

Network Design Requirements: Analysis and Design Principles

Designing large-scale networks to meet today's dynamic business and IT needs and trends is a complex assignment, whether it is an enterprise or service provider type of network. This is especially true when the network was designed for technologies and requirements relevant years ago and the business decides to adopt new IT technologies to facilitate the achievement of its goals but the business's existing network was not designed to address these new technologies' requirements. Therefore, to achieve the desired goal of a given design, the network designer must adopt an approach that tackles the design in a structured manner.

There are two common approaches to analyze and design networks:

- **The top-down approach:** The top-down design approach simplifies the design process by splitting the design tasks to make it more focused on the design scope and performed in a more controlled manner, which can ultimately help network designers to view network design solutions from a business-driven approach.
- **The bottom-up approach:** In contrast, the bottom-up approach focuses on selecting network technologies and design models first. This can impose a high potential for design failures, because the network will not meet the business or applications' requirements.

To achieve a successful strategic design, there must be additional emphasis on a business driven approach. This implies a primary focus on business goals and technical objectives, in addition to existing and future services and applications. In fact, in today's networks, business requirements are driving IT and network initiatives as shown in Figure 1-1 [6].

For instance, although compliance (as presented in Figure 1-1) might seem to be a design constraint rather than a driver, many organizations today aim to comply with some standards with regard to their IT infrastructure and services to gain some business advantages, such as compliance with ISO/IEC 27001 Information Security Management,¹ will

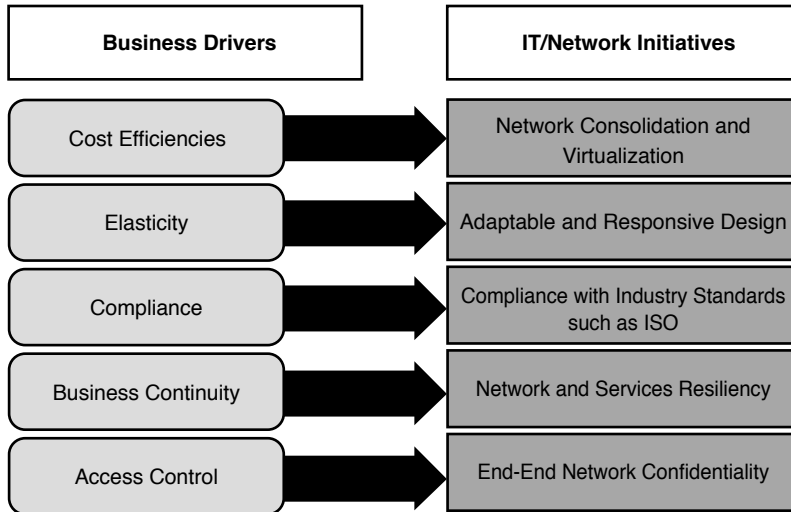


Figure 1-1 *Business Drivers Versus IT Initiatives*

help businesses like financial services organizations to demonstrate their credibility and trust. This ultimately will help these organizations to gain more competitive advantages, optimize their operational uptime, and reduce operational expenses (fewer number of incidents as a result of the reduced number of information security breaches).

Throughout this book and for the purpose of the CCDE exam, the top-down approach is considered as the design approach that can employ the following top-down logic combined with the prepare, plan, design, implement, operate and optimize (PPDIOO) lifecycle:

- Analyze the goals, plans, and requirements of the business.
- Define application requirements from the upper layers of the Open Systems Interconnection (OSI) reference model that can help to identify the characteristics of an application.
- Specify the design of the infrastructure along with the functional requirements of its components, for the network to become a business enabler.
- Monitor and gather additional information that may help to optimize and influence the logical or physical design to adapt with any new application or requirements.

Design Scope

It is important in any design project that network designers carefully analyze and evaluate the scope of the design before starting to gather information and plan network design. Therefore, it is critical to determine whether the design task is for a green field (new) network or for a current production network (if the network already exists, the

design tasks can vary such as optimization, expansion, integration with other external networks, and so on). It is also vital to determine whether the design spans a single network module or multiple modules. In other words, the predetermination of the design scope can influence the type of information required to be gathered, in addition to the time to produce the design. Table 1-1 shows an example of how identifying the design scope can help network designers determine the areas and functions a certain design must emphasize and address. As a result, the scope of the information to be obtained will more be focused on these areas.

Table 1-1 *Design Scope*

Design Scope	Detailed Design Scope Example
Enterprise campus network and remote sites	Rollout of IP telephony across the enterprise, which may require a redesign of virtual LANs (VLANs), quality of service (QoS), and so on across the LAN, WAN, data center (DC), and remote-access edge
Campus only	Add multi-tenancy concept to the campus, which requires design of VLANs, IPs, and path isolation across the campus LAN only
Optimize enterprise edge availability	Add redundant link for remote access, which might require redesign of the WAN module and remote site designs and configurations such as overlay tunnels

Note Identifying the design scope in the CCDE exam is very important. For example, the candidate might have a large network to deal with, whereas the actual design focus is only on adding and integrating a new data center. Therefore, the candidate needs to focus on that part only. However, the design still needs to consider the network as a whole, a “holistic approach,” when you add, remove, or change anything across the network (as discussed in more detail later in this chapter).

Business Requirements

This section covers the primary aspects that pertain to the business drivers, needs, and directions that (individually or collectively) can influence design decisions either directly or indirectly. The best place to start understanding the business’s needs and requirements is by looking at the big picture of a company or business and understanding its goals, vision, and future directions. This can significantly help to steer the design to be more business driven. However, there can be various business drivers and requirements based on the business type and many other variables. As outlined in Figure 1-2, with a top-down design approach, it is almost always the requirements and drivers at higher layers (such as business and application requirements) that drive and set the requirements and directions for the lower layers. Therefore, network designers aiming to achieve a business-driven design must consider this when planning and producing a new network

design or when evaluating and optimizing an existing one. The following sections discuss some of the business requirements and drivers at the higher layers and how each can influence design decisions at the lower layers.

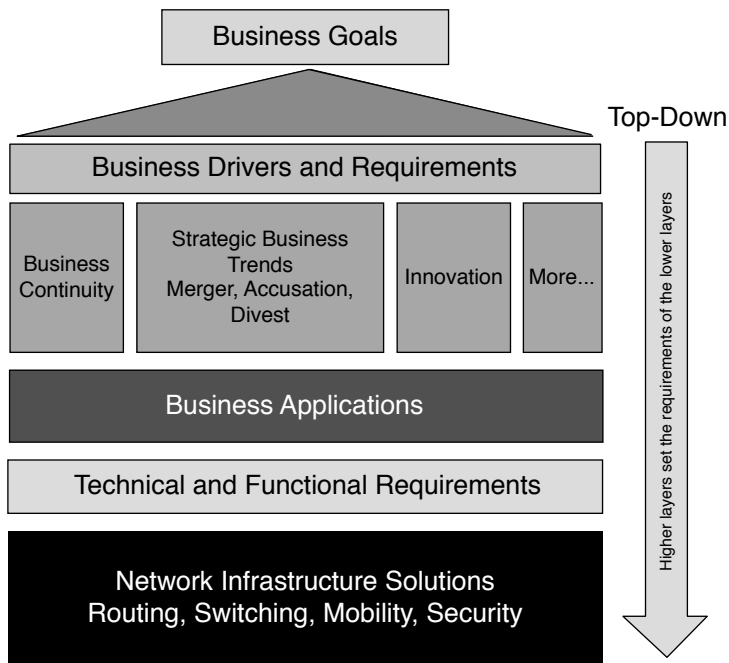


Figure 1-2 *Business-Driven Technology Solutions*

Business Continuity

Business continuity (BC) refers to the ability to continue business activities (business as usual) following an outage, which might result from a system outage or a natural disaster like a fire that damages a data center. Therefore, businesses need a mechanism or approach to build and improve the level of resiliency to react and recover from unplanned outages.

The level of resiliency is not necessarily required to be the same across the entire network, however, because the drivers of BC for the different parts of the network can vary based on different levels of impact on the business. These business drivers may include compliance with regulations or the level of criticality to the business in case of any system or site connectivity outage. For instance, if a retail business has an outage in one of its remote stores, this is of less concern than an outage to the primary data center, from a business point of view. If the primary data center were to go offline for a certain period of time, this would affect all the other stores (higher risk) and could cost the business a larger loss in terms of money (tangible) and reputation (intangible). Therefore, the resiliency of the data center network is of greater consideration for this retailer than the resiliency of remote sites [17].

Similarly, the location of the outage sometimes influences the level of criticality and design consideration. Using the same example, an outage at one of the small stores in a remote area might not be as critical as an outage in one of the large stores in a large city [11]. In other words, BC considerations based on risk assessment and its impact on the business can be considered one of the primary drivers for many businesses to adapt network technologies and design principles to meet their desired goals [5].

Elasticity to Support the Strategic Business Trends

Elasticity refers to the level of flexibility a certain design can provide in response to business changes. A change here refers to the direction the business is heading, which can take different forms. For example, this change may be a typical organic business growth, a decline in business, a merger, or an acquisition. For instance, if an enterprise campus has three buildings and is interconnected directly, as illustrated in Figure 1-3, any organic growth in this network that requires the addition of a new building to this network will introduce a lot of complexity in terms cabling, control plane, and manageability. These complexities result from the inflexible design, which makes the design incapable of responding to the business's natural growth demand.




Figure 1-3 *Inflexible Design*

To enhance the level of flexibility of this design, you can add a core module to optimize the overall design modularity to support business expansion requirements. As a result, adding or removing any module or building to this network will not affect other modules, and does not even require any change to the other modules, as illustrated in Figure 1-4. In other words, the design must be flexible enough to support the business requirements and strategic goals. If network designers understand business trends and directions in this area, such understanding may influence, to a large extent, design choices.

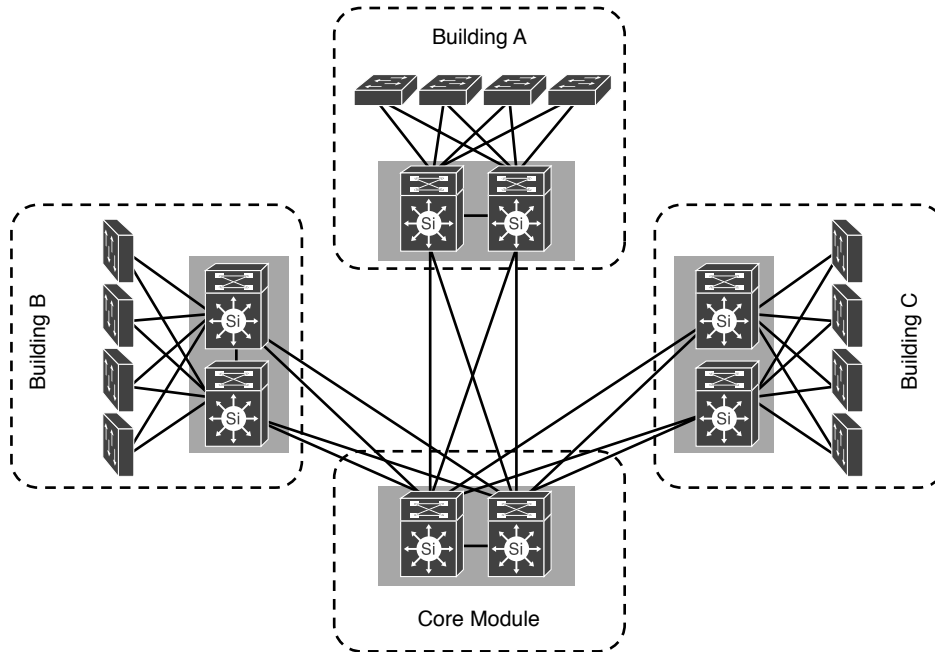


Figure 1-4 *Flexible Design*

Similarly, a flexible network design must support the capability to integrate with other networks (for examples, when mergers and acquisitions occur). With mergers and acquisitions, however, the network can typically grow significantly in size within a short period of time, and the biggest challenge, in both scenarios (mergers and acquisitions), is that network designers have to deal with different design principles, the possibility of overlapping IP address space, different control plane protocols, different approaches, and so on.

IT as a “Business Innovation” Enabler

In today’s market, many businesses understand how IT technologies enhance their services and provide innovation to their customers. Therefore, when a certain technology can serve as a business enabler that can help the organization to compete in the market or increase its customers’ satisfaction, the adoption of that technology will be supported by the business [17].

For example, nowadays, cloud-based data centers are opening new opportunities for hosting service providers to generate more revenue for the business. To offer good cloud-based services, there must be a reliable, flexible, and high-performance data center infrastructure to deliver this service. Consequently, this engenders the initiative and will drive the business to build a high-performance, next-generation data center network.

This network, by acting as a basis for cloud services, will be the enabler of the business's revenue-generation solution.

The Nature of the Business

Classifying the industry in which the business belongs or identifying the business's origins can aid in the understanding of some indirect requirements, even if these are not mentioned explicitly. For example, information security is almost always a must for a banking business whenever traffic crosses any external link. So by default, when planning a design for a business based in the banking industry, the design must support or offer security capabilities to gain acceptance from the business. In addition, industry-specific standards apply to IT infrastructure and services need to be considered. (For instance, healthcare organizations may consider complying with the IEC-80001-1 standard.²⁾)

Business Priorities

Each business has a set of priorities that are typically based on strategies adopted for the achievement of goals. These business priorities can influence the planning and design of IT network infrastructure. Therefore, network designers must be aware of these business priorities to align them with the design priorities. This ensures the success of the network they are designing by delivering business value. For example, company X's highest priority is to provide a more collaborative and interactive business communication, followed by the provision of mobile access for the users. In this example, providing a collaborative and interactive communication must be satisfied first before providing or extending the communication over any mobility solution for the end users. In sum, it is important to align the design with the business priorities, which are key to achieving business success and transforming IT into a business enabler.

Functional Requirements

Functional requirements compose the foundation of any system design because they define system and technology functions. Specifically, functional requirements identify what these technologies or systems will deliver to the business from a technological point of view. For example, a Multiprotocol Label Switching (MPLS)-enabled service provider might explicitly specify a functional requirement in a statement like this: "The provider edge routers must send VoIP traffic over 10G fiber link while data traffic is to be sent over the OC-48 link." It is implied that this service provider network needs to have provider edge (PE) routers that support a mechanism capable of sending different types of traffic over different paths, such as MPLS Traffic Engineering (MPLS-TE). Therefore, the functional requirements are sometimes referred to as *behavioral requirements* because they address what a system does.

Note The design that does not address the business’s functional requirements is considered a poor design; however, in real-world design, not all the functional requirements are provided to the designer directly. Sometimes they can be decided on indirectly, based on other factors. See the “Application Requirements” section later in this chapter for more details.

Technical Requirements

The technical requirements of a network can be understood as the technical aspects that a network infrastructure must provide in terms of security, availability, and integration. These requirements are often called *nonfunctional requirements*. Technical requirements vary, and they must be used to justify a technology selection. In addition, technical requirements are considered the most dynamic type of requirements compared to other requirements such as business requirements because, based on technology changes, they change often. Technical requirements include the following:

- Heightened levels of network availability (for example, using First Hop Redundancy Protocol [FHRP])
- Support the integration with network tools and services (for example, NetFlow Collector, or authentication and authorization system “RADIUS servers”)
- Cater for network infrastructure security techniques (for example, control plane protection mechanisms or infrastructure access control lists [iACLs])

Note The technical requirements help network designers to specify the required technical specifications (features and protocols) and software version that supports these specifications and sometimes influence the hardware platform selection based on its technical characteristics.

Application Requirements

From a business point of view, user experience is one of the primary, if not the highest, priority that any IT and network design must satisfy. The term *end users* can be understood differently according to the type of business. The following are the most common categories of end users:

- **Customers:** Customer can be individuals, such as a customer of a bank, or they can be a collective, such as the customers of an MPLS service provider. From a business point of view, customer satisfaction can directly impact the business’s reputation and revenue.

- **Internal users:** In this category, the targeted users are internal users. Productivity of these users can translate to business performance efficiency, which has a direct relation to business success and revenue.
- **Business partners:** Partners represent those entities or organizations that work together to achieve certain goals. Therefore, efficient interaction between partners can enhance their business success in the service of strategic goal achievement.

Therefore, a network or a technology that cannot deliver the desired level of the users' expectation (also known as *quality of experience*) means a failure to achieve either one of the primary business goals or failure to satisfy a primary influencer of business success. Consequently, networks and IT technologies will be seen by the business as a cost center rather than as a business enabler.

On this basis, networks design must take into account how to deliver the desired level of experience. In fact, what influences users' experience is what they see and use. In other words, from a user's perspective, the quality of experience with regard to applications and services used by different types of users is a key deterministic factor.

Deploying network applications or services without considering the characteristics and network requirements of these applications will probably lead to a failure in meeting business goals. For example, a company providing modern financial services wants to distinguish itself from other competitors by enabling video communication with its customer service call center agents. If the network team did not properly consider video communication requirements as a network application, the application will probably fail to deliver the desired experience for the end users (customers in this example). Consequently, this will lead to a failure in achieving the company's primary business goal. In other words, if any given application is not delivered with the desired quality of experience to the end users, the network can simply be considered as not doing its job.

Furthermore, in some situations, application requirements can drive functional requirements. For instance, if a service provider has a service level agreement (SLA) with its clients to deliver Voice over IP (VoIP) traffic with less than 150 ms of one-way delay and less than 1 percent of packet loss, here VoIP requirements act as application requirements, which will drive the functional requirements of the PE devices to use a technology that can deliver the SLA. This technology may include, for example, a Class-Based Tunnel Selection (CBTS) MPLS-TE protected with fast reroute (FRR) to send VoIP traffic over high-speed links and provide network path protection in case of a failure within 50 ms. In addition, network designers should also consider the answers to the following fundamental questions when evaluating application requirements:

- How much network traffic does the application require?
- What is the level of criticality of this application and the service level requirement to the business?
- Does the application have any separation requirements to meet industry regulations and corporate security policies?

- What are the characteristics of the application?
- How long does the application need after losing connectivity to reset its state or session?

Design Constraints

Constraints are factors or decisions already in place and cannot be changed and often lead to a direct or indirect impact on the overall design and its functional requirements. In reality, various design constraints affect network design. The following are the most common constraints that network designers must consider:

- **Cost:** Cost is one of the most common limiting factors when producing any design; however, for the purpose of the CCDE practical exam, cost should be considered as a design constraint only if it is mentioned in the scenario as a factor to be considered or a tiebreaker between two analogous solutions.
- **Time:** Time can also be a critical constraint when selecting a technology or architecture over another if there is a time frame to complete the project, for example.
- **Location:** Location is one of the tricky types of constraints because it can introduce limitations that indirectly affect the design. For instance, a remote site may be located in an area where no fiber infrastructure is available and the only available type of connectivity is over wireless. From a high-level architectural point of view, this might not be a serious issue. From a performance point of view, however, this might lead to a reduced link speed, which will affect some sensitive applications used by that site.
- **Infrastructure equipment:** A good example here is that of legacy network devices. If a business has no plan to replace these devices, this can introduce limitations to the design, especially if new features or protocols not supported by these legacy platforms are required to optimize the design.
- **Staff expertise:** Sometimes network designers might propose the best design with the latest technologies in the market, which can help reduce the business's total cost of ownership (TCO) (for example, in the case of virtualization in the data center and the consolidation of data and Fibre Channel [FC] storage over one infrastructure Fibre Channel over Ethernet [FCoE]). This can be an issue, however, if the staff of this company has no expertise in these technologies used to operate and maintain the network. In this case, you have two possible options (with some limitations applying to each, as well):
 - **Train the staff on these new technologies:** This will be associated with a risk, because as a result of the staff's lack of experience, they may take a longer time to fix any issue that might occur, and at the same time, data center downtime can cost the business a large amount of money.
 - **Hire staff with experience in these technologies:** Normally, people with this level of expertise are expensive. Consequently, the increased operational cost might not justify the reduced TCO.

Note In some situations, if the proposed solution and technologies will save the business a significant amount of money, you can justify the cost of hiring new staff.

Crafting the Design Requirements

This section demonstrates how different types of requirements collectively can lead to the achievement of the desired network design, which ultimately will facilitate achieving business goals. This demonstration is based on an example that goes through the flow of information gathering (the top-down approach, as illustrated in Figure 1-5) to build up the requirements starting from the business goals (top) to the technical requirement level, such as the required features (bottom).

A national retail company is currently expanding their business to add several international sites within the next 12 months. However, with their current IT infrastructure, they face a high number of expenses in managing and maintaining their two current data and voice networks. In addition, the business wants to invest in technologies that offer enhancements to business activities and increase employee productivity.

The following is a typical classification of the requirements, some of which might be provided directly and some of which can be implied or indicated based on other requirements and goals:

- **Business goals:**
 - Reduce operational cost
 - Enhance employees' productivity
 - Expand the business (adding more remote sites)
- **Business requirements:**
 - Reduce the cost of maintaining multiple networks for voice and data
 - Improve employee productivity by enhancing and integrating internal communications through video and mobile devices, without compromising the company's security policy
 - Support the business expansion (the rollout of the new remote sites)
- **Functional requirements:**
 - A unified infrastructure that supports voice, video, data, and wireless
 - Ability to provide isolation between the traffic of guests and internal staff (for both wired and wireless) to comply with the standard security policy of the organization
 - Capability to support introducing new remote sites to the network without any redesign

- **Application requirements:** In this particular example, we assume that no specific application requirements were given. In fact, they do not need to be provided because it is obvious that this organization is going to add VoIP and video as applications or services. The network must provide the required level of network efficiency, performance, and availability to meet VoIP/video requirements (real-time, delay, and jitter sensitive) to achieve one of the business goals.
- **Technical requirements:** To achieve the above network's functional and application requirements considering the ultimate business goals, the design must cater to the following:
 - **High availability:** Increase high availability to support critical traffic such as voice (for example, redundant nodes, control plane tuning, FHRP)
 - **Quality of service:** Improve users' quality of experience by introducing QoS to achieve high-quality voice and video services (for example, queuing, traffic shaping)
 - **Security:** Optimize the network security to accommodate the new design, including voice and wireless security (for example, access control with 802.1X authentication)
 - **Traffic isolation:** Provide traffic isolation to meet the information security requirements (for example, tunneling such as generic routing encapsulation [GRE] with MPLS/VRFs [virtual private network routers/forwarders])
 - **Scalability:** Scalable network (WAN) design to support the projected business growth during the next 12 months (for example dynamic multipoint VPN [DMVPN] versus point-to-point GRE)

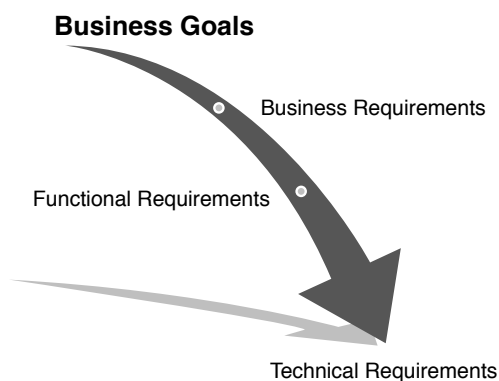


Figure 1-5 *Design Requirements Flow*

Table 1-2 summarizes the mapping between the technical requirements, business priorities, and technical implementation planning priorities.

Table 1-2 *Design Requirements Mapping*

Technical Requirements	Features/Tools to Enable the Technical Requirement	Business Priority Level, 1–5 (1 = highest, 5 = lowest)	Implementation Priorities
High availability	Modular and hierarchal network design. Network node and path redundancy. Tune control plane for fast convergence.	1	1
QoS	Traffic marking, classification, and prioritization to provide QoS.	2	3
Security	Add stateful firewalls with filtering to control communications between networks.	3	2
Traffic isolation	Use network virtualization techniques.	3	2
Scalability	WAN design must be scalable enough to support the business expansion such avoiding point-to-point mesh types of connectivity over the WAN.	1	1

Note that in Table 1-2 there are two levels of priorities:

- **Business priority level:** This level reflects the priority from business point of view.
- **Implementation priority level:** This level reflects the priorities from a technical implementation point of view. Although business priorities must always be met first, sometimes from a technical point of view there may be prerequisites to enable services or features across the network first to achieve a specific requirement. In the preceding example, from a technical design and implementation point of view, the network first needs to be designed and connected in the desired way (for example, hierarchical LAN, hub-and-spoke WAN). Following, the network designer can apply virtualization techniques to meet traffic separation requirements. At the same time, security appliances and functions must be integrated (holistic approach not siloed approach). Finally, QoS can be applied later to optimize traffic flows over both the physical and logical (virtualized) networks. In other words, the described sequence of the implementation plan in the preceding example was driven by both the business and technical priorities, considering that the business priorities must always be given the preference when technically possible.

Planning

There is an enduring adage: “If you do not have a plan, you are planning to fail.” This adage is accurate and applicable to network design. Many network designers focus on implementation after obtaining the requirements and putting them in a design format. They sometimes rely on the best practices of network design rather than focusing on planning: “What are the possible ways of getting from point A to point B?” This planning process can help the designer devise multiple approaches or paths (design options). At this point, the designer can ask the question: Why? Asking *why* is vital to making a business-driven decision for the solution or design that optimally aligns with the business’s short- or long-term strategy or objective. In fact, the best practices of network design are always recommended and should be followed whenever applicable and possible.

However, reliance on best practices is more common when designing a network from scratch (greenfield), which is not common with large enterprises and service provider networks.

In fact, IT network architectures and building constrictions and architectures are quite similar in the way they are approached and planned by designers or consultants.

For example, five years ago, an investment company built a shopping mall in one of the large cities in Europe, which was architected and engineered based on the requirements at that time. Recently, the stakeholders requested the design to be reviewed and optimized to overcome the increased number of people visiting this shopping mall (tourist and locals), because this increase was not properly projected and planned for during the original design five years ago.

Typically, the architects and engineers will then evaluate the situation, identify current issues, and understand the stakeholders’ goals. In other words, they gather business requirements and identify the issues. Next, they work on optimizing the existing building (this may entail adding more parking space, expanding some areas, and so forth) rather than destroying the current building and rebuilding it from scratch. However, this time they need to have proper planning to provide a design that fits current and future needs.

Similarly, with IT network infrastructure design, there are always new technologies or broken designs that were not planned well to scale or adapt to business and technology changes. Therefore, network designers must analyze business issues, requirements, and the current design to plan and develop a solution that can optimize the overall existing architecture. This optimization might involve the redesign of some parts of the network (for example, WAN), or it might involve adding a new data center to optimize BC plans. To select the right design option and technologies, network designers need to have a *planning approach* to connect the dots at this stage and make a design decision based on the information gathering and analysis stage. Ultimately, the planning approach leads to the linkage of design options and technologies to the gathered requirements and goals to ensure that the design will bring value and become a business enabler rather than a cost center to the business. Typical tools network designers use at this stage to facilitate and simplify the selection process are the decision tree and the decision matrix.

Decision Tree

The decision tree is a helpful tool that a network designer can use when needing to compare multiple design options, or perhaps protocols, based on specific criteria. For example, a designer might need to decide which routing protocol to use based on a certain topology or scenario, as illustrated in Figure 1-6.

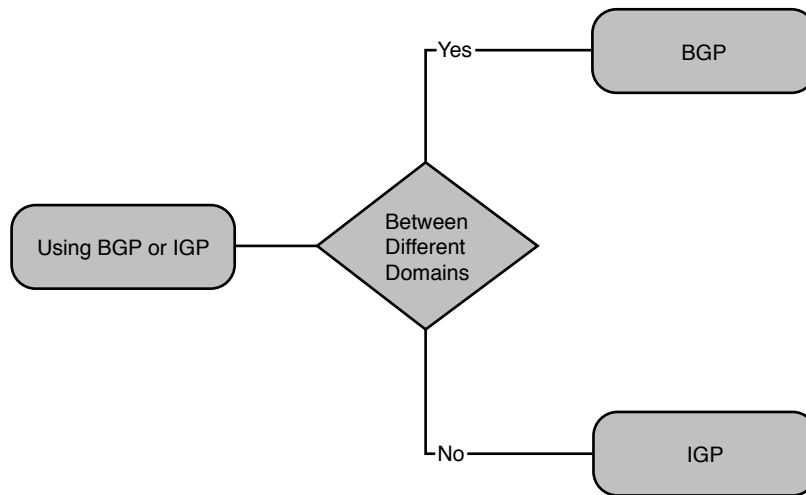


Figure 1-6 *Decision Tree*

Decision Matrix

Decision matrixes serve the same purpose as decision trees; however, with the matrix, network designers can add more dimensions to the decision-making process. In Table 1-3, there are two dimensions a network designer can use to select the most suitable design option. In these two dimensions, both business requirements and priorities can be taken into account to reach the final decision, which is based on a multidimensional approach.

Table 1-3 *Decision Matrix*

Business Requirements	Priority	Design Option 1	Design Option 2
Cost savings	4	Moderate	Low
Scalable	1	High	High
Secure	2	Low	High

When using the decision matrix as a tool in the preceding example, design option 2 is more suitable based on the business requirements and priorities. The decision matrix is not solely reliant on the business requirements to drive the design decision; however,

priorities from the business point of view were considered as an additional dimension in the decision-making process, which makes it more relevant and focused.

Planning Approaches

To develop a successful network design, a proper planning approach is required to build a coherent strategy for the overall design. Network designers can follow two common planning approaches to develop business-driven network designs and facilitate the design decisions:

- **Strategic planning approach:** Typically targets planning to long-term goals and strategies. For example, a service provider needs to migrate its core from legacy (ATM) to be based on MPLS instead. The design decision in this approach has to cater to long-term goals and plans.
- **Tactical planning approach:** Typically targets planning to overcome an issue or to achieve a short-term goal. For instance, an enterprise might need to provide remote access to some partners temporally (for example, over the Internet) until dedicated extranet connectivity is provisioned. Design decisions in this approach generally need to provide what is required within a limited period of time and are not necessarily required to consider long-term goals.

Strategic Balance

Within any organization, there are typically multiple business units and departments. Each has its own strategy, some of which are financially driven, whereas others are more innovation driven. For example, an IT department is more of an in-house service provider concerned with ensuring service delivery is possible and optimal, whereas the procurement department is cost driven and always prefers the cheaper options. The marketing department, in contrast, is almost always innovation driven and wants the latest technology. Consequently, a good understanding of the overall business strategy and goals can lead to compromise between the different aims of the different departments. In other words, the idea is that each business unit or entity within an organization must have its requirements met at a certain level, so that all can collectively serve the overall business goals and strategies. For instance, let's consider a case study of a retail business wanting to expand its geographic presence by adding more retail shops across the globe with low capital expenditure (capex). Based on this goal, the main point is to increase the number of remote sites with minimal cost (expansion and cost):

- **IT point of view:**
 - The point of sales application used does not support offline selling or local data saving. Therefore, it requires connectivity to the data center to operate.
 - The required traffic volume from each remote site is small and non-real time.
 - Many sites are to be added within a short period of time.
 - **Optimum solution:** IT suggested that the most scalable and reliable option is to use a MPLS VPN as a WAN.

- **Marketing point of view:** If any site cannot process purchased items due to a network outage, this will impact the business's reputation in the market.
- **Optimum solution:** High-speed, redundant links should be used.
- **Finance point of view:** Cost savings.
- **Optimum solution:** One cheap link, such as an Internet link, to meet basic connectivity requirements.

Based on the preceding list, it is obvious that the consideration for WAN redundancy is required for the new remote sites; however, cost is a constraint that must be considered as well.

When applying the strategic balance (alignment) concept, each departmental strategy can be incorporated to collectively achieve the overall optimum business goal by using the suboptimum approach from each department's perspective.

In this particular example, you can use two Internet links combined with a VPN overlay solution to achieve the business goal through a cost-effective solution that offers link redundancy to increase the availability level of the remote sites, meeting application bandwidth requirements while at the same time maintaining the brand reputation in the market at the desired level.

Network Design Principles

Network design can be defined as the philosophy that drives how various components, protocols, and technologies should be integrated and deployed based on certain approaches and principles to construct a cohesive network infrastructure environment that can facilitate the achievement of tactical or strategic business goals. Previous sections in this chapter described the end-to-end process of a network design (PPDIOO) and the approaches that you can use to tackle different network designs (top down and bottom up).

This section, however, focuses on the design principles that network designers must consider when designing a network and translating business, functional, or application requirements into technological requirements. It is important to understand that these principles are not independent of each other. Therefore, to generate an effective design and achieve its intended goals, network designers should consider how each of these principles can integrate into the overall architecture; however, not all of these principles may be applicable or required by every design. By following the top-down approach, network designers can evaluate and decide what to consider and focus on.

Reliability and Resiliency

A reliable network delivers nearly every packet accepted by the network to the right destination within a reasonable amount of time [2]. In today's modern converged networks, maintaining a highly reliable network is extremely important because modern businesses rely heavily on IT services to achieve their business goals. For these businesses (especially service providers and modern financial institutions), it can be even more

critical that their network is considered as a revenue-generating entity. For example, a five-minute network outage that might affect x number of customers can cost the business hundreds of thousands of dollars. Therefore, having a highly reliable and available network architecture that can survive during any network component failure without any operator intervention (also known as *resiliency*) is a key design principle. It is considered a top priority of most of today's modern businesses. For the network to achieve the desired level of resiliency, a redundant component must exist to take over the role of the failed component following a failure event.

Note In spite of the fact that network availability is considered a “top priority of most of today's modern businesses,” not every network needs high availability, and not every part of a network requires the same level of availability considerations. For more information about high-availability considerations, see Chapter 9, “Network High-Availability Design.”

Modularity

Modular design is commonly used in software development, where an application can be built from multiple blocks of codes that collectively integrate to form the desired application. These “building blocks” of modular structure enhance and simplify the overall application architecture. For example, if an issue exists in one block or module of that software, it can easily be isolated from the other modules and fixed separately without impacting other parts. Furthermore, from the perspective of ongoing enhancements, it is easier to add additional modules or blocks to this structure if new features are required. This makes the overall application architecture more structured and manageable.

Similarly, modularity is one of the fundamental principles of a structured network. In a structured network, the network architecture can be divided into multiple functional modules, with each module serving a specific role in the network and represented by an individual physical network. The individual physical network is also known as the places in the network (PIN), such as the enterprise campus, WAN, or the data center.

Consequently, these functional modules are easy to replicate, redesign, and expand. Studies show that when building an IT network, about 20 percent of the budget goes to acquiring the hardware, and 80 percent goes to operational costs [9]. For instance, if a network is designed in a way (for example, flat) that cannot isolate security breaches, system upgrades, or failures in certain parts of the network, it will not be “responsive enough” to adapt to future business requirements such as scalability and fast network convergence.

Whereas with the modular design approach, if any given module faces an issue such as a security breach or the addition or removal of modules, there should be no need to redesign the network or introduce any effect to the other modules. Furthermore, breaking complex parts in the network based on a modular approach into manageable blocks will optimize the overall network manageability. In other words, from a network design point of view, modularity can promote design simplicity, flexibility, fault isolation, and

scalability, as shown in Figure 1-7 [10]. At the same time, modularity reduces operational costs and complexities.

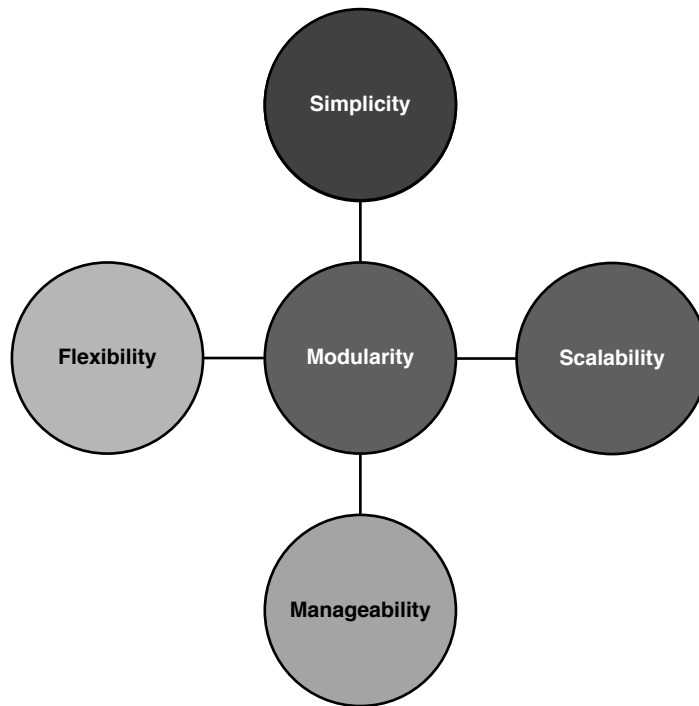


Figure 1-7 *Modularity*

Reliable and Manageable Scalability

Designing for scalability is one of the most important aspects that a network designer should consider. In fact, what can be considered as a successful scalable network is not only measured by the network size factor but also by how stable and reliable the network is and how easily this network can be managed and troubleshot. For example, a network may be designed in a way that can expand to thousands of routers, but at the same time, a failure in one link or device can introduce a high degree of processing and CPU utilization across the network, which may lead to instability. This design cannot be considered a successful scalable design, even though it can grow to a large number of routers. Similarly, a network may be designed to grow to a large scale while at the same time being difficult to manage and troubleshot, with any issue potentially taking a long time to be fixed. This network cannot be considered a successful scalable design either.

Therefore, a successful scalable design must offer a manageable and reliable network at the same time. In other words, the added complexity of the configurations and troubleshooting with the increased size to an unmanageable scale will outweigh any advantage of the ability to expand by the size factor only, as shown in Figure 1-8 [8].

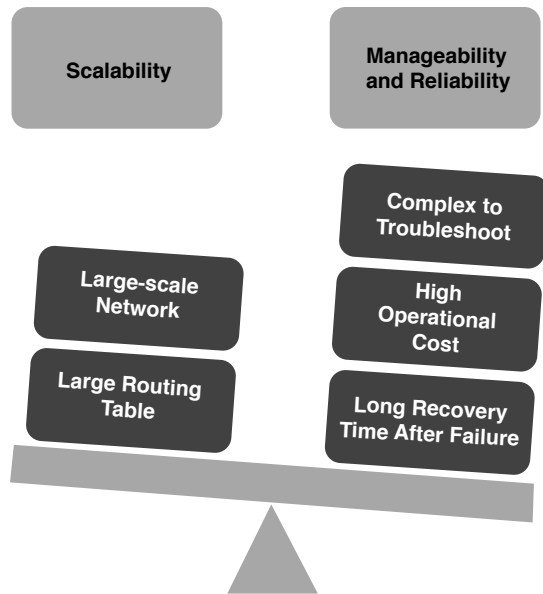


Figure 1-8 *Scalability Versus Manageability and Reliability*

Fault Isolation and Simplicity

Fault isolation boundaries are usually designed to provide fail-stop behavior for the desired fault model. The term *fail-stop* implies that incorrect behavior does not propagate across the fault isolation boundary [3]. In other words, network design must take into account how to prevent a failure in one domain from being signaled or propagated across all other domains in the network. This is necessary to avoid any unnecessary processing and delay across the entire network due to a link or device failure in one domain of the network. For example, in routing design, you can use summarization and logical flooding domains to mitigate this issue, where it is always advised that complicated networks (either physically or at the control plane layer) be logically contained, each in its own fault domain [10]. Consequently, you will have a more stable, reliable, and faster converging network.

Furthermore, this concept introduces simplicity to the overall architecture and reduces operational complexities. In general, fault isolation prevents network fault information from being propagated across multiple network areas or domains, which facilitates achieving more stable network design and optimizing network recovery time after any network failure event. For instance, with regard to the design in Figure 1-9, if any failure occurs on any network node or link, its impact will be propagated to all devices across the campus network. A high amount of unnecessary processing can occur after this failure.

In contrast, with regard to the design in Figure 1-10, if there is any failure at any network node, there will be limited impact (typically contained within one fault domain). This limited impact is facilitated by the separation of the “fault domains” principle, which normally can be achieved by logical separation (for example, hiding reachability and topology information of Open Shortest Path First [OSPF] between the flooding

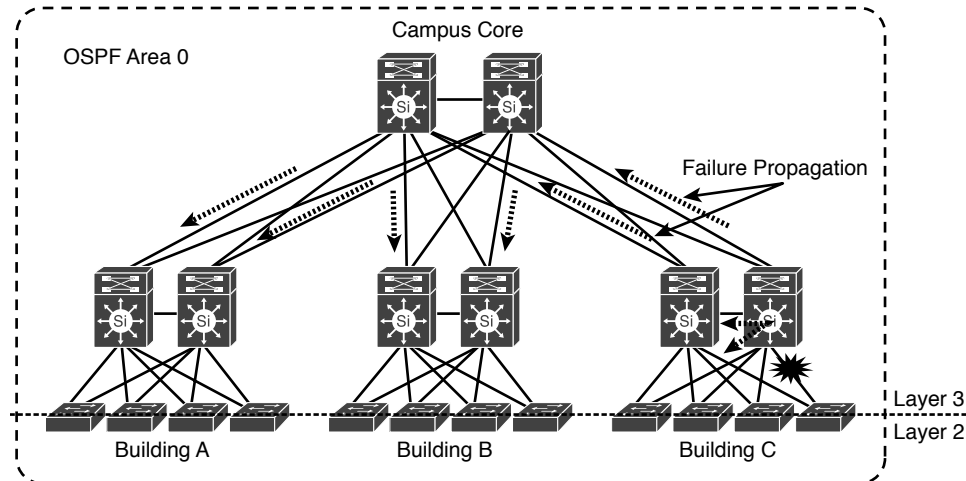


Figure 1-9 *Single-Fault Domain Design*

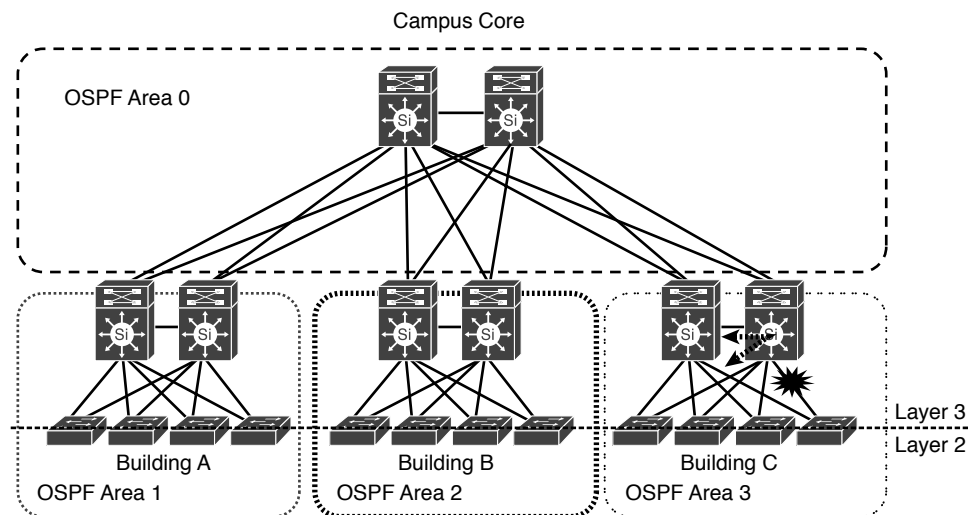


Figure 1-10 *Multi-Fault Domain Design*

domains or blocks). Consequently, the design will be able to offer a higher level of scalability, stability, and manageability.

Hierarchy

Adding hierarchy to the overall network design and within each domain or block can significantly simplify the design and enhance its flexibility to introduce more efficient places to isolate fault domains, because the network designer will have more chokepoints

to aggregate traffic, topology information, and reachability information, which will facilitate the logical (functional) or physical isolation. For instance, Figure 1-11 shows a typical service provider network constructed of two tiers (core and POPs). Each POP is built from multiple tiers as well. This hierarchal structure provides the ability to control and minimize failure propagation between POPs over the core and between the different layers within each POP. Therefore, this design offers a level of stability (fault isolation) and control (chokepoints between the different tiers) to this service provider that makes it more responsive to any future growth requirements (scalable) without adding complexity to the overall design.

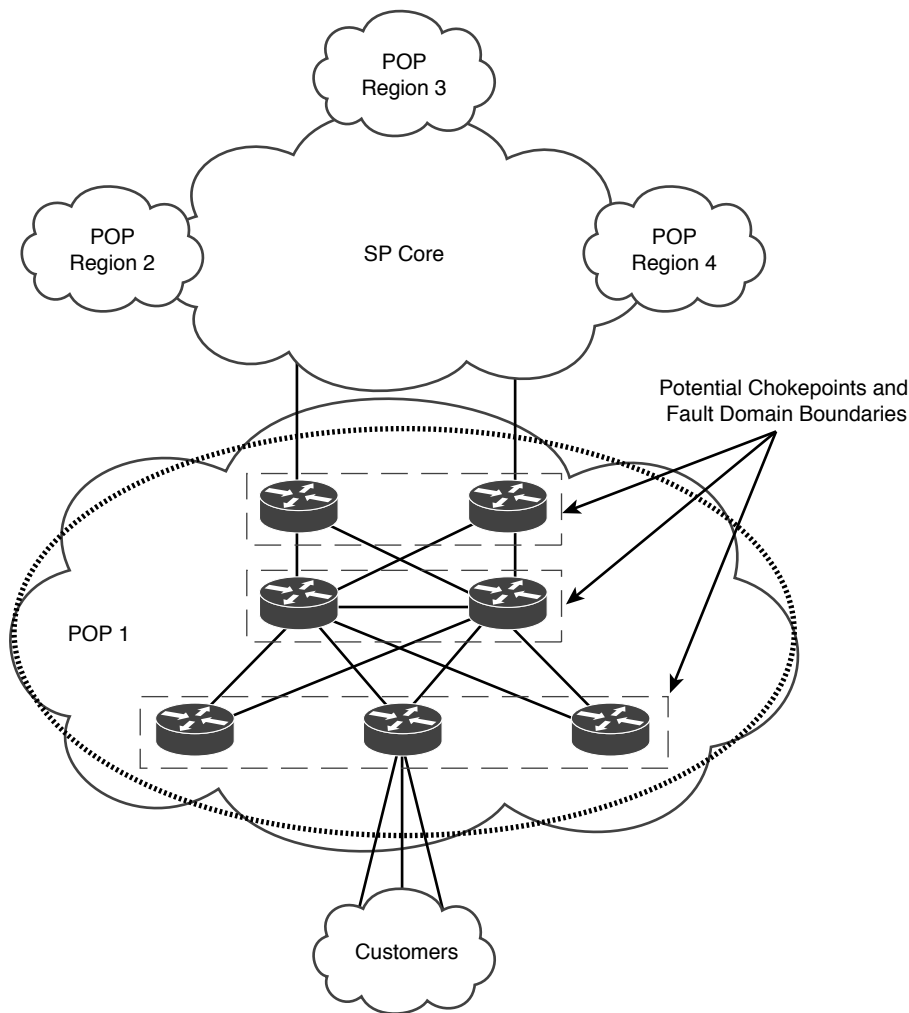


Figure 1-11 *Hierarchal Architecture*

Responsiveness

The term *responsiveness* refers to the set of design characteristics where a flexible architecture needs to respond to the changing business, application, and functional requirements. For example, the design should have the capability to support the rollout of secure mobility for internal staff and guests as well. In this case, a flexible architecture is required to respond to this new requirement and support rapid deployment and integration with different systems without any major change to the overall design. A modular and resilient design combined with the right forwarding and control plane architectures can collectively achieve this goal from the network point of view. When the overall architecture offers a high degree of flexibility, this will significantly facilitate the adoption and integration of the new business and technology requirements in the future, along with more simplified implementations to enable across the network. As a result, the network will be seen as a true service delivery entity from the business point of view.

Holistic Design Approach

A holistic design is a common design principle that considers a system being designed as an interconnected “whole” entity while at the same time potentially being part of something larger [4]. In other words, the holistic design approach emphasizes the big picture of any system, such as an enterprise network. Throughout this book, there are many design options and considerations of different types of networks, whether an enterprise network or service provider network, where each can fit a different set of requirements or solve certain issues. It is important to always consider the big picture using the holistic design concept before applying any of the designs discussed in this book in any given part of the network. Otherwise, you will be approaching the design using a siloed approach that will deal with the network as “communication islands.” The result will possibly be a suboptimal architecture with lower performance, higher costs, and less flexibility. Most important, communications between these islands will probably be inefficient to a large extent because of the limited vision of this approach (not to mention the added complexity to manage networks that were not engineered as a single architecture).

To illustrate the importance of the holistic approach, consider the network in Figure 1-12. This network belongs to a manufacturing company that has a requirement to provide traffic separation between the internal staff and contractor staff from the remote sites across the WAN. A network designer suggested using two VRFs per site with a GRE tunnel per VRF across the WAN network. At the time of the design, there were only five remote sites. After six months, the number of remote sites doubled, and the company’s IT staff started facing many complexities and traffic engineering issues. In addition, there was a requirement to extend this traffic separation to the data center services.

Technically, this is a valid design, and it does provide traffic isolation as per the business requirements; however, this design did not consider the big picture using a holistic design approach. Instead, it was designed based on communication islands (siloed approach), which probably will lead to a nonscalable and inflexible design.

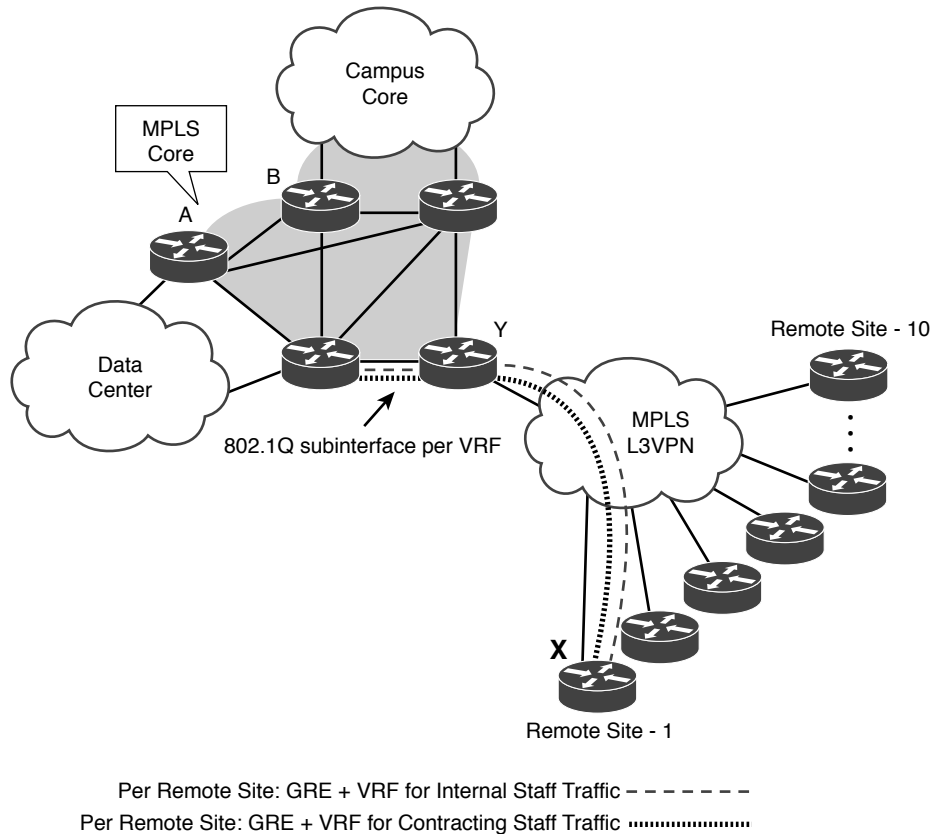


Figure 1-12 *Holistic Approach*

In this particular example, the holistic approach could have helped the designer to consider the business growth requirements. It could have also helped the designer to consider how this traffic isolation can integrate and work as a whole interconnected system with the “MPLS-enabled enterprise core” for optimal and smooth end-to-end communications and integration. Instead, the designer focused only on how to meet this requirement between point X and Y. However, with the holistic approach, the designer can look at the entire architecture to see how point X and Y can optimally communicate with point A and B as well. If the number of remote sites increases, the designer can determine how this may affect the manageability and scalability of the entire network architecture, rather than focusing on only one side or block of the network.

Physical Layout Considerations

In building construction, the foundation carries the load of the structure on top of it, in addition to any future anticipated loads, such as people and furniture. Everything goes on top of the foundation. In other words, a solid and reliable foundation can be

considered the basis of everything that comes after. Similarly, in network architectures, the foundation of the network (which is the underlying physical infrastructure) is where the entire traffic load is handled. It can be a critical influencer with regard to the adoption of certain designs or goals. An example of this is the interconnection of the service provider's core nodes over a physical fiber infrastructure in a ring topology. On top of this physical ring, there is a full mesh of Multiprotocol Border Gateway Protocol (MP-BGP) peering sessions between the PEs at the service provider edges (POPs), as illustrated in Figure 1-13.

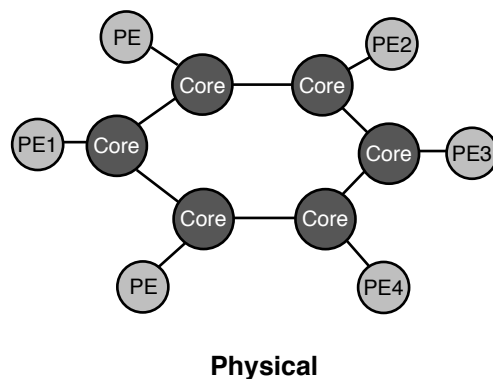


Figure 1-13 *Network Physical Layout View*

From the control plane and logical architectural point of view, these PEs are seen as directly interconnected, as shown in Figure 1-14.




Figure 1-14 *Network Logical Layout View*

However, the actual traffic load and forwarding is handled by the physical core ring network, which means if any core router in the path has a congested link, it can affect all the POPs/PEs communicating over this path.

Consequently, it is important that network designers understand the foundation of the network to evaluate the load, traffic flow, and other aspects that sometimes cannot be

seen from a protocol and the logical design point of view. In addition, understanding the underlying physical infrastructure can help network designers to steer traffic over the desired links based on different criteria such as expense of a given link, bandwidth requirements, and reliability of a given link to recover from a failure. (For example, protected fiber links can be more reliable than nonprotected links.) Throughout this book, the implications of using different Layer 2 and Layer 3 control protocols and overlay technologies are covered, considering the physical layout (the foundation of the network) when applicable. In other words, the logical layout of a network should never be designed in isolation of the underlying physical topology.

Note Sometimes it is impossible, and unreasonable, to have visibility of the physical topology, such as routing or tunneling over the Internet.

Table 1-4 provides a summarized comparison between the most common physical network layouts.

Table 1-4 *Physical Topologies Comparison*

Topology Layout	Scalability	Cost	Manageability	Flexibility	Redundancy
Ring	Limited	Low	Moderate	Limited due to the high level of dependencies	Limited
Fully meshed	Limited (introduce limitations on higher layers [for example, control plane])	High	Complex	Limited	Very high
Partially meshed	Moderate	Moderate	Moderate	Moderate	Moderate
Multiplane core	High	High	Simple	High	Moderate
Leaf and spines	Can offer a high degree of scale-out capability	Moderate	Simple	High	High
Hierarchal	Can offer high scale up	Moderate	Simple	High	High
Hub and spoke	Limited	Moderate	Moderate (depends on the number of spokes)	Limited	Moderate

Note *Redundancy* in Table 1-4 refers to the nature of the topology to support redundancy. Network designers must decide on the level of redundancy and in which areas redundancy is to be placed, based on the design requirements and goals.

No Gold Plating

Gold plating in software engineering or project management (or time management in general) refers to continuing to work on a project or task well past the point where the extra effort is worth the value it adds (if any)[4]. This concept is 100 percent applicable to network design and can be considered one of the critical principles of network design. In fact, overengineering or adding more features and capabilities that go beyond the requirements (whether business, functional, or application) can be a double-edged sword. In other words, a network designer might think that adding extra features to the design will make the customer more satisfied. However, these extra features may lead to additional implementation time and cost (for example, higher end-product prices or more licenses), whereas the design without these features still meets the requirements. Therefore, network designers must produce a design that focuses on how to meet the requirements and become a business enabler rather than add unnecessary features, which can lead to a negative impact.

Note Requirements (business, functional, or application) define what is necessary (and what is not). As discussed earlier, the requirements can take different forms, such as fixing a current limitation or providing a design that will align with the business's future vision.

Summary

To achieve a network design that delivers value to the business and aligns with its goals and directions, network designers must follow a structured approach. This approach must start from the top level, focusing on business needs, drivers, and directions, to generate a business-driven design. In addition, with the top-down approach, network designers and architects can always produce business-driven network architecture that facilitates, at later stages, the selection of the desired hardware platforms and technology features and protocols to deploy the intended design. This makes the network design more responsive to any new business or technology requirements. Furthermore, considering the different design principles discussed in this chapter and considering the different requirements (business, functional, and mission-critical applications) can help network designers to plan and make the right design choices that ultimately will make the network be seen as a “business enabler.”