I-NET ©
Draft and Copyright

Design Principles used in an Enterprise Network

I-NET©
Draft and Copyright

A Company

- Think about a company and its structure.
- If the company is small or based in one location, the network might be rather small and have a modest number of users. The company might have a server room or a small data center.

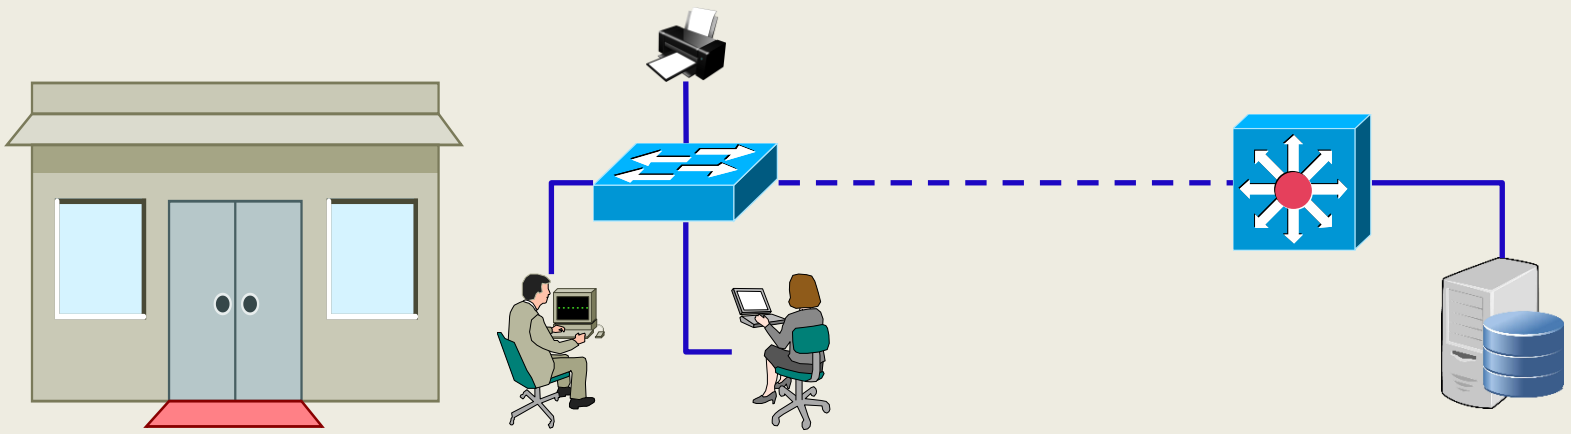

I-NET©
Draft and Copyright

A Company

- If the company is large, it might be spread over many geographic locations, with many different campuses and remote sites. The data center might be located in one large space, or it may be distributed across several redundant locations. The company might have an enterprise wide IP telephony deployment, video over IP, network security devices, a growing teleworker community, multiple Internet connections, and several network connections to partner companies.

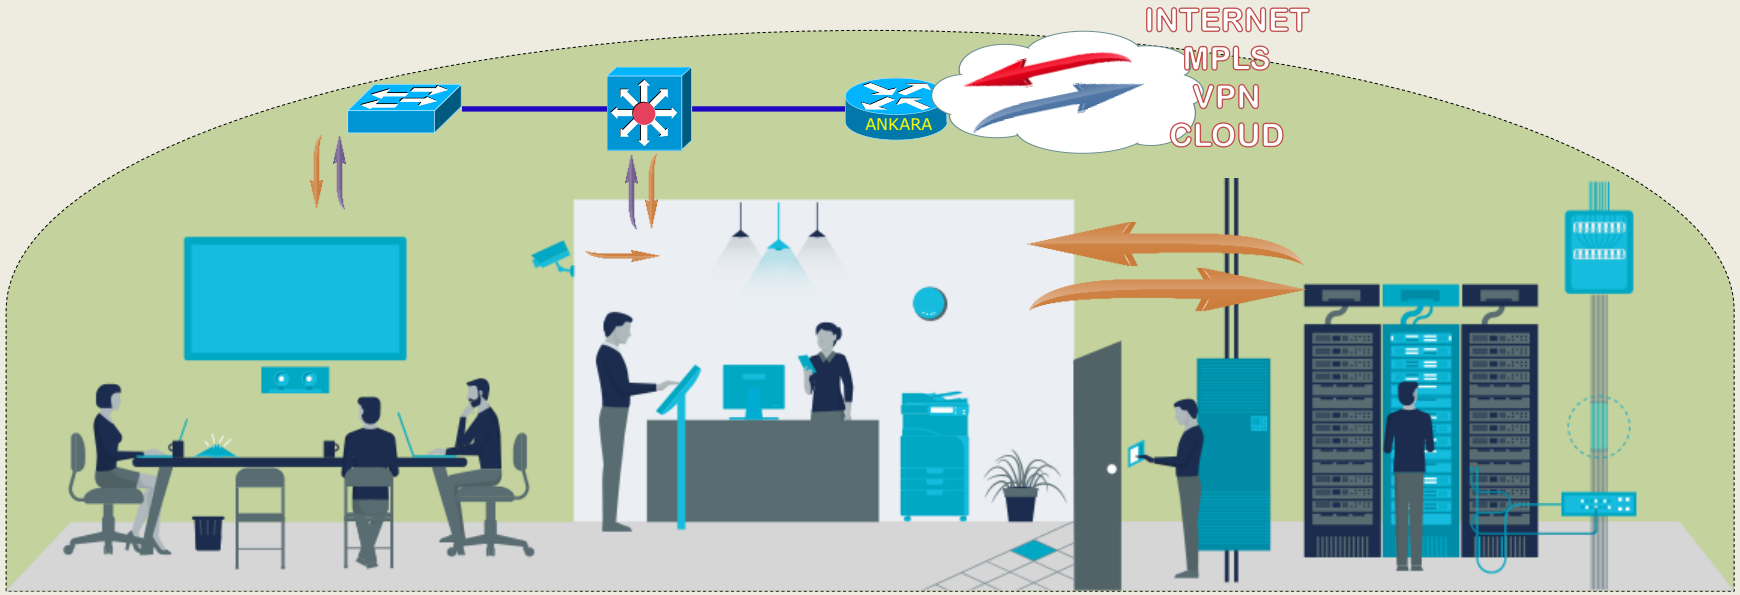

I-NET ©
Draft and Copyright

Networking Staff

- Regardless of the size of the company, the IT needs are basically the same.
- The roles in the network and the type of work done by the people in those roles:

❑ **HELPDESK**
- first level of support for problems
- taking a general problem statement
- diagnosis of network health

❑ **NETWORK ENGINEER**
- third level of support for problems,
- focus on project work;
  - detailed planning for new configurations
  - new network features

❑ **NETWORK TECHNICIAN**
- second level of support for problems
- reacting to calls from the help desk
- monitoring the network proactively
- implements changes on behalf of the Network engineers

❑ **NETWORK DESIGNER**
- gathering requirements from internal and external customers,
- translating those requirements into a network design

I-NET ©
Draft and Copyright

Enterprise Campus Network Design

I-NET©
Draft and Copyright

Network Design

- **A vital part of the campus network design** is an understanding of traffic flow including;
  - Speed
  - Bandwidth

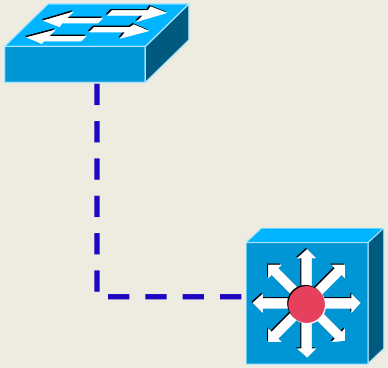

- The network traffic can then be effectively deployed and managed,
- You can also scale the campus network to support future needs.

I-NET ©
Draft and Copyright

Hierarchical Network Design

I-NET©
Draft and Copyright

Simple Shared Ethernet Network

- All devices on this network segment must share the available bandwidth.

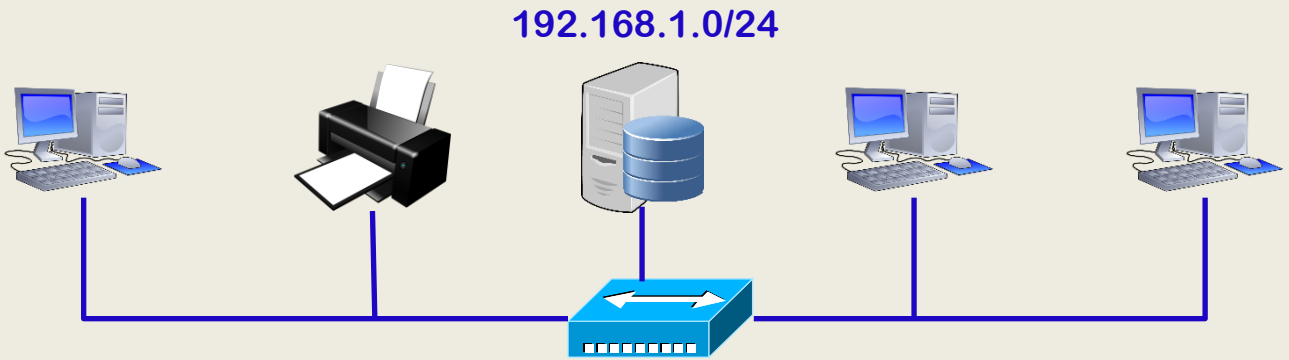

CSMA/CD

- Collision/Collision Domain ??
- Broadcast/Broadcast Domain ??
- Problem starts when you add more users to the network.

I-NET©
Draft and Copyright

Simple/Flat Ethernet Network

- Each host should connect to a dedicated switch interface so that they can operate in full-duplex mode, preventing collisions altogether.

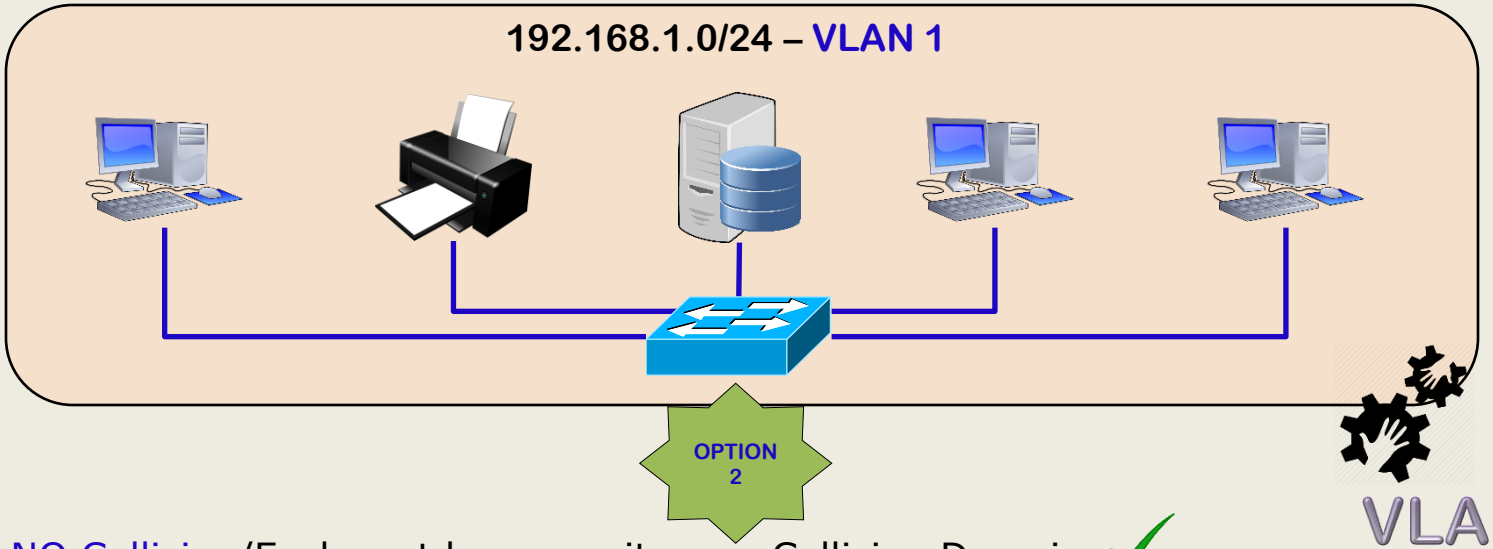

- NO Collision/Each port becomes its own Collision Domain ✓
- Broadcast/Broadcast Domain ?
- A VLAN defines the scope of a broadcast domain.

I-NET©
Draft and Copyright

Network Segmentation

- To reduce the size of a broadcast domain, break it up into smaller Layer-2 VLANs. The smaller VLANs must be connected by a Layer-3 device, such as a router or a multilayer switch.

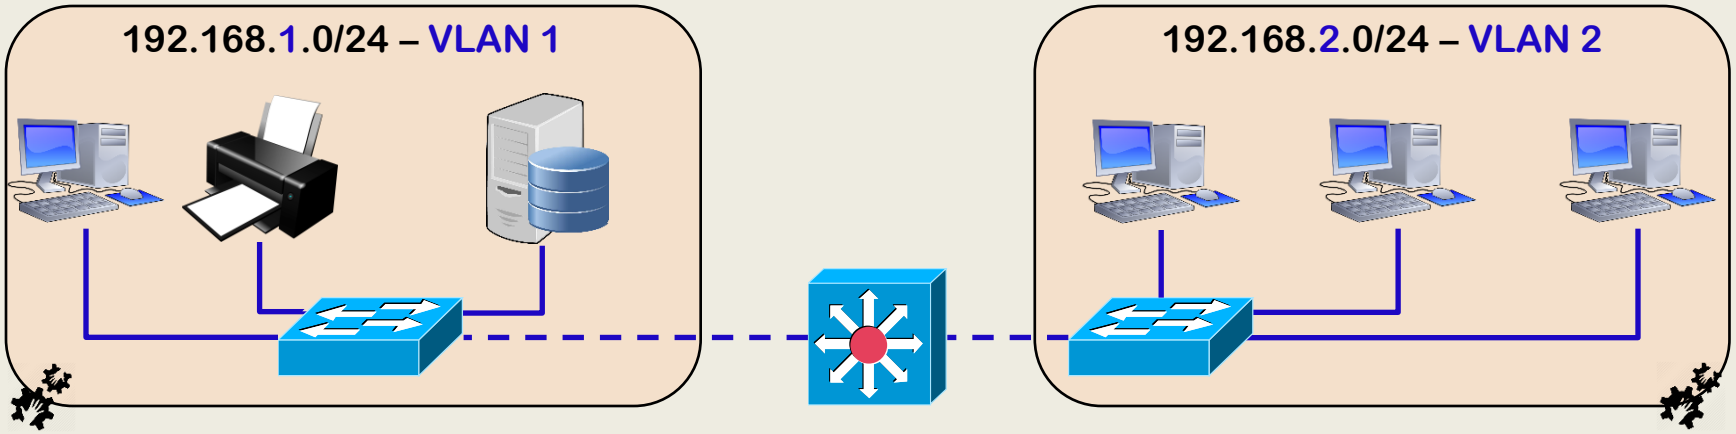

  - Limited number of ports on switches
  - Add more switches to the network

I-NET©
Draft and Copyright

Network with new VLANs

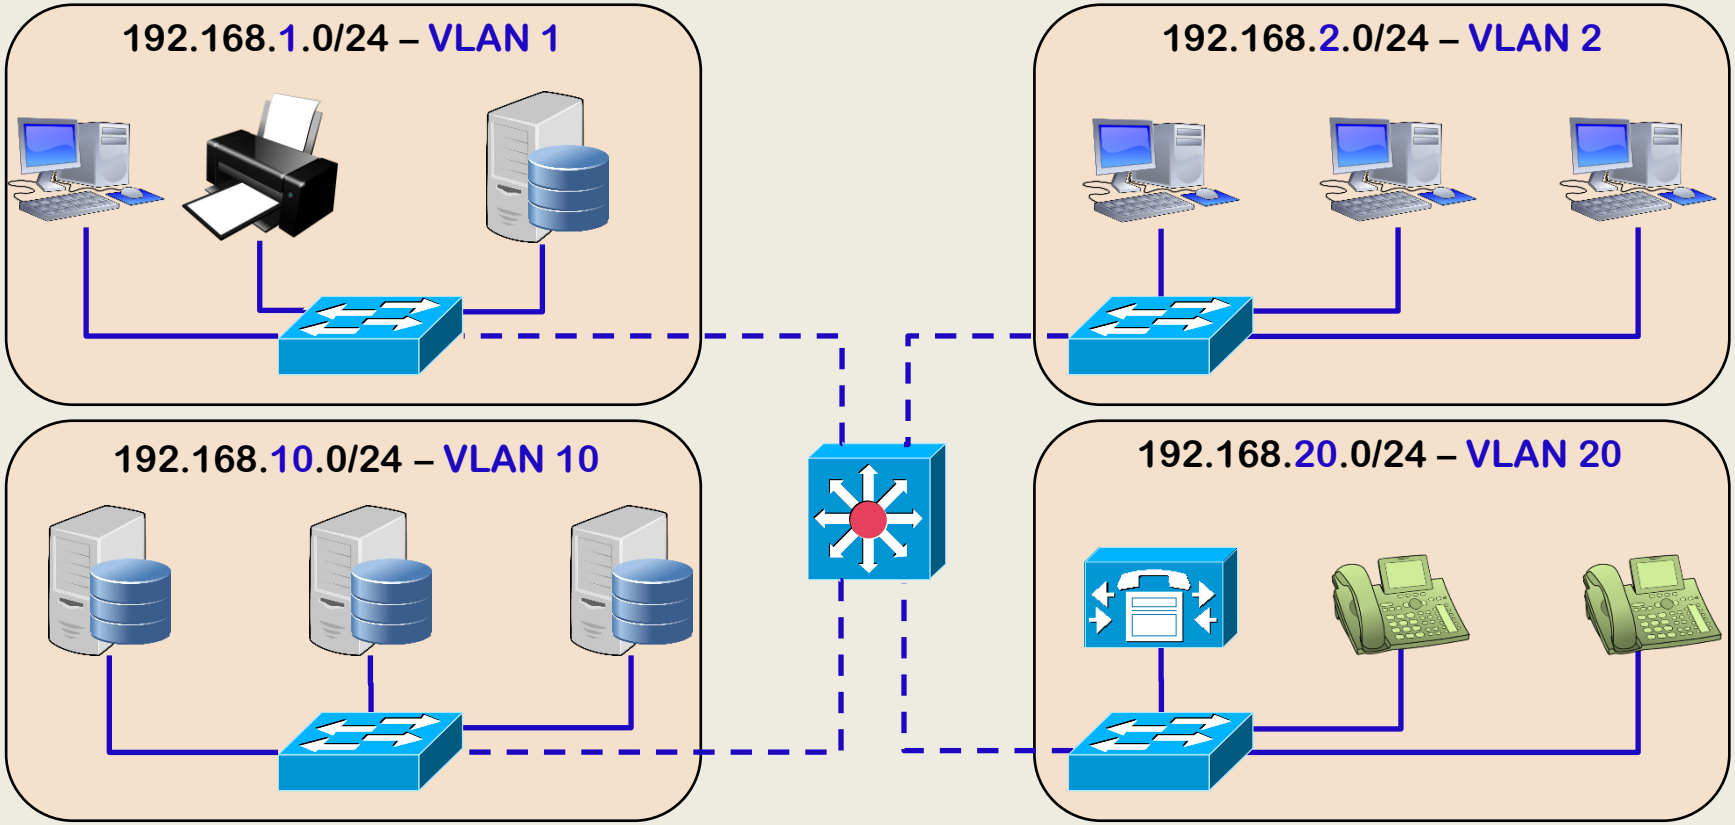

I-NET ©
Draft and Copyright

Hierarchical Network Design

- During the hierarchical network design, the network should;
  - Be in a predictable behavior in mind;
    - A network needs to recover from failures and topology changes quickly and in a predetermined manner.
  - Offer low maintenance and high availability.
  - Support future expansions and upgrades.
  - Support variety of multiprotocol and multicast traffic;
    - The network should be capable of efficiently connecting users with the resources they need, regardless of location.

Trunking, STP, EtherChannel, HSRP, ARP, TCP,....

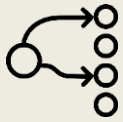

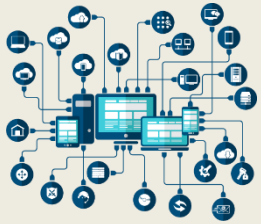

I-NET ©
Draft and Copyright

Hierarchical Network Design

- Cisco has well-known hierarchical approach to network design that enables network designers to organize the network into different layers of devices.
- The resulting network is efficient, intelligent, scalable, and easily managed.

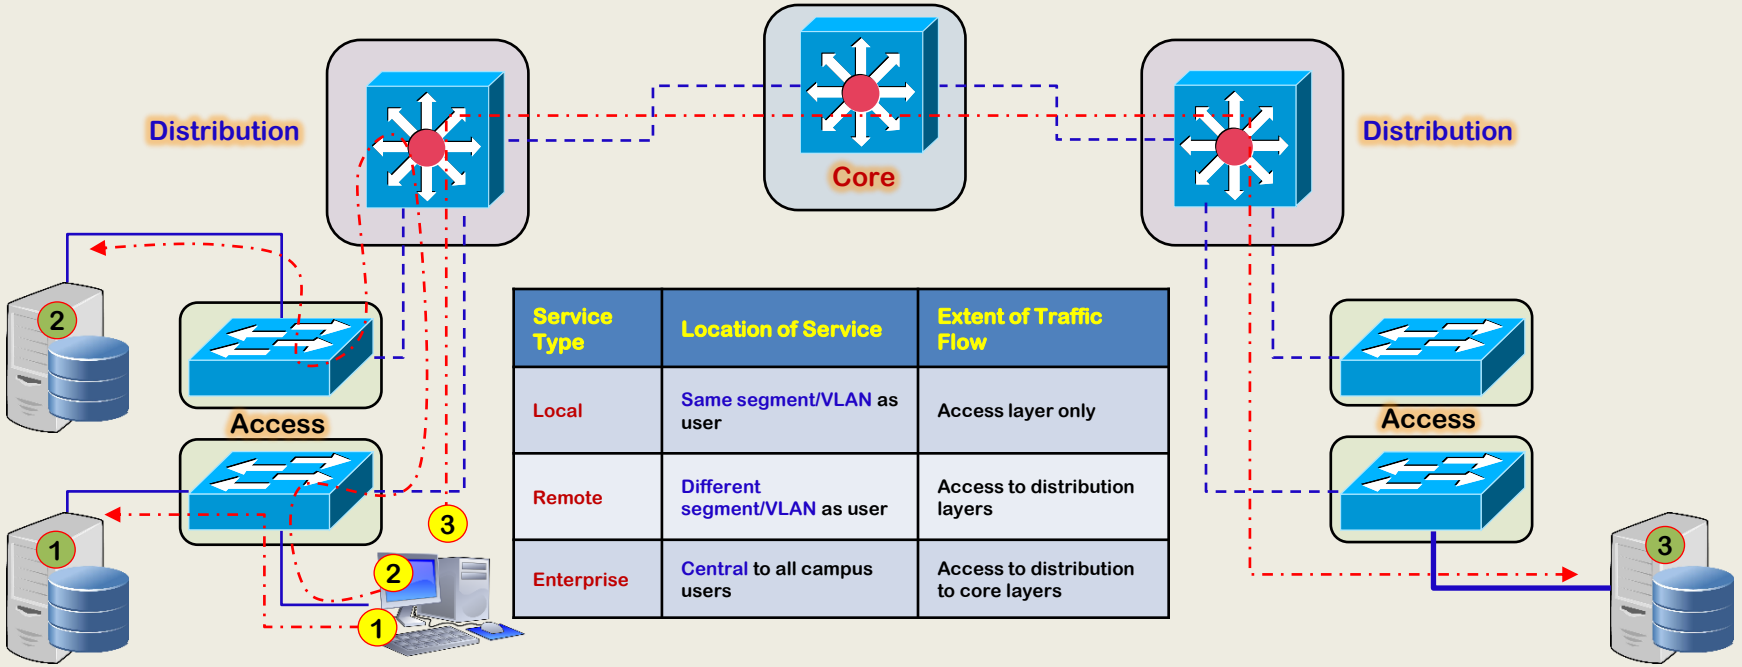

Service Type	Location of Service	Extent of Traffic Flow
Local	Same segment/VLAN as user	Access layer only
Remote	Different segment/VLAN as user	Access to distribution layers
Enterprise	Central to all campus users	Access to distribution to core layers

I-NET©
Draft and Copyright

Access Layer

- The **access layer** exists where the end users are connected to the network.
- Access switches usually provide Layer-2 (VLAN) connectivity between users.
- Devices in this layer, sometimes called building access switches, should have the following capabilities:
  - Low cost per switch port
  - High port density
  - Scalable uplinks to higher layers
  - High availability
  - Ability to converge network services (for data, voice, video)
  - Security features and quality of service (QoS)

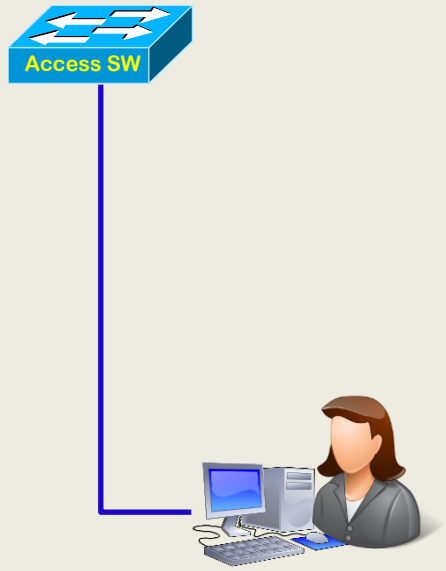

I-NET©
Draft and Copyright

Distribution Layer

- The **distribution layer** provides interconnection between the access and core layers, should have the following capabilities:
  - Aggregation of multiple access layer switches: In the distribution layer, uplinks from all access layer devices come together.
  - High Layer-3 routing throughput for packet handling: Must be capable of processing the total volume of traffic from all the connected devices.
  - Security and policy-based connectivity functions: VLANs and broadcast domains converge at the distribution layer, requiring routing, filtering, and security.
  - QoS features
  - Scalable and redundant high-speed links to the core and access layers

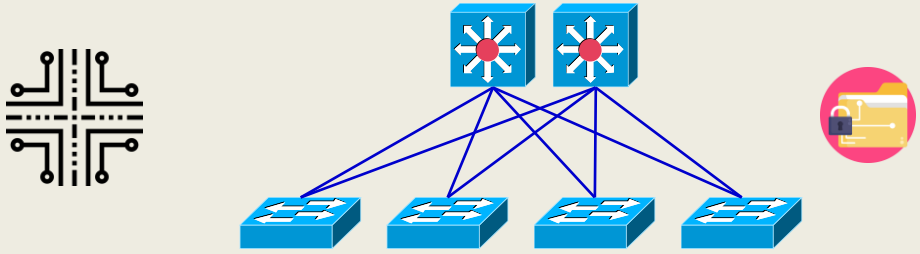

I-NET©
Draft and Copyright

Core Layer

- A campus network's **core layer** provides connectivity of all distribution-layer devices. The core, must be capable of switching traffic as efficiently as possible. Core devices, sometimes called campus backbone switches, should have the following attributes:
  - Very high throughput at Layer 3
  - No costly or unnecessary packet manipulations (access lists, packet filtering)
  - Redundancy and resilience for high availability
  - Advanced QoS functions
- Devices in a campus network's core layer or backbone should be optimized for high-performance switching.
- The core layer should be designed with simplicity and efficiency in mind.

I-NET©
Draft and Copyright

Collapsed Approach

- Although campus network design is presented as a three-layer approach (access, distribution, and core layers), the hierarchy can be collapsed or simplified in certain cases.
- **Small or medium-size campus networks** might **not have** the size or volume requirements that would require the functions of all three layers.
- In that case, combine the distribution and core layers for simplicity and cost savings.
- When the distribution and core layers are combined into a single layer of switches, it is called collapsed core.

Three-layer approach

Collapsed Core

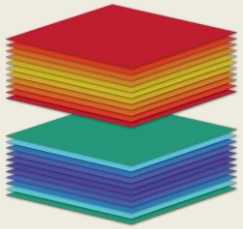

I-NET ©
Draft and Copyright

Modular Network Design

I-NET©
Draft and Copyright

Modular Network Design

- Designing a new network that has a hierarchy with three layers is fairly straightforward. The resulting network is organized, efficient, and predictable.
- However, a simple hierarchical design does not address other best practices like redundancy, and scalability.

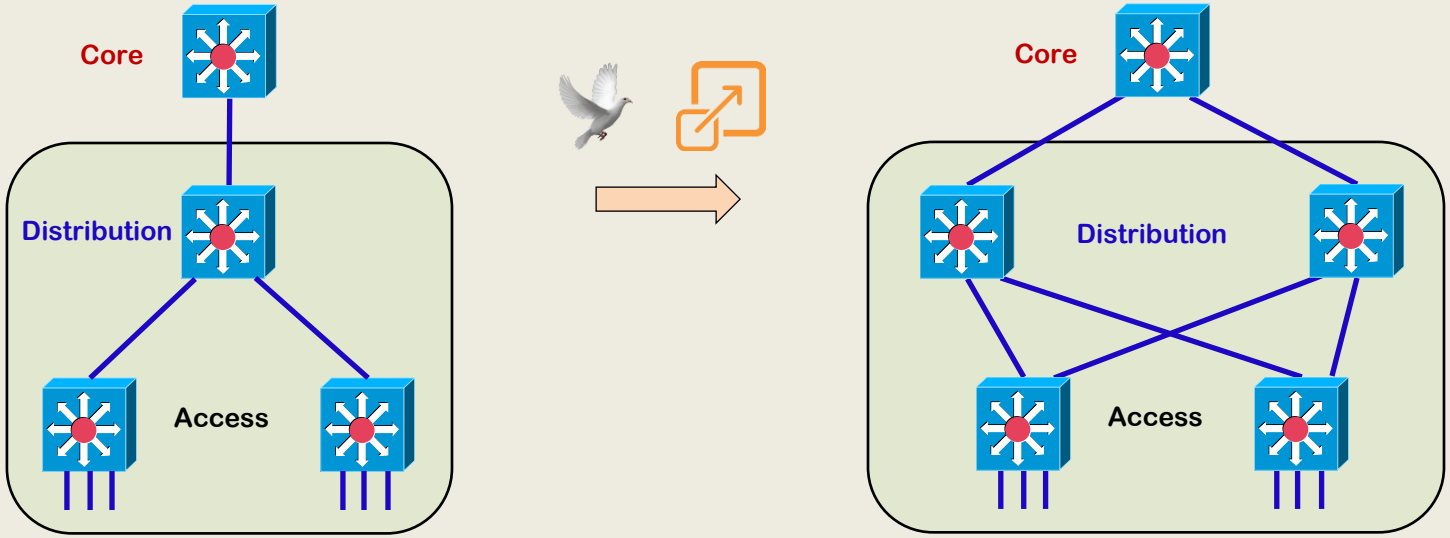

I-NET©
Draft and Copyright

Modular Network Design

- One weakness is still present in the redundant design. The core layer has only one switch. If that core switch fails, users in the access layer will still be able to communicate with each other. However, they will not be able to reach other areas of the network, such as a data center, the Internet, and so on.
- To mitigate the effects of a core switch failure, add a second, redundant core switch.

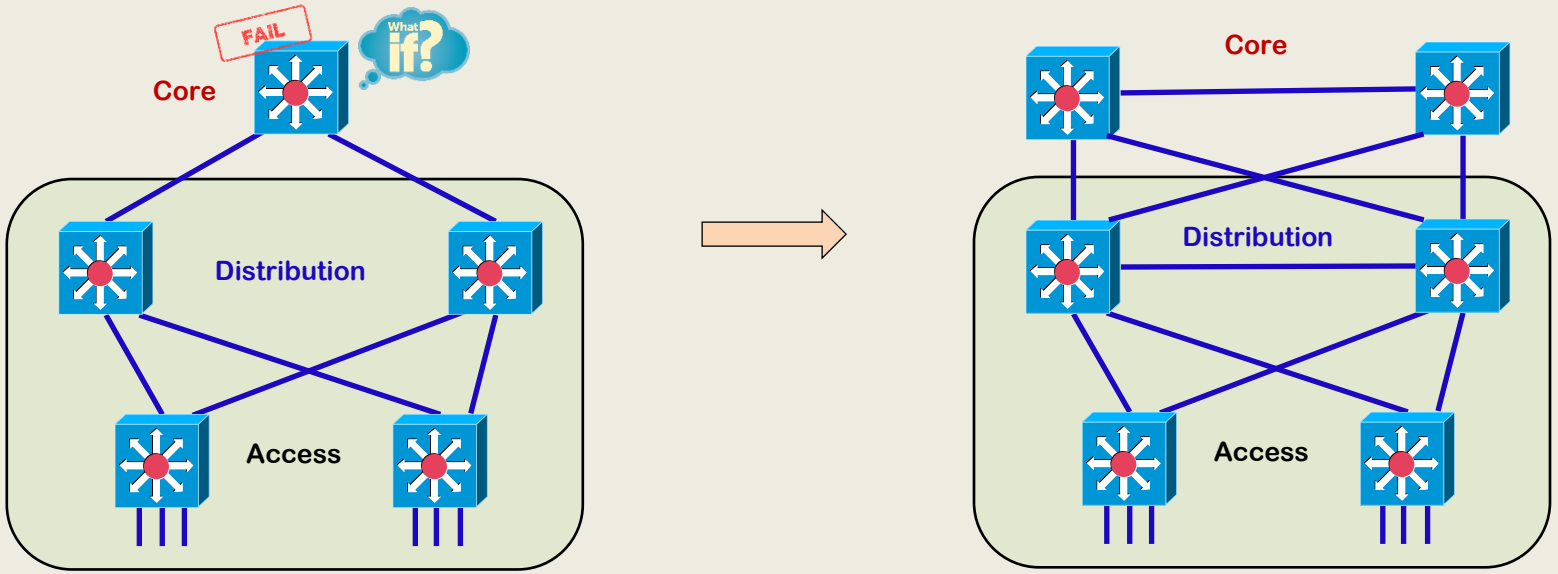

I-NET©
Draft and Copyright

Modular Network Design

- To maintain organization, simplicity, and predictability; design a campus network in a logical manner, using a modular approach.
  - In this approach, each layer of the hierarchical network model can be broken into basic functional units. These units, or modules, can then be sized appropriately and connected, while allowing for future scalability and expansion.
- Enterprise campus networks can be divided into the following basic elements or building blocks:
  - Switch block: A group of access layer switches, together with their distribution switches.
  - Core: The campus network's backbone, which connects all switch blocks.
- Other related elements can exist; A **data center** containing enterprise resources or services can have its own access and distribution layer switches, forming a switch block that connects into the core layer.
- If the data center is very large, it might have its own core switches.

I-NET©
Draft and Copyright

Modular Network Design

- As the network grows, add a completely new Access/distribution switch block that contains the areas of new growth.

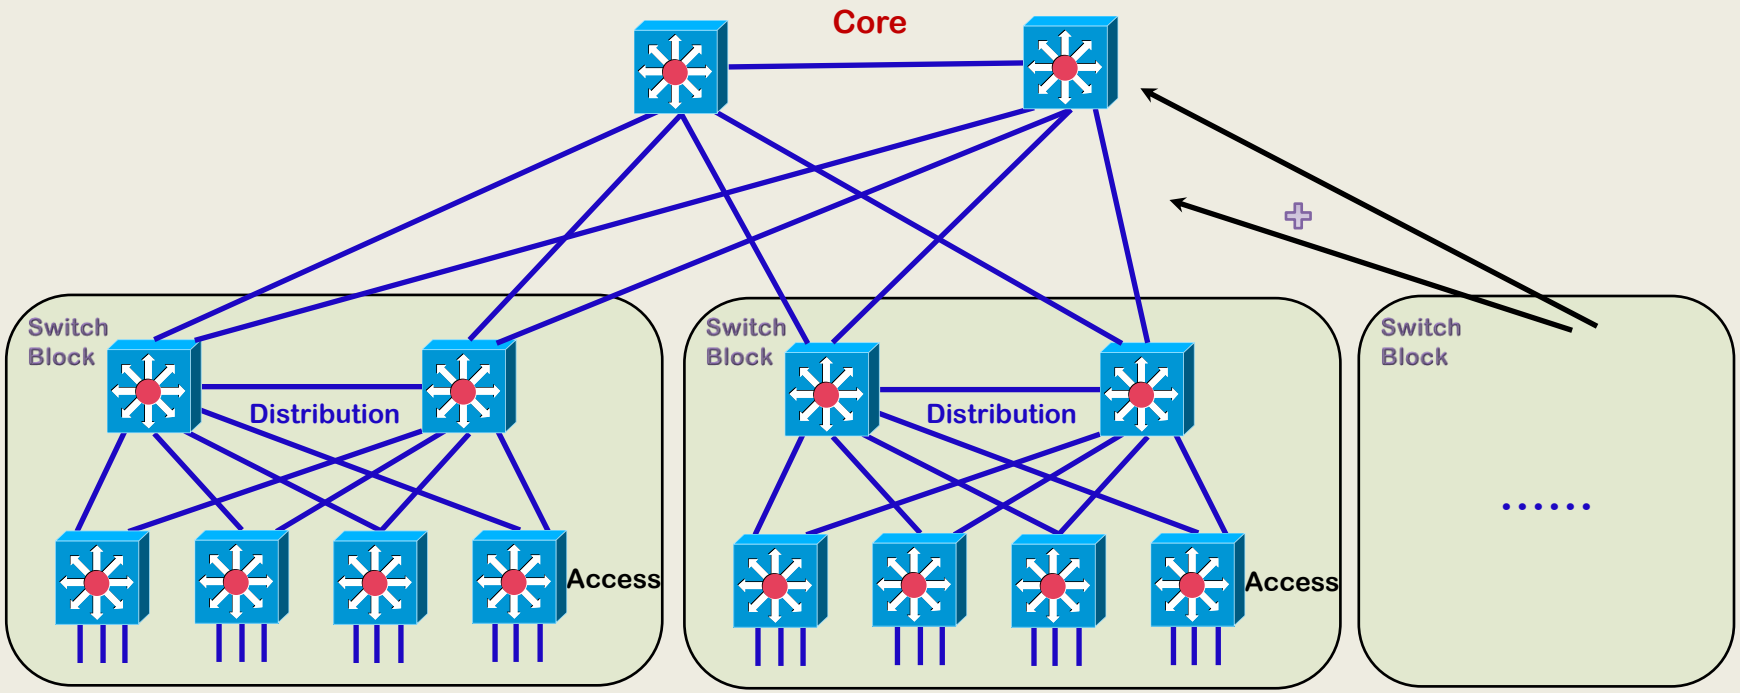

I-NET©
Draft and Copyright

Sizing a Switch Block

- Containing access and distribution layer devices, the switch block is simple in concept.
  - At the access layer, switch selection is usually based on port density or the number of connected users.

Architecture Metric	Cisco Catalyst 9400	Cisco Catalyst 9600
Intended Layer	Access / Building Aggregation	Campus Core / Distribution
Switching Capacity	Up to ~9 Tbps	Up to ~25.6 Tbps
Fabric Bandwidth	Lower fabric capacity optimized for PoE access	High-capacity switching fabric
Packet Forwarding Rate	Hundreds of Mpps	Several billion packets per second
Maximum Uplink Speeds	10G / 25G / 40G	10G / 25G / 100G / 400G
Routing Table Scale	Moderate campus routing	Massive campus core routing scale

Forwarding Architecture Comparison

- The distribution layer must be sized according to the number of access layer switches that are aggregated into a distribution device. Consider the following factors:
  - Traffic types and patterns
  - Amount of Layer-3 switching capacity at the distribution layer
  - Total number of users connected to the access layer switches
  - Geographic boundaries of subnets or VLANs

I-NET ©
Draft and Copyright

Sizing a Switch Block

- Generally, **a switch block is too large if the following conditions are observed:**
  - The routers (multilayer switches) at the distribution layer become traffic bottlenecks. This congestion could be because of the volume of inter-VLAN traffic, intensive CPU processing, or switching times required by policy or security functions (access lists, queuing, and so on).

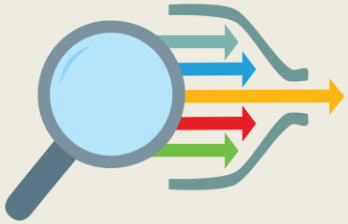

  - Broadcast or multicast traffic slows the switches in the switch block. Broadcast and multicast traffic must be replicated and forwarded out many ports simultaneously. This process requires some overhead in the multilayer switch, which can become too great if significant traffic volumes are present.

I-NET©
Draft and Copyright

Network Core

- A core should consist of two multilayer switches that connect two or more switch blocks in a redundant fashion. A redundant core is sometimes called a dual core because it is usually built from two identical switches.
- A core layer is required to connect two or more switch blocks in a campus network. Because all traffic passing to and from all switch blocks must cross the core, the core layer must be as efficient and resilient as possible.
- The core is the campus network's basic foundation and carries much more traffic than any other switch block.
- The links between layers should be designed to carry the amount of traffic load handled by the distribution switches, at a minimum.
- The links between core switches should have sufficient size to carry the aggregate amount of traffic coming into one of the core switches.
- Consider the average link utilization, but allow for future growth.

I-NET©
Draft and Copyright

Network Core

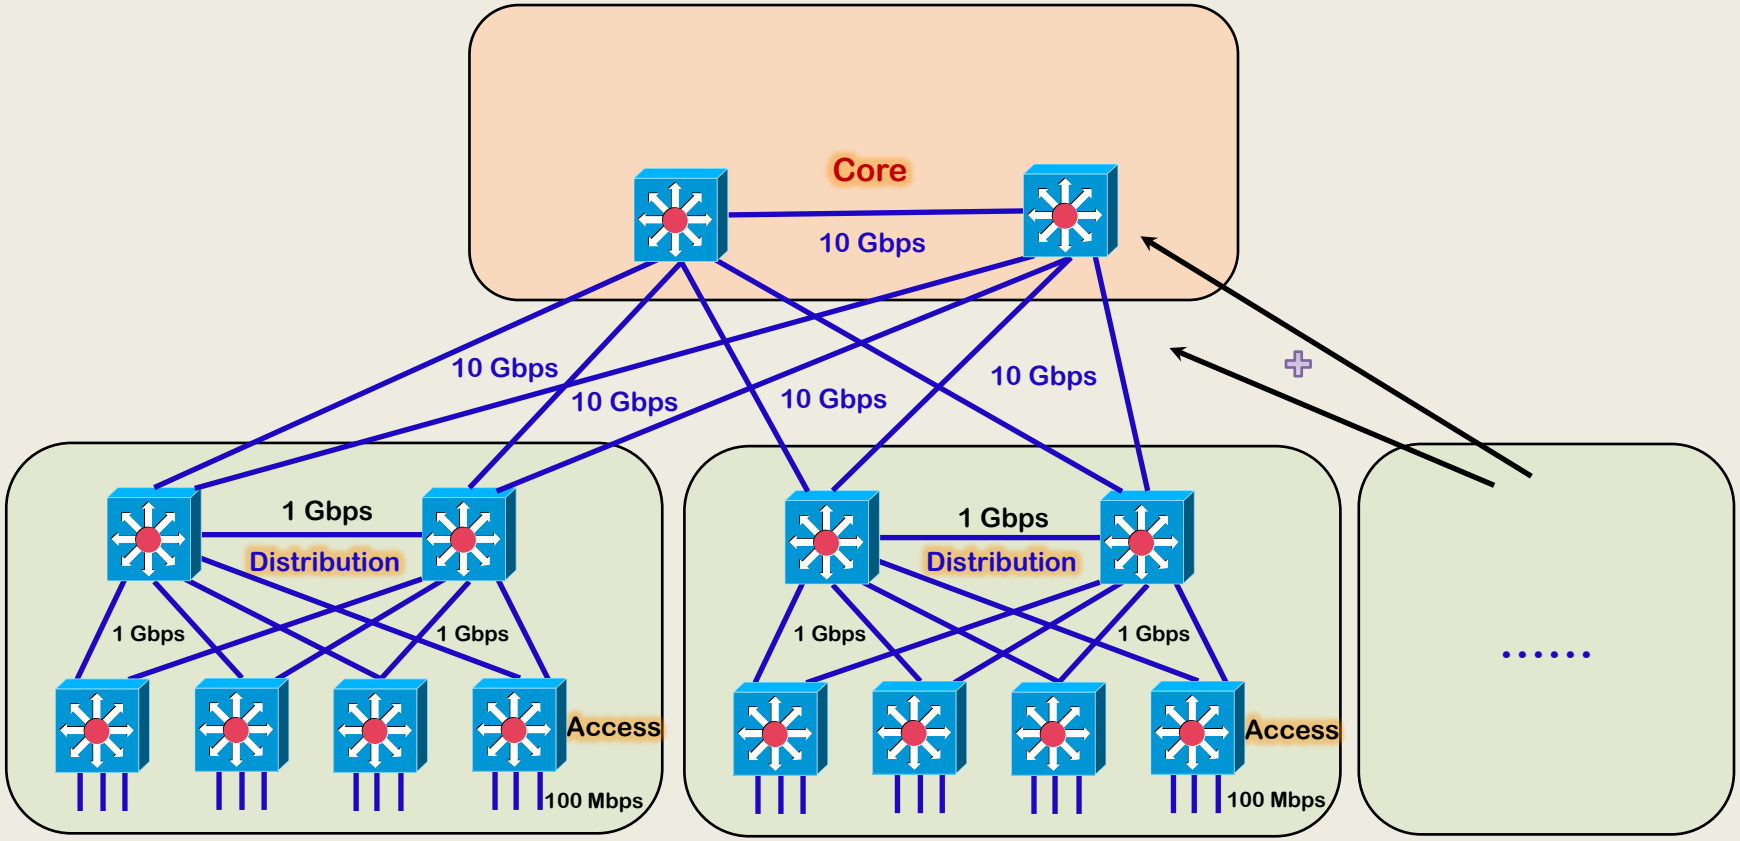

I-NET ©
Draft and Copyright

Best practices Hierarchical Network Design

- Best practices that can help with a redundant hierarchical network design:
  - Design each layer with pairs of switches.
  - Connect each switch to the next higher layer with two links for redundancy.
  - Connect each pair of distribution switches with a link, but do not connect the access layer switches to each other (unless the access switches support some other means to function as one logical stack or chassis).
  - Do not extend VLANs beyond distribution switches. The distribution layer should always be the boundary of VLANs, subnets, and broadcasts. VLAN traffic should not traverse the network core.

I-NET©
Draft and Copyright

Best Practice Loop-Free Switch Block Topology

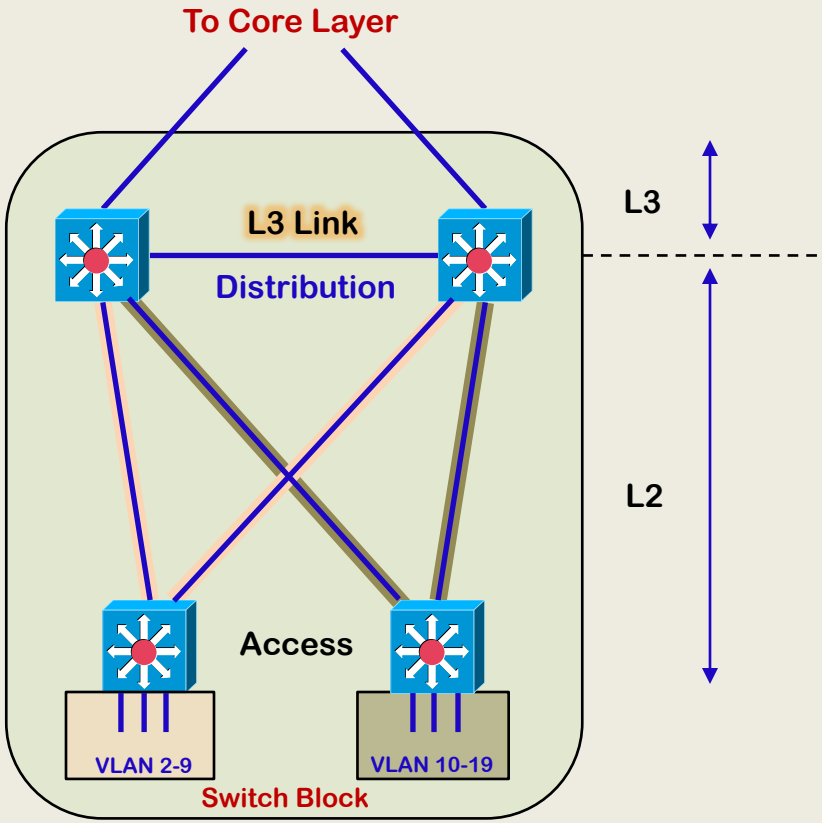

I-NET©
Draft and Copyright

Best Practice A Completely Routed Switch Block

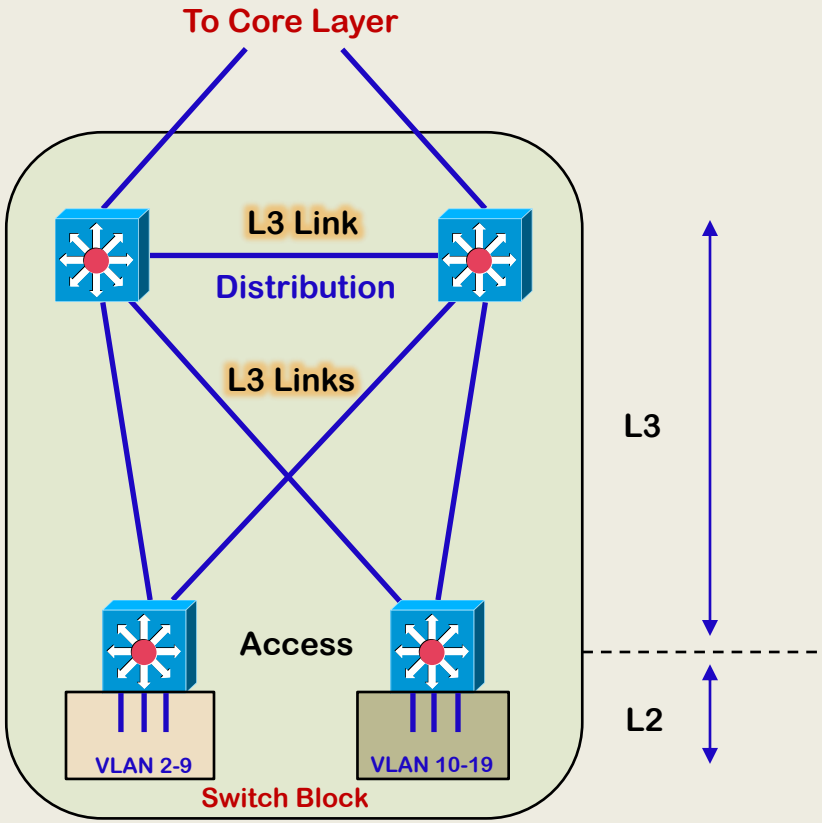

I-NET ©
Draft and Copyright

Data Centers Network

Types

Architecture

Topology

I-NET ©
Draft and Copyright

Data Centers Network

I-NET©
Draft and Copyright

Spine-Leaf Architecture

- Spine-leaf architecture is a modern network topology **widely used in data centers and cloud environments.** It's designed to offer improved **scalability, high availability, and consistent low-latency communication.**
- Unlike traditional three-tiered hierarchical networks (core, distribution, access layers), the spine-leaf architecture **consists of only two layers of switches:**

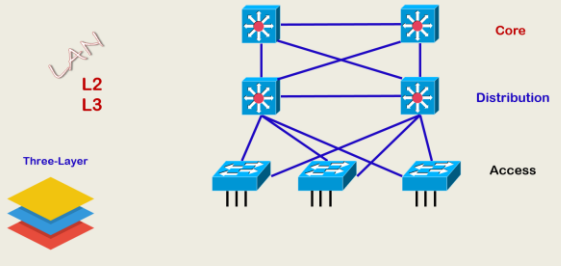

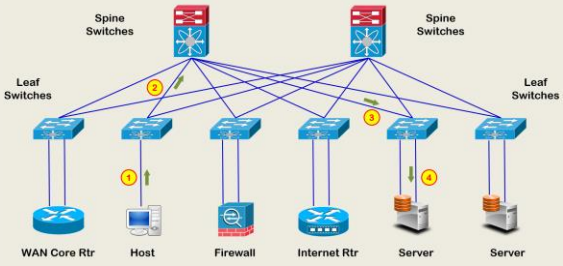

- Spine and leaf switches:
  - **The spine layer** consists of switches that perform routing and work as the core of the network.
  - **The leaf layer** involves access switches that connect to servers, storage devices, and other end-users.
- This structure helps data center networks reduce hop count, reduce network latency and prevent bottlenecks, which are main issues of three-tier architecture.

I-NET ©
Draft and Copyright

Cisco SD-Access

I-NET©
Draft and Copyright

Network Landscape Challenges

- **This massive increase of networked devices produces new network requirements.** The one thing that all users demand is **access anytime, anywhere, to all the resources that they require, which is a serious challenge. The fact that user locations can be mobile makes this networking landscape even more difficult.** How can you ensure that the **right people and devices are connecting to the right resources in a dynamically changing environment? How can you block devices that should not communicate with each other?**

I-NET©
Draft and Copyright

Network Management Scalability Challenges

- With the expanding and changing network landscape, network management scalability is also becoming a challenge.

95%
Network Changes Performed Manually

75%
Of OpEx Spent on Changes And Troubleshooting

70%
Of Policy Violations are Due to Human Error

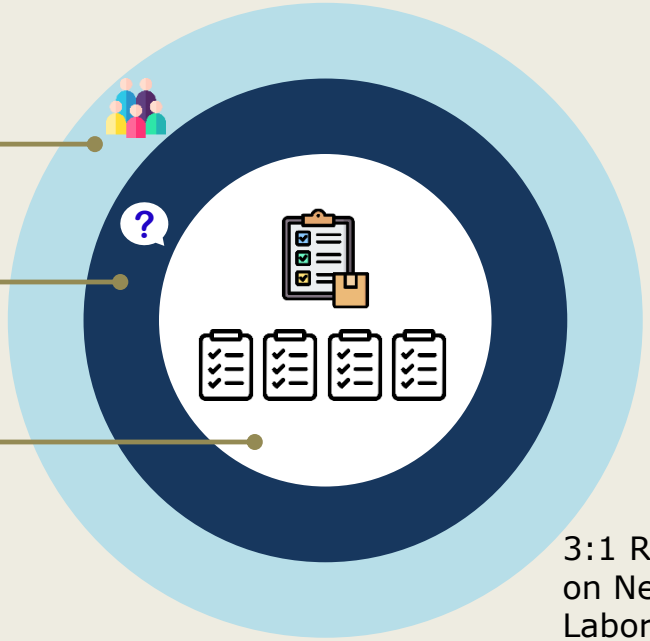

3:1 Ratio of OpEx to CapEx on Network Operations, Labor, and Tools

I-NET©
Draft and Copyright

Transition to Software-Defined Networking

- As your business plans for this new dynamic landscape, you may want to transition your traditional enterprise network to software-defined networking technology to address the many challenges that IoT presents.

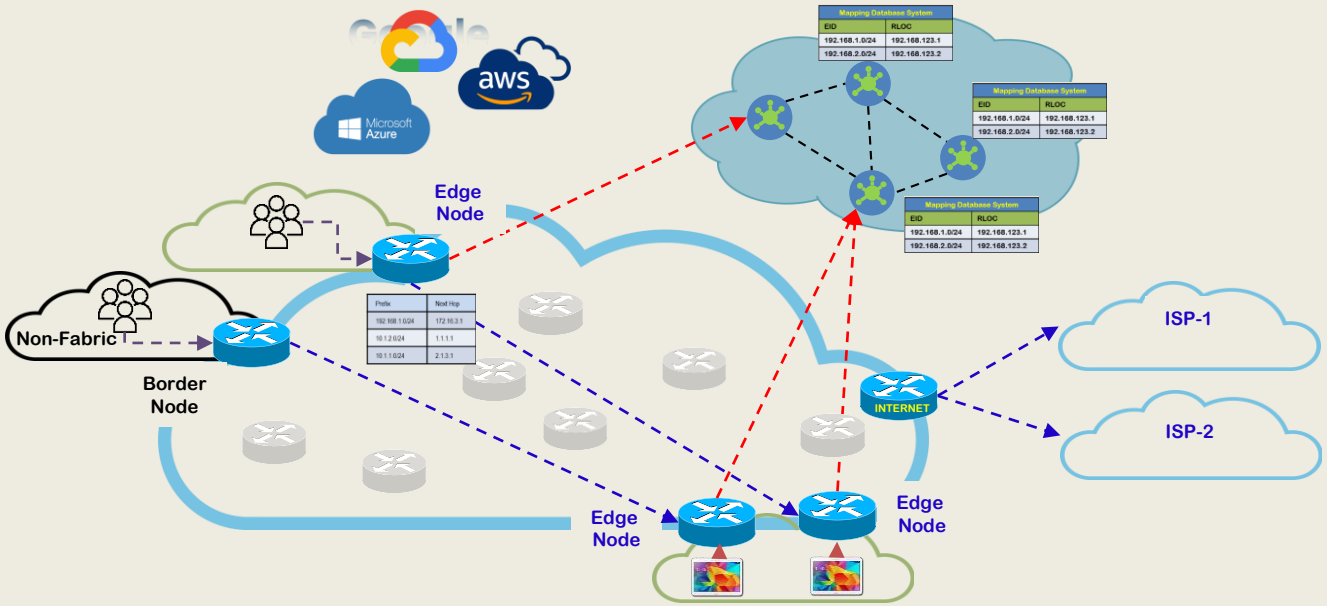

Software-Defined Access

- **Cisco Software-Defined Access** works by decoupling a network's logical functions from its physical infrastructure through a fabric-based architecture. This allows for policy-based automation, where network access is determined by a user's identity rather than their physical location or IP address.
- The solution relies on four integrated foundational elements:
  - Controller-Based Architecture: **Cisco Catalyst Center** (formerly DNA Center) acts as the centralized "brain" for management and orchestration, automating device configuration and providing network visibility.
  - Policy Enforcement Engine: **Cisco Identity Services Engine** (ISE) manages authentication and authorization, assigning Scalable Group Tags (SGTs) to users and devices to enforce granular access policies.
  - Network Fabric: **The physical** and **virtual infrastructure** that carries traffic. It is divided into an underlay (physical transport) and an overlay (logical tunnels) for user traffic.
  - Programmable Infrastructure: **Standardized hardware** (like Catalyst 9000 Series) that supports APIs and telemetry to enable automation and continuous monitoring.

I-NET©
Draft and Copyright

Cisco SD-Access Architecture

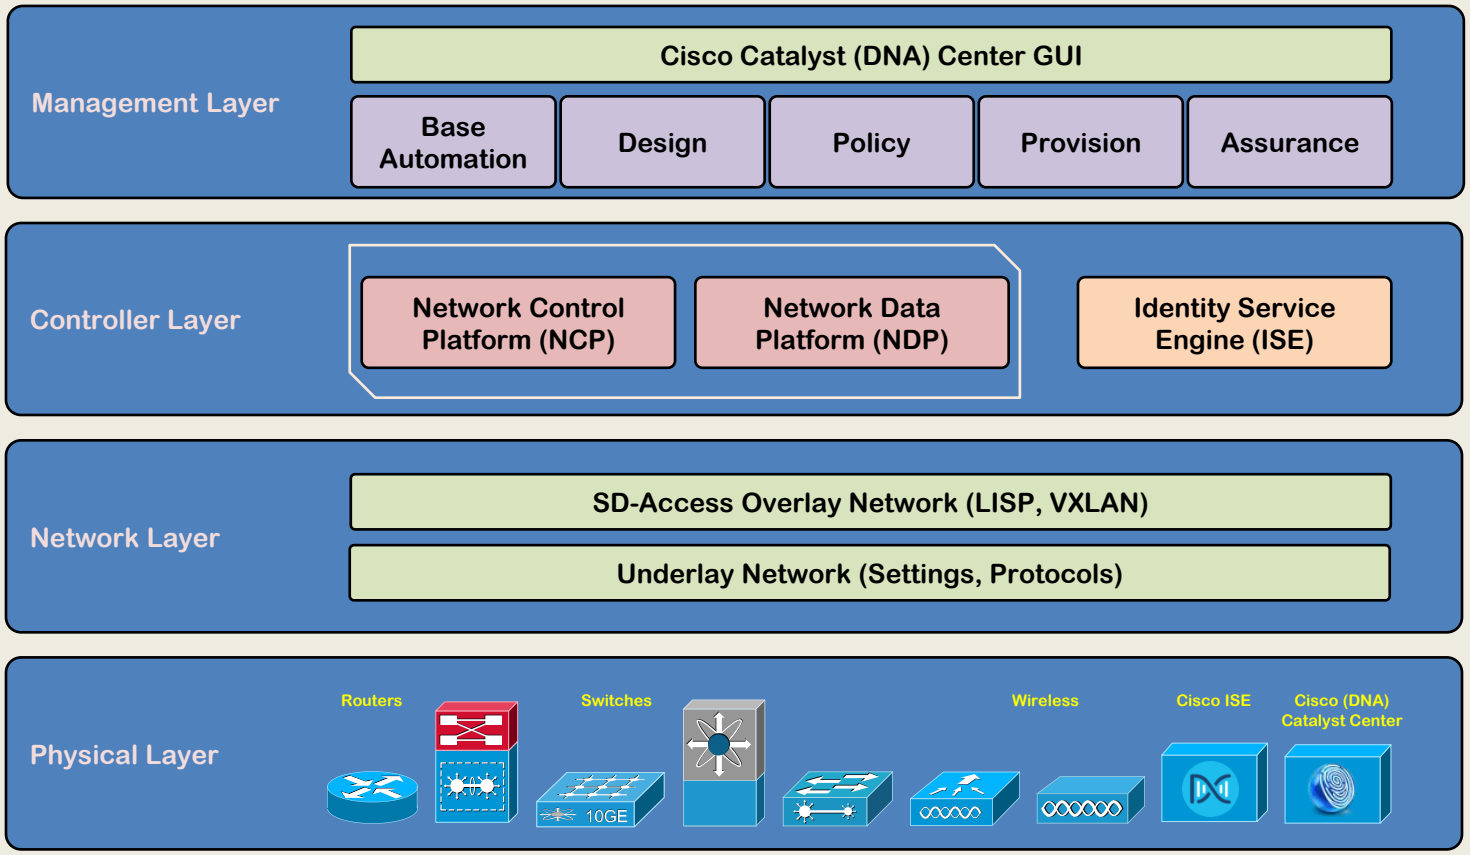

I-NET ©
Draft and Copyright

Cisco
SD-WAN

I-NET©
Draft and Copyright

Cisco Catalyst SD-WAN Solution

- Cisco Catalyst SD-WAN (formerly Viptela) operates on the principle of separating the management, control, and data planes into a centralized architecture. This transformation replaces complex, box-by-box manual configuration with a secure, scalable, and automated virtual IP fabric overlay.

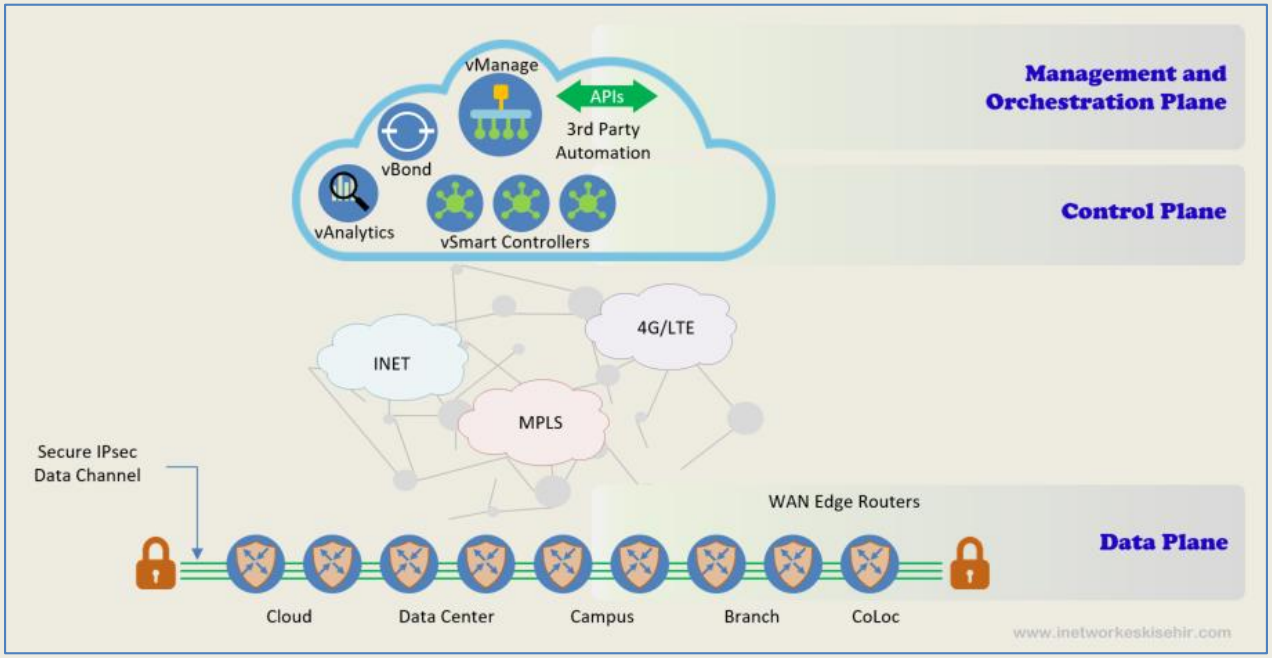

I-NET©
Draft and Copyright

Architectural Separation (The Planes)

- The solution divides network functions into four distinct layers to simplify management and improve scale:
  - Orchestration Plane (Validator/vBond): Coordinates the initial bring-up and authentication of devices. It ensures only authorized routers join the fabric.
  - Management Plane (Manager/vManage): Provides a single centralized dashboard (GUI) for configuration, monitoring, and troubleshooting. Admins push policies to the entire network with a few clicks instead of using CLI on individual devices.
  - Control Plane (Controller/vSmart): Acts as the "brains" of the network. It manages routing intelligence, security, and policy distribution using the Overlay Management Protocol (OMP).
  - Data Plane (Edge Routers/vEdge/cEdge): Responsible for the actual forwarding of user traffic based on policies received from the control plane.

I-NET ©
Draft and Copyright

Core Working Principles

- The fabric is built on five fundamental steps that modernize traditional WAN deployments:
  - Transport Independence: It abstracts the underlying physical transport (MPLS, Broadband, LTE, 5G) into a single logical "transport-agnostic" network.
  - Centralized Routing Intelligence: All edge devices advertise prefixes to the centralized vSmart controllers, which then reflect that information to other routers, eliminating the need for a full-mesh routing adjacency.
  - Automated Security: The fabric automatically establishes secure IPsec tunnels between sites. It handles authentication and encryption key exchanges automatically via the controllers, removing the need for manual VPN configurations.
  - Policy-Driven Reachability: Network administrators use centralized policies to influence how traffic flows; for example, directing critical VoIP traffic over MPLS while sending bulk backups over standard internet.
  - Application-Aware Routing: The system identifies applications (like Office 365 or Zoom) and monitors link performance (latency, jitter, loss) in real-time to select the best path for each specific app.

I-NET ©
Draft and Copyright

Key Protocols and Mechanisms

- OMP (Overlay Management Protocol): The primary protocol for sharing routing, security, and policy information across the fabric.
- BFD (Bidirectional Forwarding Detection): Runs within the secure tunnels to continuously monitor link health and performance for rapid failover.
- Zero-Touch Provisioning (ZTP): Automates the onboarding of new edge routers, allowing them to be configured and secured as soon as they are plugged into an internet connection.

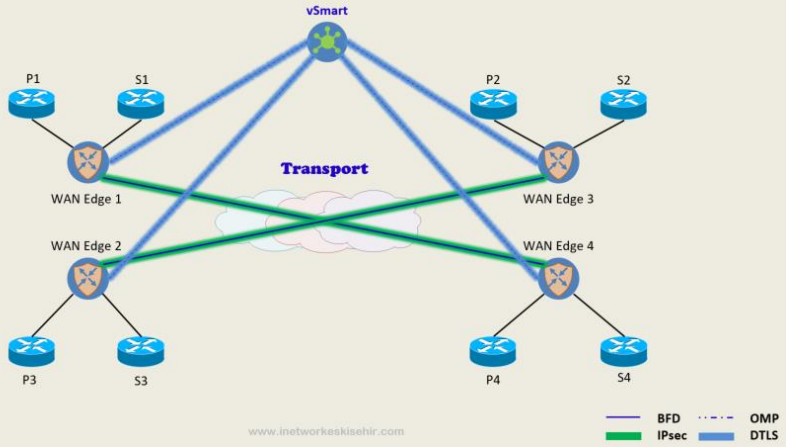

I-NET©
Draft and Copyright

QoS Configurations

I-NET ©
Draft and Copyright

QoS configurations

- Network devices don't really care about the type of traffic they have to forward. Switch receives an Ethernet frame, looks for the destination MAC address and forwards the frame towards the destination. The same thing applies to router, it receives an IP packet, looks for the destination in the routing table and it forwards the packet towards the destination.
- This forwarding logic is called best effort or FIFO (First In First Out).
- Sometimes, this can be an issue.

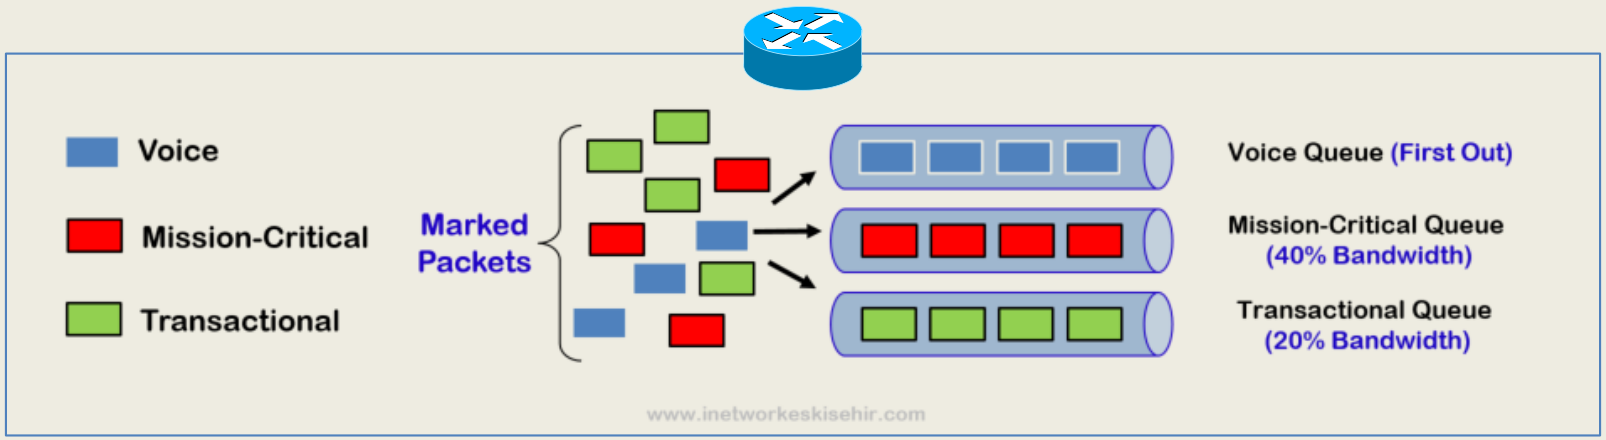

I-NET ©
Draft and Copyright

QoS configurations

- Cisco QoS configurations use the Modular QoS CLI (MQC) framework;
  - Class Maps,
  - Policy Maps,
  - Service Policies
- To classify, mark, and prioritize traffic, ensuring bandwidth for critical applications.

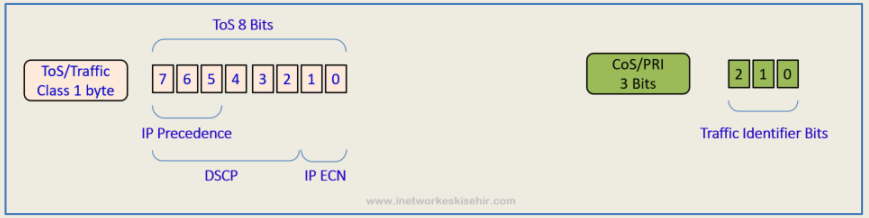

- Core Components of Cisco QoS (MQC)
  - Class Map (class-map): Defines the criteria to identify traffic (e.g., matching ACLs, DSCP values, or input interfaces).
  - Policy Map (policy-map): Defines actions for classified traffic, such as bandwidth guarantees, priority queuing, or dropping traffic.
  - Service Policy (service-policy): Applies the policy map to an interface, either incoming (input) or outgoing

I-NET©
Draft and Copyright

QoS configurations

- Key QoS Actions in Policy Maps;
  - Priority (LLQ): Provides low-latency queuing for voice/video.
  - Bandwidth Guarantee: Ensures minimum bandwidth for a class.
  - Traffic Policing: Drops or re-marks traffic exceeding a set rate.
  - Traffic Shaping: Buffers excess traffic to smooth traffic bursts.
  - Weighted Random Early Detection (WRED): Proactively drops TCP packets to avoid congestion.

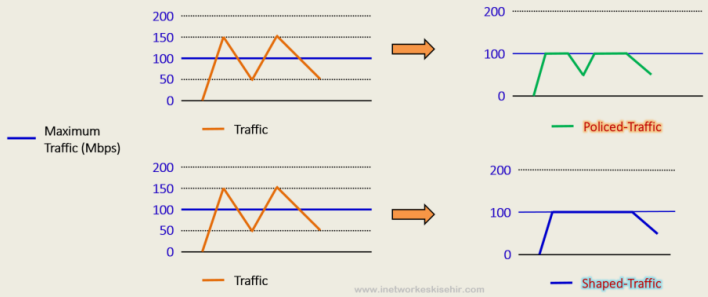

I-NET©
Draft and Copyright

Interpreting Sample Configuration

```
class-map match-any VOICE_CLASS
   match dscp ef
!
policy-map QOS_POLICY
   class VOICE_CLASS
      priority percent 30                    // Gives voice 30% priority bandwidth
   class class-default
      fair-queue                             // Queues all other traffic fairly
!
interface GigabitEthernet0/1
   service-policy output QOS_POLICY          // Applies policy outgoing
```

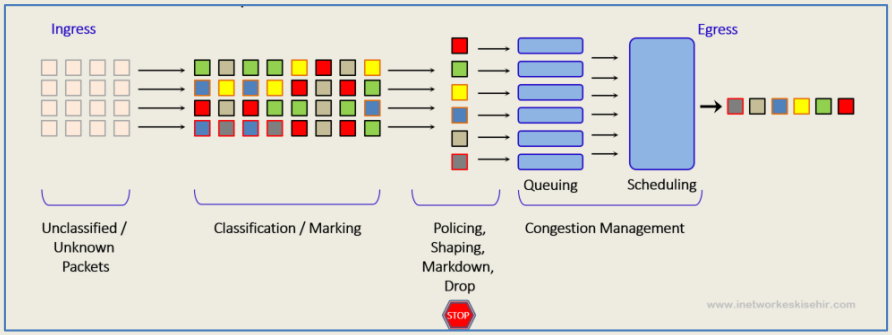

I-NET ©
Draft and Copyright

Cloud Networking

I-NET ©
Draft and Copyright

Cloud Networking

- **Cloud networking** offers connectivity to and between applications and workloads across clouds, cloud services, on-premises data centers, and edge networks.
- It's vital to performance, security, and efficient management of hybrid cloud and multi-cloud environments.
- Most application interactions begin at or terminate beyond an organization's own network, so cloud networking is becoming the new enterprise core.
- Designing, deploying, and maintaining a cloud network for agility, performance, security, and efficiency is vital to companies using public cloud services.

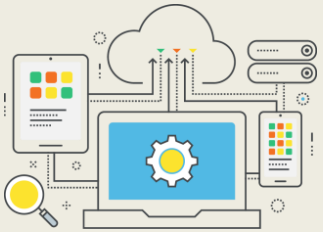

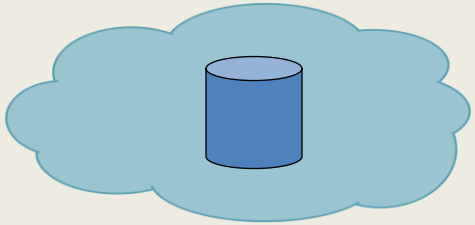

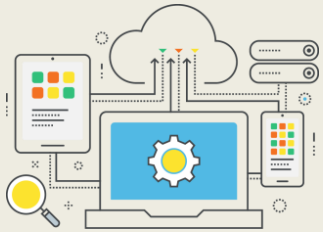

I-NET ©
Draft and Copyright

On-Premises and Cloud

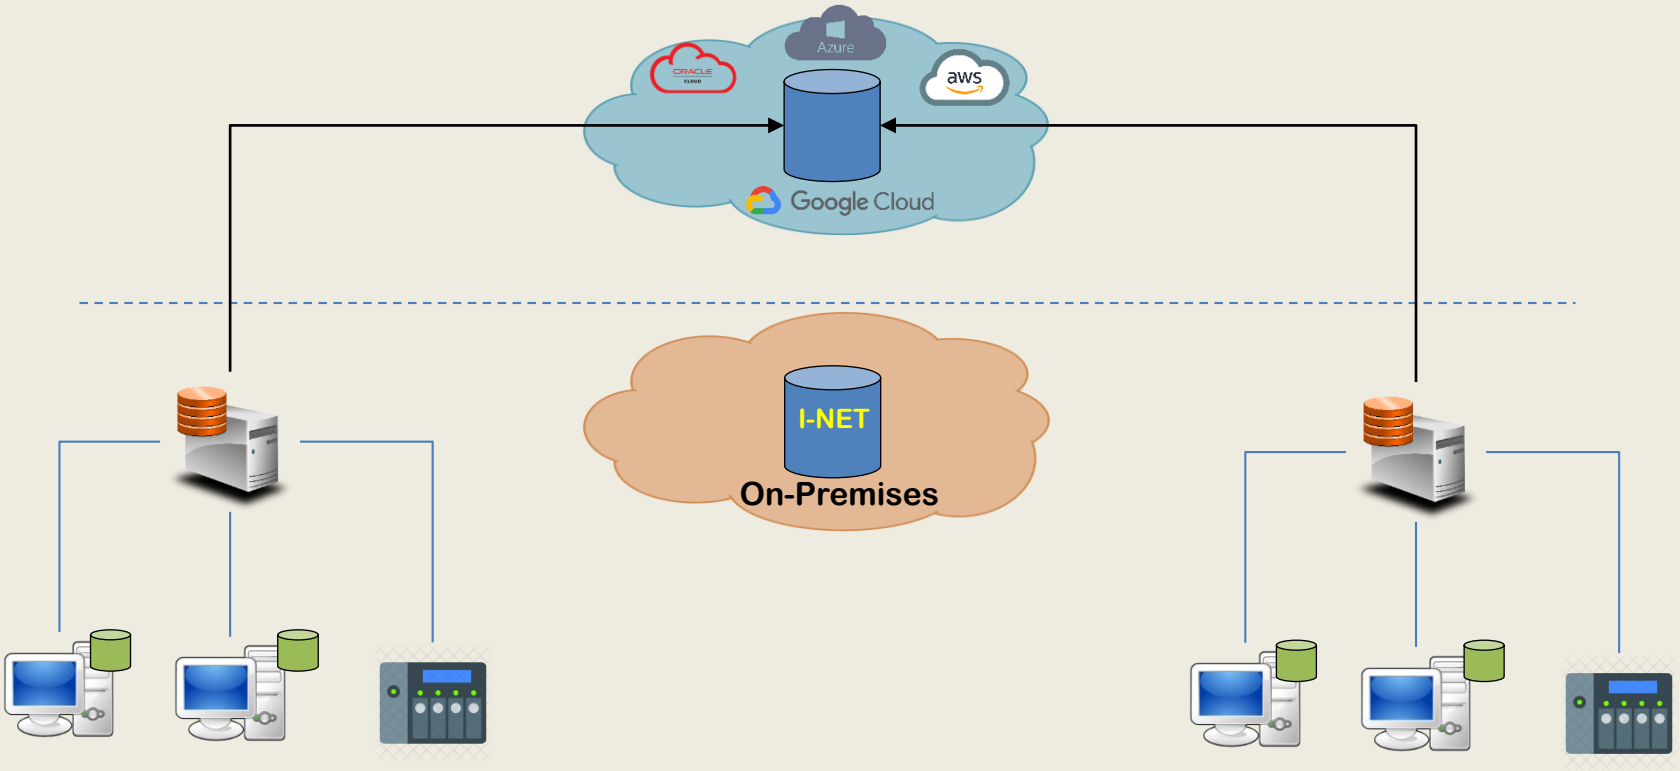

I-NET©
Draft and Copyright

Cloud Networking

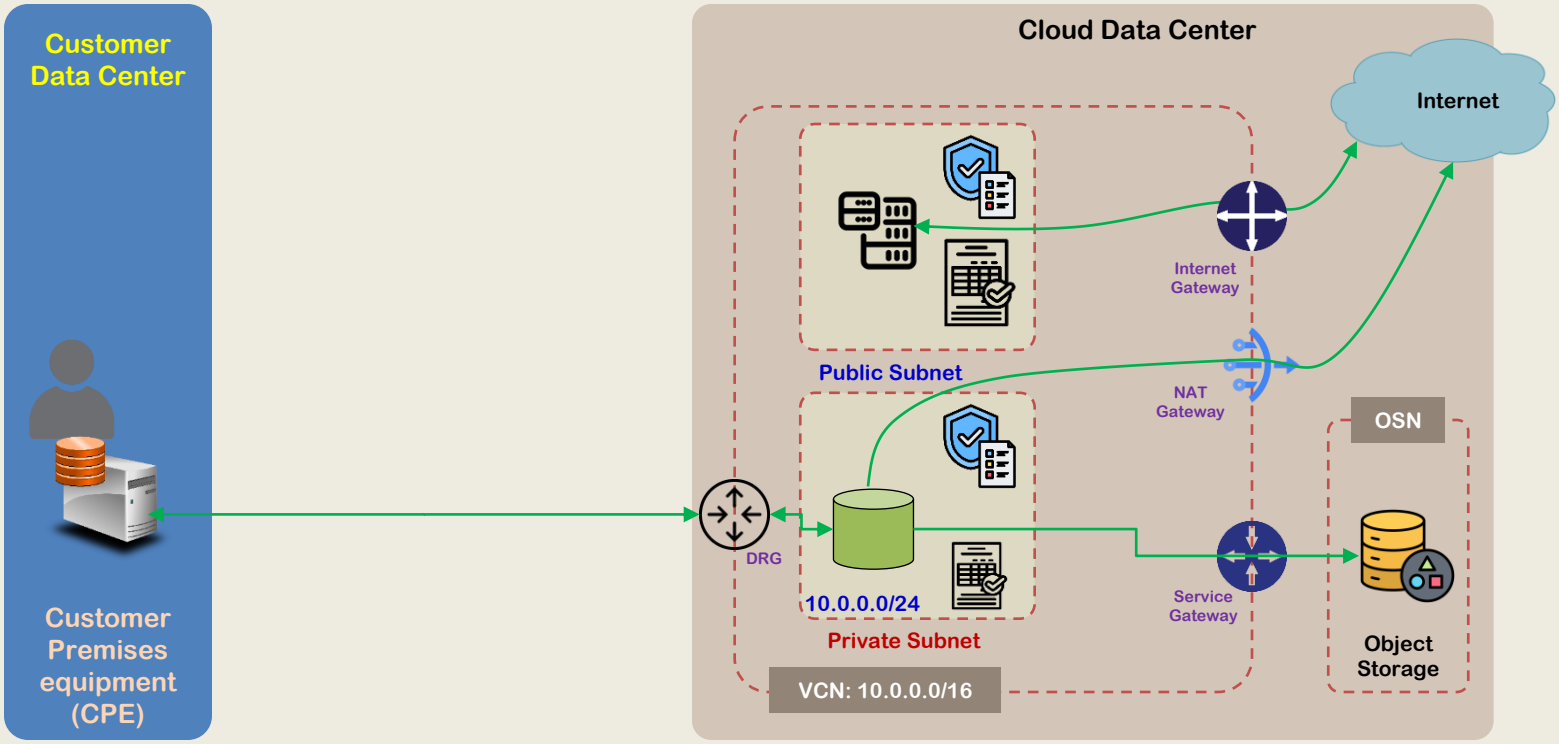

I-NET©
Draft and Copyright

Cloud Networking

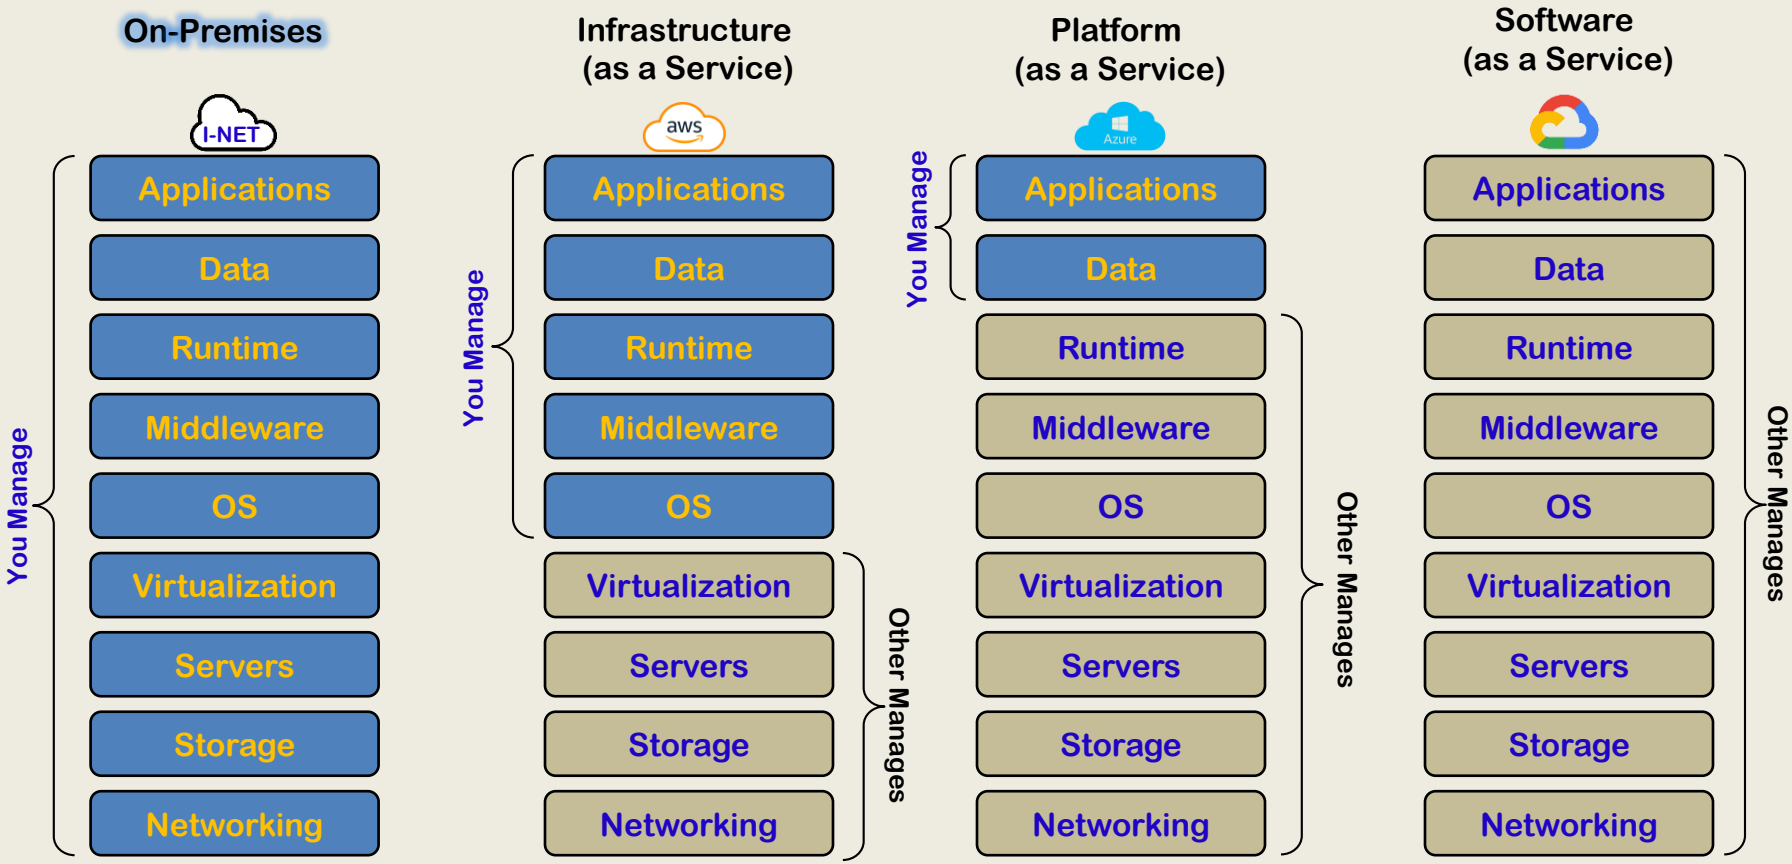

High Availability

I-NET ©
Draft and Copyright

What is Network Redundancy?

- **Network redundancy** is the practice of designing and configuring a network with;
  - backup components,
  - alternate data paths,
  - failover mechanisms

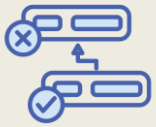

- so that connectivity and services remain available even if one part of the network fails.

- Instead of relying on;
  - a single device,
  - link,
  - provider,

- redundant networks distribute risk across multiple resources to maintain uptime and business continuity.

Different Types of Network Redundancy

- There are several types of network redundancy, each designed to protect against specific failure scenarios:
  - Link redundancy: Uses multiple physical or logical connections between devices to prevent outages caused by cable or interface failures.
  - Device redundancy: Deploys backup hardware such as switches, routers, or firewalls to maintain operations if a primary device fails.
  - Path redundancy: Creates multiple routing paths within the network so traffic can reroute automatically if a primary path becomes unavailable.
  - Power redundancy: Implements backup power sources such as dual power supplies, UPS systems, or generators to keep network components running during power loss.
  - ISP redundancy: Connects to multiple internet service providers to protect against external connectivity outages.
  - Geographic redundancy: Distributes infrastructure across different physical locations to reduce risks from environmental events or regional disruptions.

I-NET ©
Draft and Copyright

First Hop Redundancy Protocol

- A single point of failure (SPOF) is a component in the infrastructure that, if it fails, causes an outage for an entire section of the network. These points are undesirable in network design because they compromise the infrastructure's overall reliability, availability, and fault tolerance.
- Imagine a LAN network with multiple servers and switches. Everything inside the LAN is redundant. However, take a look at the components with red arrows. If any of these components (1-5) fail, the servers lose connectivity to the outside world.

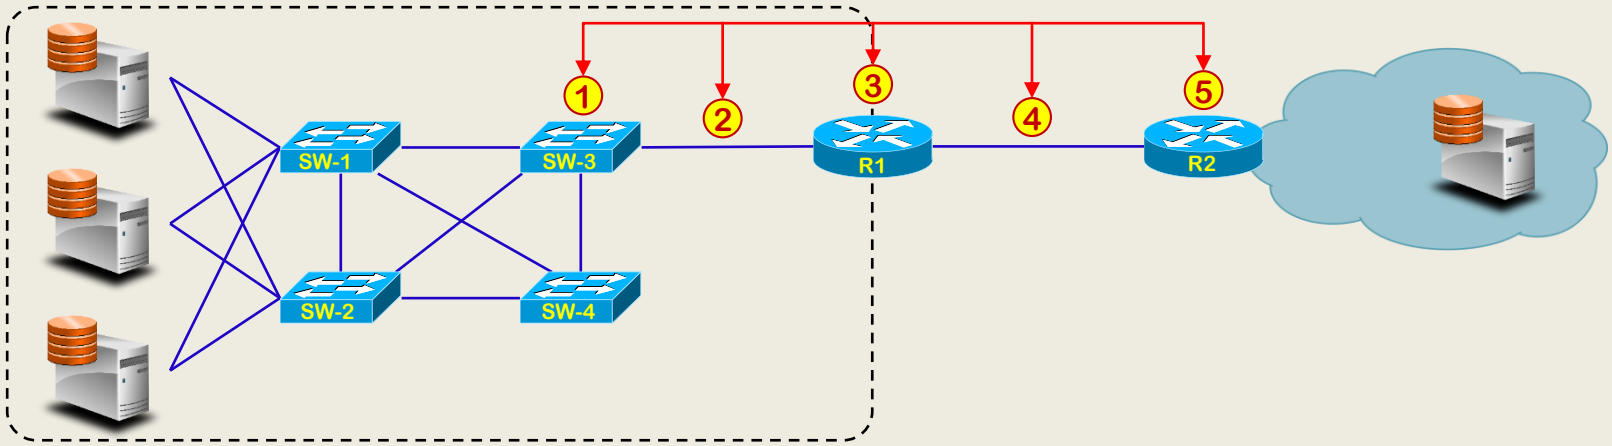

I-NET ©
Draft and Copyright

First Hop Redundancy Protocol

- The network must have redundancy at every component to be highly available and resilient. There are no single points of failure. Any element in the network, such as a power supply, cable, or optical module, might break, or even an entire device might lose power, and the network still works fine.

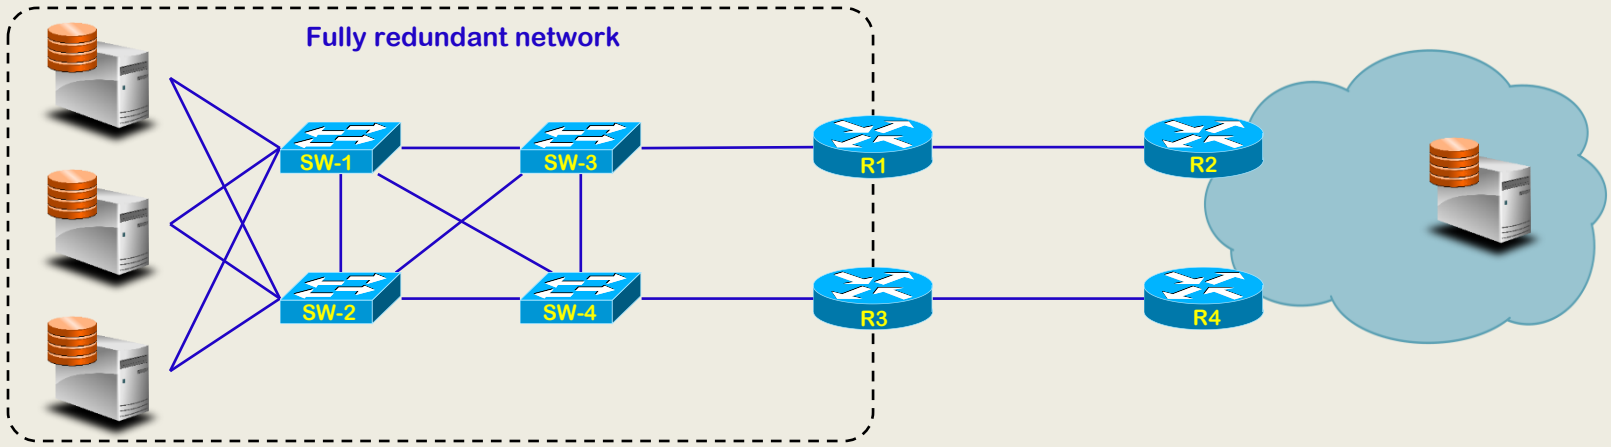

- Physical redundancy is very easy to understand. The network must have at least two of every physical component to be highly available; two power supplies, two switches, two cables, two WAN routers, two links, etc.

I-NET ©
Draft and Copyright

First Hop Redundancy Protocol

- Physical redundancy alone is not enough. The network needs **protocols** that can utilize the redundancy. If you have two WAN links but don't have a dynamic routing protocol, if one link goes down, the traffic might not be forwarded to the other link. Having multiple switches in the LAN is also not enough on its own. You must use the Spanning-Tree Protocol (STP) to calculate a loop-free topology and recalculate the new paths if one switch goes down.

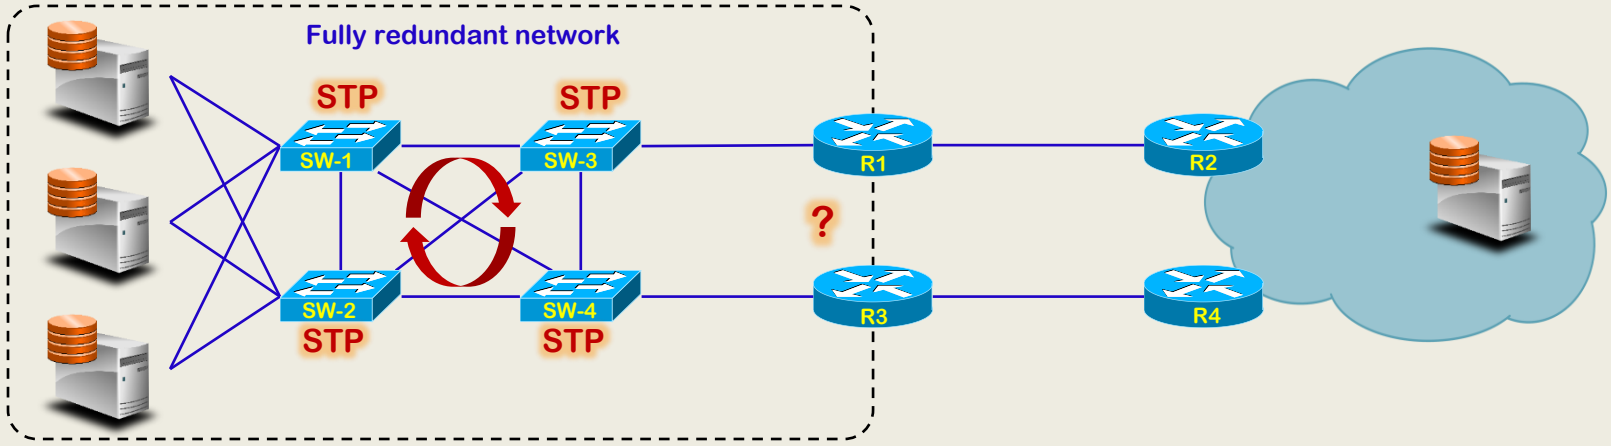

I-NET ©
Draft and Copyright

First Hop Redundancy Protocol

- The IPv4 protocol does not utilize redundancy by default. IPv4 focuses on delivering packets from a source to a destination but does not include built-in mechanisms for ensuring redundancy. IP hosts rely on a single configuration with one default gateway IP address that never changes.

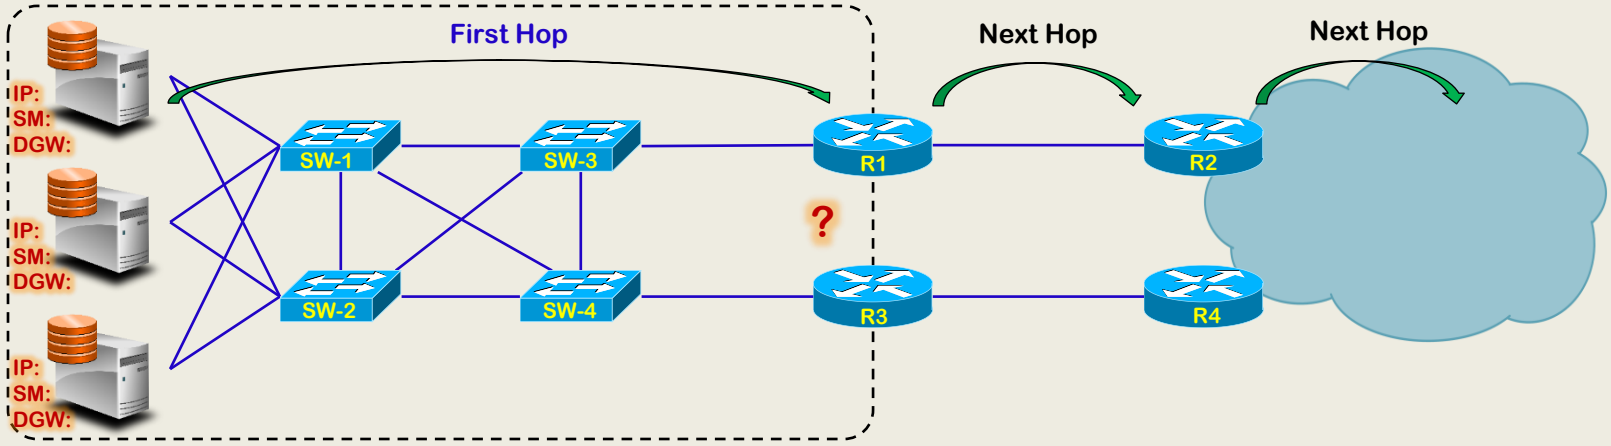

I-NET©
Draft and Copyright

First Hop Redundancy Protocol

- All operating systems, such as Windows, MacOS, and Linux, allow only one default gateway address to be configured on a host. Therefore, it is the network's responsibility to ensure that this default gateway address is highly available.
- First Hop Redundancy Protocol (FHRP) is a protocol that ensures high availability for hosts' default gateway. It allows multiple routers to appear as a single gateway address to hosts.

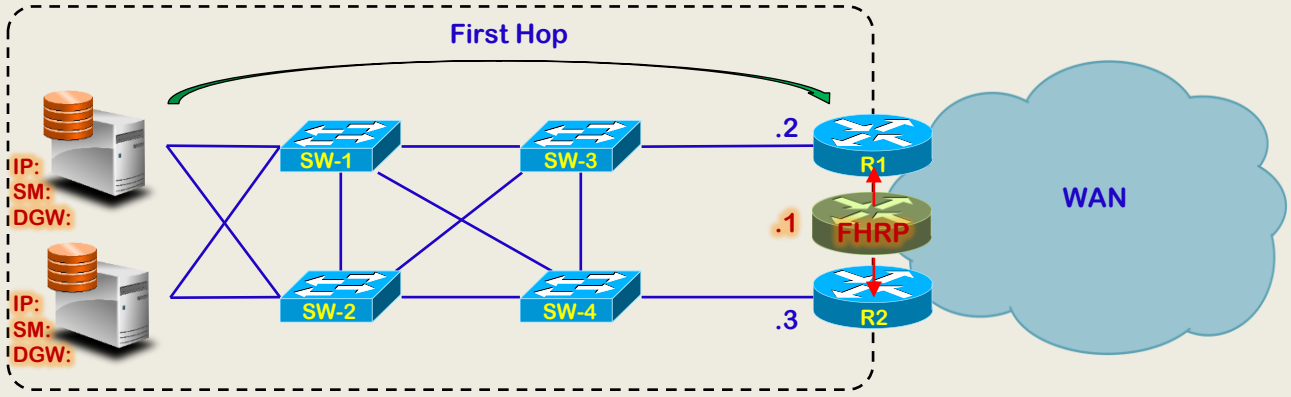

- FHRP protocols work by creating a group of routers that collectively act as a single virtual gateway for devices in the network. These routers share a virtual IP address (VIP) and a virtual MAC address.

I-NET ©
Draft and Copyright

Types of FHRP Protocols

- Hot Standby Router Protocol (HSRP)
  - Cisco's own proprietary protocol
  - Uses an active/standby router model
  - Provides virtual IP and MAC addresses for the router group
  - Supports authentication for security
- Virtual Router Redundancy Protocol (VRRP)
  - Industry-standard (RFC 5798)
  - Similar to HSRP but with some technical differences
  - Works across a few different vendors' devices
  - Has slightly faster convergence than HSRP

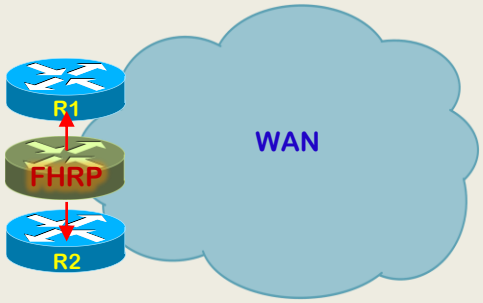

- Gateway Load Balancing Protocol (GLBP)
  - Another Cisco proprietary protocol
  - Provides load balancing across multiple routers
  - Can distribute traffic using round-robin, weighted, or host-dependent methods
  - Uses resources more efficiently compared to HSRP/VRRP

I-NET ©
Draft and Copyright

Stateful Switchover (SSO)

- Stateful Switchover (SSO) is a high-availability technique that synchronizes state information (sessions, configurations) from an active to a standby controller or processor, enabling seamless failover.
- It ensures that if the active device fails, the standby takes over without disrupting user connectivity, typically used alongside Non-Stop Forwarding (NSF) to minimize downtime. Key High Availability SSO Techniques:
  - AP and Client SSO (Wireless): On Cisco WLCs, this mirrors the access point (AP) database and client information to a standby unit. When the active WLC fails, the standby takes over without forcing APs to re-discover or clients to re-authenticate, resulting in zero downtime.
  - Stateful Switchover (Routing/Control Plane): In routers with dual Route Processors (RPs), the active RP synchronizes protocol state information (e.g., routing tables, line card info) with the standby. This prevents sessions from resetting during a failover.
  - Redundancy Management Interface (RMI): In Cisco Catalyst 9800 WLCs, RMI is used to monitor the peer and ensure that the standby unit is in a "Hot Standby" state, allowing it to take over instantaneously.

I-NET ©
Draft and Copyright

Stateful Switchover (SSO)

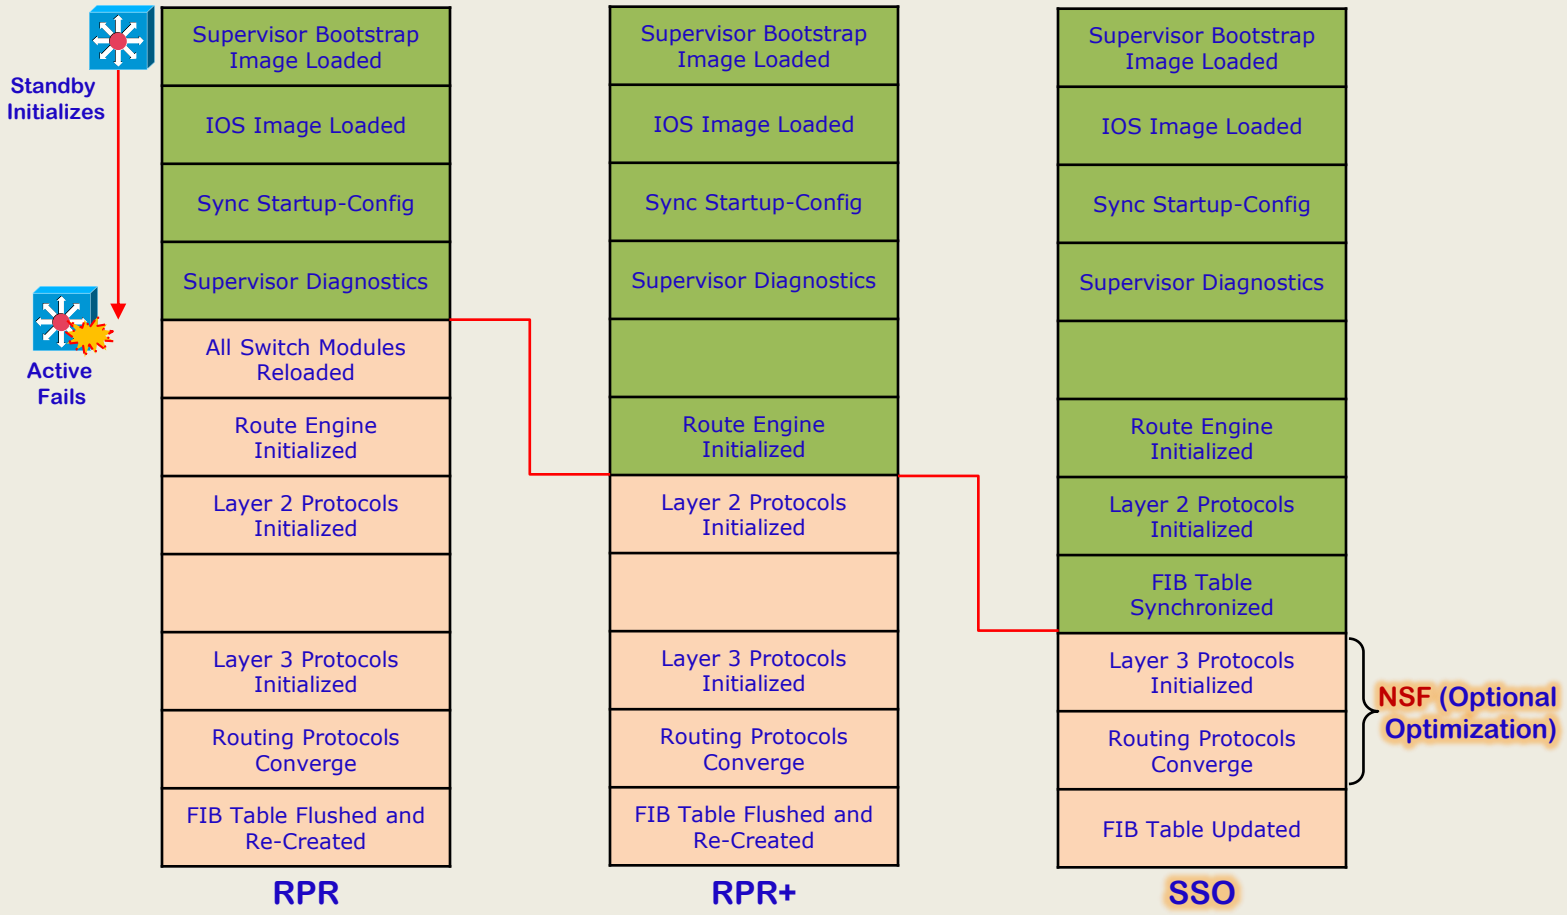

I-NET ©
Draft and Copyright

Stateful Switchover (SSO)

- Nonstop forwarding (NSF) is an interactive method that focuses on quickly rebuilding the Routing Information Base (RIB) table after a supervisor switchover.
- The RIB is used to generate the Forwarding Information Base table for CEF, which is downloaded to any switch modules or hardware that can perform Cisco Express Forwarding.
- NSF is supported by the BGP, EIGRP, OSPF, and IS-IS routing protocols.

Routing Protocol	Configuration Commands
BGP	Switch(config)#router bgp as-number Switch(config-router)#bgp graceful-restart
EIGRP	Switch(config)#router eigrp as-number Switch(config-router)#nsf
OSPF	Switch(config)#router ospf process-id Switch(config-router)#nsf
IS-IS	Switch(config)#router isis [tag] Switch(config-router)#nsf [cisco | ietf] Switch(config-router)#nsf interval [minutes] Switch(config-router)#nsf t3 {manual [seconds] | adjacency } Switch(config-router)#nsf interface wait seconds

Depends on IOS

I-NET©
Draft and Copyright

QUESTIONS?

Prepared by Ayhan KOŞAR
ayhankosar@gmail.com
ayhan.k@inetworkeskisehir.com
www.inetworkeskisehir.com

linkedin.com/in/ayhankosar