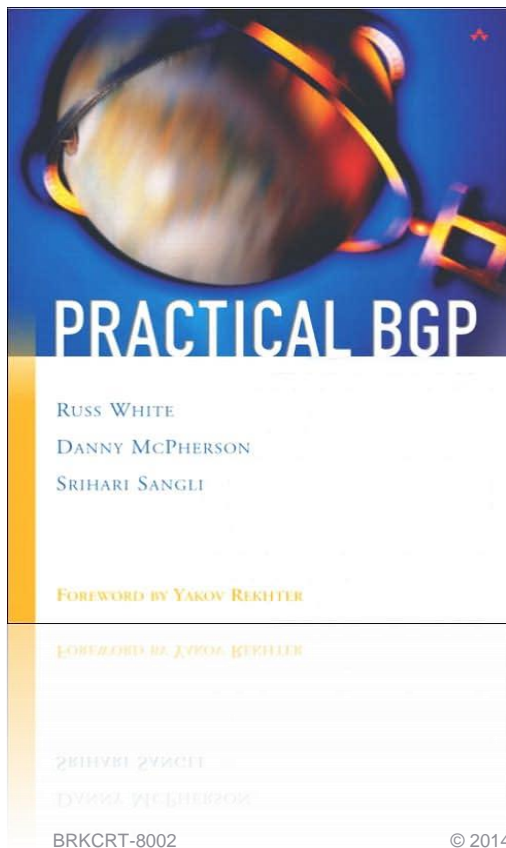
© 2014 Cisco and/or its affiliates. All rights reserved.



Cisco Public

Cisco *live!*

Preparation Reading





CISCO™