Elsevier US Job code: NAAD Ch10-P370480 3-5-2007 10:26 a.m. Page:384 Trim:7.5 in×9.25 in TS: Integra, India

CHAPTER CONTENTS

Font: Bembo & Helvetica Size:11/13pt Margins:Top:4pc Gutter:68pt Width:30pc Depth: 38 lines COP: Facing 1 Color

10

Network Design

Network design adds product, vendor, location, and configuration detail to the architecture. Our analysis and architecture work done to date helps make the design process straightforward, reproducible, and well documented. In fact, when the analysis and architecture processes are done well, the design process becomes relatively easy.

Some readers may feel that the reason the design process is straightforward is that we have already done most of what is traditionally termed "design" during the analysis and architecture. This is somewhat true; however, *design decisions* made without benefit of the analysis and architecture processes are ad hoc decisions, made without a sufficient understanding of what we are trying to achieve (problem statements and objectives), what we need (requirements), and how network components will work together (architecture). Ad hoc decision making results in a lot of work at the back end of the project, in terms of resolving conflicts and trying to understand, justify, and explain the design.

By moving much of the work forward in the project life cycle to the analysis and architecture, we are better prepared to make design decisions, procure equipment and services, and provide the background material to support our decisions, and we can couple design decisions to the architecture, requirements, and problem statements. My experience is that this is always a much better position to be in, and well worth the investment in time and resources.

Unfortunately, there is a tendency for engineers to jump to a technical solution without the discipline of analysis and architecture. I have seen this happen many times, always with less than positive results. For example, I was asked to participate in a network project that entailed the development of a national WAN to replace one that is out of date. It turned out that a replacement WAN had already been developed—done quickly with lots of money spent—but its design was ad hoc, without any requirements or architectural forethought. Although it had been implemented and in place for over a year, it had never been used and would never be used. Using the processes in this book we were able to develop a replacement WAN that was highly successful, exceeding the performance and scalability expectations of the customer.

10.1 Objectives

In this chapter you will learn the network design process. This process focuses on two major areas of work for your design: evaluating vendors and service providers, along with their products and services; and diagramming the design—that is, developing detailed blueprints that provide the physical layout for all of your information developed so far. You will learn that a critical result of this work is a set of design decisions that are traceable back to your architecture decisions, requirements, and problem statements. You will also learn how to trace design decisions back to architectural decisions and how to apply metrics to design decisions.

10.1.1 Preparation

To be able to understand and apply the concepts in this chapter to your projects, you should be familiar with the analysis and architecture processes discussed throughout this book; current and emerging network and systems technologies; network topologies and their associated hierarchy and diversity characteristics; and equipment-vendor and service-provider characteristics (although we will go through much of this in this chapter).

10.2 Design Concepts

Network design is the ultimate target of our work, the culmination of network analysis and architecture processes. Whereas network analysis provides understanding, and network architecture provides conceptual (technology and topology) descriptions of the network, network design builds upon these to add physical detail and vendor, product, and service selections to the network.

The upper part of Figure 10.1 shows an example WAN architecture describing the general topology (a multiple-ring topology); strategic locations in the network (Carrier-Independent Exchange Facilities (CIEFs) and demarcation points for each Center and Center LAN); and locations of major types of equipment (switches [S], routers [R], and monitoring devices [M]. Along with such a diagram you would have descriptions of CIEFs and equipment types when needed. In addition, technology choices could be shown on this diagram and included in the corresponding description. This architecture conveys the general concepts of this WAN, including the general structure of the network (a set of carrier facilities interconnected with carrier services in multiple rings, with each CIEF acting as a hub for point-to-point

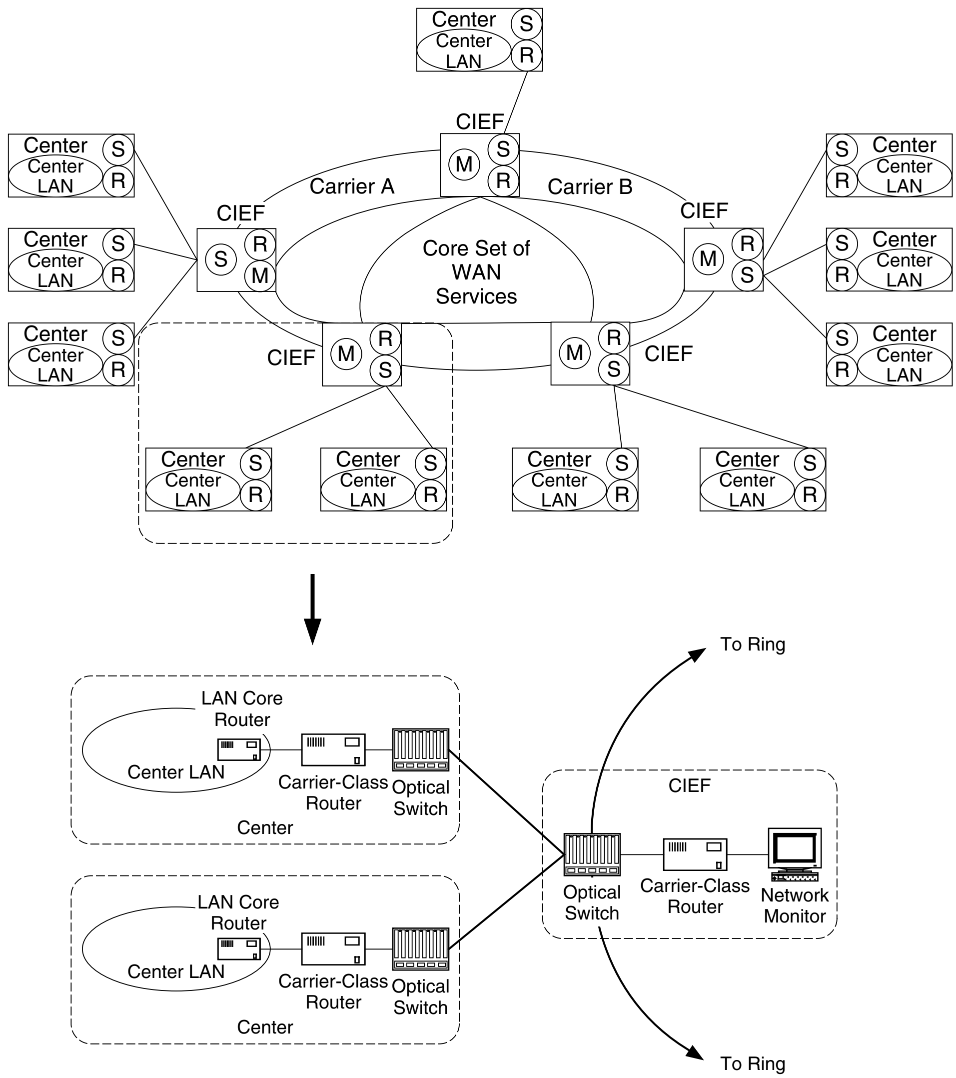

FIGURE 10.1 Network Design Adds Detail to the Architecture

connections to a number of Centers), and general locations for switches, routers, and monitors.

The lower part of Figure 10.1 is an example of how design begins to add detail to the architecture. This figure shows how Centers and CIEFs interconnect. Using generic terms like *carrier-class switches* and *routers*, the diagram indicates what these devices could actually be, based on vendor and product selections made during the design process. Some details can also be shown in terms of the configuration of the devices and how they are connected to each other. However, this diagram does not show all of the possible detail in the products, or the physical locations of equipment. It is useful to note that there can be different products of the design, providing different levels of detail depending on who is the recipient of the product.

What is shown in Figure 10.1 can be considered a *first-order product* of the design. It is not intended to be used to install or configure devices, but rather as a product for management, providing greater detail than the architecture but not so much as to be confusing. What is shown in Figure 10.1 helps describe what is being evaluated and how many devices will be needed for a particular location, and is useful for procurement and general connectivity planning.

Figure 10.2 shows a *second-order product* for one center from the same design. Note the greater amount of detail, including product types, hardware and software revision levels, device configurations, and a more explicit layout of the connectivity between devices, as well as with service providers. Second-order design products provide enough detail to fully understand the network, where devices are in relation to one another, their general locations, and where services such as QoS and VoIP should be enabled.

What is missing from a second-order product such as in Figure 10.2 is the actual location of each piece of hardware in the network design. While a second-order product may provide general locations (buildings, floors, rooms), a *third-order product* adds location detail. For example, the design could show rack layouts and where each device is located in a rack (or equivalent piece of hardware). An important part of a third-order product is showing the explicit connectivity between devices,

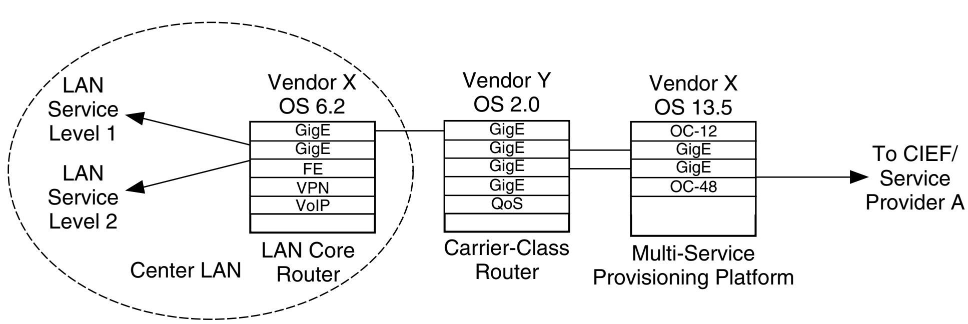

FIGURE 10.2 A Second-Order Design Product

including descriptions and diagrams of cable runs. With a third-order product you should have all of the information you need from the design to install, configure, and operate the network.

It is important to note that third-order products are the ultimate targets for the design. You may end up developing first- and second-order products as by-products while working on third-order products. Thus, it may not be more work developing first-, second-, and third-order design products, and by organizing your design products into these three categories you can support multiple recipients of the design. For example, in the process of developing highly detailed third-order design products for network implementation, you can have first-order design products for management and second-order design products for review.

10.2.1 Analogy to a Building Design

A comparison can be made between the architecture/design of a network and that of a building. For example, in designing a building, the designer/architect needs to know the general purpose for that building (e.g., residential, commercial, industrial), what type(s) of occupants the building is likely to have, approximately how many there will be, and what they will be doing (their requirements). The building design would include the building's physical layout, how space will be used, how occupants will move around within the building, and their needs for building resources (HVAC, lighting, power, water, restrooms, exits, etc.). The product of the building design is a set of blueprints that provide a physical layout for all of this information and show the relationships among components of the building (e.g., where trade-offs were made among space, power, and other resources).

A network design has these same characteristics. Instead of building occupants we have network users; what they will be doing are their requirements (their applications and devices, their work and traffic flows). The building's architecture becomes the network architecture: its topology and technologies. While a building architecture describes relationships among building functions (space, HVAC, power, plumbing, water, lighting, elevators, stairs, etc.), the network architecture describes relationships among network functions (security, network management, routing/addressing, performance, etc.). And as the product of a building's design is a set of blueprints providing the physical layout for all of its information, the network design is also a set of blueprints and diagrams, along with supporting textual descriptions for the network.

One big difference between a building design and a network design is that there is often an artistic component to a building design, while there is rarely one for a network design. There can be an artistic component to a network, such as when it is a part of an exhibit, show, or conference; however, the vast majority of networks is not seen by its users.

10.2.2 Design Products

The key products of a network design are:

- Network blueprints
- A component plan
- Vendor, vendor equipment, and service-provider selections
- Traceability
- Metrics for measuring design success

Network blueprints describe the physical aspects of your network design: locations of network devices, servers, the cable plant, physical security and secure locations; how devices are interconnected, their interface types and speeds; as well as device-specific and service-specific information. Blueprints usually also show the support infrastructure for the cable plant: building entrances, conduit, cable runs, and the like.

Network blueprints may consist of a single diagram or several diagrams of the network, depending on your needs and the size of your project. If the network design is large, you may need to have one or more high-level diagrams that show the entire network in some detail, along with more detailed diagrams that focus on specific areas of the network, such as the WAN, campus LANs, individual buildings, even floors or rooms of a building (Figure 10.3). As part of your detailed diagrams you may want to focus on strategic locations of your network, developed during the architecture process. This is the traditional and most common form of network blueprint.

For other cases you may want to have several detailed diagrams that focus not on specific areas of your network, but instead on specific functions of your network (as we discussed during the architecture process). When this is the case, it is useful to have the diagrams as overlays (diagrams on clear sheets of plastic), where you can place one diagram on top of another, in order to see how and where the functions are applied within the network, and how and where they overlap (Figure 10.4). You can use the results of the architecture process as the starting point to develop such overlays.

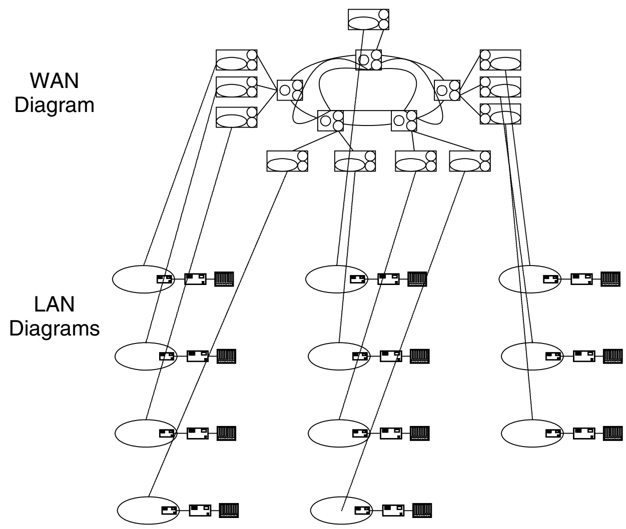

FIGURE 10.3 Diagrams Focus on Geographical Areas of a Network

If you decide to describe the network design as overlays of multiple functions, then each function can have a *component plan* which describes the mechanisms associated with that function, internal interactions among those mechanisms, and external interactions among functions. Component plans can be complementary to the network blueprints, often providing function-specific information that normally would not be on blueprints. For example, a component plan that describes security for the network might include configuration information for security devices, such as sets of ACLs or VPNs, and where in the network each set would be applied.

Each way of describing the network has its advantages. Describing geographical areas of a network (e.g., WANs, LANs) is intuitive and useful for installation, operation, troubleshooting, and modification of that network. Since such diagrams have all relevant devices depicted, it is easy to see connectivity between devices, and to trace a path across the network.

Describing functions of a network has the advantage of showing all devices and services of a particular function (e.g., security) together on the same diagram. This allows network personnel to more easily see where those devices are, how that function will be provided in the network, and how those devices are interconnected. For example, in a layered, defense-in-depth, security strategy you may have different security techniques and/or degrees of security at each layer. Functional diagrams can show these security layers, which devices/services are at each layer, and how the devices (and layers) interconnect and interact.

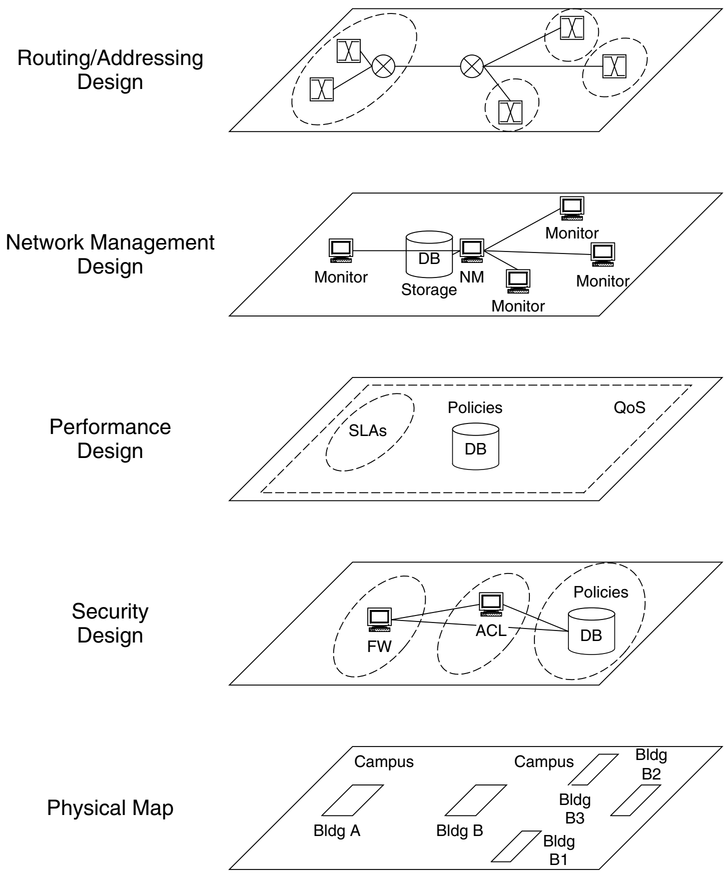

FIGURE 10.4 Diagrams Focus on Logical Functions of Network

Along with blueprints and component plans, the analysis and architecture processes provide information for you to make *vendor, vendor equipment, and service-provider selections*. In making these selections, we apply a process similar to that used to make architecture decisions. Using products from network analysis and architecture we develop an initial set of options and then develop complete sets of candidate options and evaluation criteria; gather data to apply to the evaluations; refine our evaluation criteria and develop ratings; and then apply criteria and ratings to prioritize the candidate options, reducing the number of options to one or a few optimal choices.

Once you know the equipment, vendors, and service providers for your design, you will apply configuration details to your design, which will be used during network implementation and testing. These data consist of general and vendor-specific configuration information and protocol selections (if necessary). For example, sets of ACLs mentioned earlier as part of the security component plan would be considered configuration details. Routing protocols, AS numbers, and specific routing and peering information would also be considered configuration details.

Along with these products you will also be able to show *traceability* between design decisions, architecture decisions, requirements, and problem statements. This is a powerful capability that will help you to address any questions and challenges that may arise regarding your design.

As part of the decisions regarding your design you will have associated *metrics* that are used to describe how you measure success of the design, in the same way that metrics were coupled to requirements. As part of your traceability you may also show how design metrics trace back to requirements metrics. Design metrics can also be used to validate your design.

All of the analysis and architecture work done up to this point provide you with an excellent foundation for making design decisions. This input to the design is discussed next.

Experience has shown that many network projects start at this point, without the underlying analysis and architecture work. The resulting ad hoc design decisions can be disastrous. In addition, I have found that, by applying the analysis and architecture processes, making design decisions becomes the easiest part of the project. Knowing the problems that need to be solved, the requirements for the project, and the network architecture, the design decisions can actually *flow* from this information. Of course, this should be expected: throughout the analysis and architecture processes we are preparing to make design decisions, so that when we get to the design process, we will have already put a fair bit of thought into it.

10.2.3 Input to the Design

At this point in the process we have lots of information to use for our design. These analysis and architecture processes provide products that serve as input to our design. The better these products are, the better our input is, and the better our design will be. The term "garbage in, garbage out" is quite applicable here. In Figure 10.5 architecture products feed into the two parts of the design process—evaluations and layout—producing a number of design products.

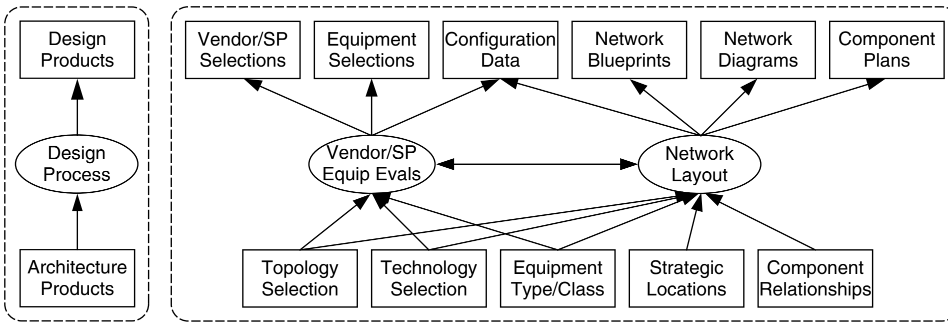

FIGURE 10.5 Architecture and Design Products

Throughout this chapter we discuss how these products are used in the design process.

10.3 Design Process

The design process consists of vendor, equipment, and service-provider evaluations and network layout. These processes build upon and add detail to architectural decisions.

Architecture products include:

- Selected topology
- Selected technologies
- Equipment types/classes
- Supporting material (e.g., traceability diagrams)
- Architectural relationships
  - Routing and addressing
  - Security
  - Network management
  - Performance
  - Others (as needed)

All of these architecture products can be combined into the reference architecture for the project.

Analysis products include:

- Requirements
- Flow information
- Problem statements
- Service information

Vendor, equipment, and service-provider evaluations build on technology and equipment type/class selections made during the architecture process, working toward vendor, service provider, and equipment selections for the design. Such information is often formalized in documents such as requests for proposal (RFPs), which are useful during equipment and service procurement. Recall that during the architecture process we developed a similar document, a request for information (RFI), which is useful in gathering the detailed product and service information needed to prepare for evaluations.

Network layout combines topology, technology, equipment types, relationships, and strategic locations to develop blueprints and component plans for your network design. To some degree, both parts of the design process can be done at the same time and using each other's products (thus the arrow shown between the evaluation and layout processes in Figure 10.5). You may want to incorporate equipment and service-provider information from your evaluations into the blueprints and component plans. Thus, the network layout is usually the last part of the design process.

We now examine each part of the design process in detail.

10.4 Vendor, Equipment, and Service-Provider Evaluations

The evaluation process presented here can be applied to vendors and service providers, their equipment and services. In general, this process consists of using products from network analysis and architecture to develop an initial set of options (termed *seeding the evaluation process*); conducting discussions to develop a complete set of candidate options, along with criteria to evaluate those options; gathering and developing data to apply to the evaluations; refining evaluation criteria and developing ratings; applying criteria and ratings to prioritize the candidate options; and modifying the set of candidate options, with the goal of selecting the optimal candidate. This process is shown in Figure 10.6.

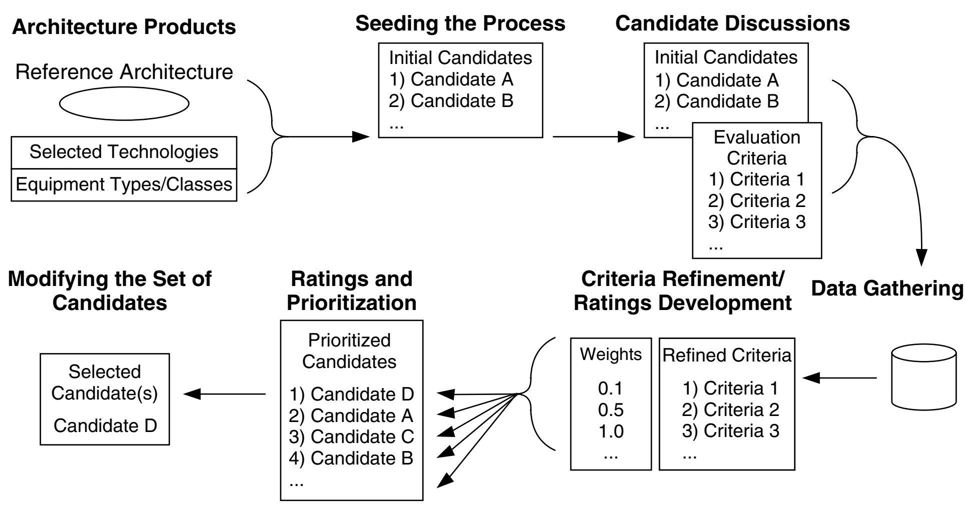

FIGURE 10.6 Vendor, Equipment, and Service-Provider Evaluation Process

This is an iterative process. There are times when an optimal candidate can be found with a single iteration of the process; at other times it may take two or three iterations to produce an optimal result. The number of iterations depends in part on the complexity of the evaluations, how well prepared you are, and how well you and your evaluation team perform the evaluation. As you get used to the process, you will be able to get results with fewer iterations.

Ideally, a template can be developed that can then be applied to each and every evaluation, making the process straightforward, predictable, and reproducible. That is what this section strives to provide. However, we also recognize that, depending on the project, some evaluations may have their own unique characteristics. In all cases I have found that this process can still provide excellent results.

Although this process can be applied to evaluate vendors, vendor equipment, and service providers, it is important to note that we do not want to combine these evaluations into one. Keep each evaluation separate—do not mix vendors, vendor equipment, and service providers—as this will confuse the evaluators and overly complicate the process. For example, some of the evaluation criteria for vendors will be different from those for service providers, and trying to apply criteria to both concurrently would be problematic.

An important part of this process is that you will develop a detailed set of information regarding how your evaluation decisions were made. This information can be used to help reduce or eliminate disagreements that may arise regarding

Elsevier US Job code: NAAD Ch10-P370480 3-5-2007 10:26 a.m. Page:397 Trim:7.5 in×9.25 in TS: Integra, India

your vendor, service provider, or equipment selections. Such disagreements tend to be more likely as the design budget increases.

Example 10.1.

It has been my experience that a large part of the design process is often spent resolving protests brought by vendors, service providers, and even a subset of the evaluators, disagreeing with a selection (often a vendor selection). The evaluation process presented here will provide you with plenty of documentation with which you can resolve or even avoid such protests.

For example, I have participated in vendor evaluations where Vendor A is the popular choice, while Vendor B has the best technical solution. Without a process for evaluating and selecting vendors, Vendor A would have been chosen as the *de facto* choice, without regard to its inferior technical solution. I have used this process to help make it obvious that the technically superior vendor should be chosen.

10.4.1 Seeding the Evaluation Process

The purpose of seeding an evaluation is to get the process started quickly. Seeding consists of generating an initial list of candidates for discussion. One person, such as the project manager, or a few select people can seed the evaluation process. Since the goal is to rapidly kick-start the evaluation process, we do not want to spend much time on this, but rather quickly put together a short list of candidates (vendors, service providers, or equipment, depending on the evaluation) for discussion. This list often consists of the most obvious candidates.

Why is this needed? I have found that it is much easier to generate discussion and develop options when some options are already on the table. This tends to focus the group and gives them something to work toward (or against).

Some or all of the analysis and architecture products shown in Section 10.3 should be directly applicable to start the evaluation process. We can use these products to determine an initial set of candidates. For example, the architecture process provides us with network technologies and network topology that were selected for the project. The selected network technologies may be available only from a subset of vendors (and their equipment) and may be supported by a subset of service providers. Similarly, the selected network topology may be supported only by a subset of service providers. In addition, only some service providers may be able to reach the strategic locations selected for the network architecture. For example, a strategic location may lie in a city that is not readily available to some providers, or it may be prohibitively expensive for them to provide service to those locations.

Font: Bembo & Helvetica Size:11/13pt Margins:Top:4pc Gutter:68pt Width:30pc Depth: 38 lines COP: Facing 1 Color

Architectural relationships, along with strategic locations, may have resulted in specific equipment types or classes being chosen. For example, strategic locations that require high-performance routing, security, and management indicate the need for larger-scale equipment (e.g., carrier-class equipment) where we can combine these functions, or the need for multiple instances and types of equipment to distribute these functions. The class or type of equipment chosen may indicate a particular set of vendors, maybe also particular pieces of equipment.

We also have the products of the network analysis—the requirements and flow specifications—which can be used to help determine initial candidates for evaluation.

The seeding of the evaluation process results in having some candidate vendors, vendor equipment options, or service providers to take into the next step in this process—discussions with project participants.

10.4.2 Candidate Discussions

Having developed a short list (seed) of candidates, we want to use that list to generate a more complete list of candidates, as well as begin to develop a set of criteria that we can use for evaluation.

Whereas one or a few persons can develop the seed list, the complete list should involve a much larger group of participants, including the project team, select potential users of the network, technical staff, management, and/or stakeholders. Your evaluation team should represent the scale and scope of your project.

Using the seed as a starting point, discussions are held, often involving a white-board or Web equivalent, to expand this list. At this point in the process it is better to include all candidates, including obviously inferior ones, and reject them during the prioritization and selection, rather than rejecting them outright. Even if the list becomes unwieldy, it can be quickly pared down once your criteria and ratings have been developed. In this way you will be able to show that you considered a sizable number of options, along with reasons for rejecting them. This can be useful later if someone contests your selection with one that happens to be on your list of rejections. Additionally, you may find that one of your not-so-obvious choices turns out to be one of the better ones.

Along with the complete set of candidates you will need to develop your set of evaluation criteria. Common examples of evaluation criteria are shown in Figure 10.7. Your set will also likely contain evaluation criteria specific to your project, your network, and your organization.

At this point in the process you should not use these criteria to evaluate your candidates. We still need to refine the criteria, as well as develop weights for the

Elsevier US Job code: NAAD Ch10-P370480 3-5-2007 10:26a.m. Page:399 Trim:7.5in×9.25in TS: Integra, India

Evaluation Criteria	
1	Costs
2	Technologies
3	Performance
4	Risks

FIGURE 10.7 An Example Set of Initial Evaluation Criteria

criteria, before applying them. But even before that we need to perform some research and data gathering to support the evaluation.

10.4.3 Data Gathering

Now that you have developed a complete (or nearly complete) set of (vendor, vendor equipment, service provider) candidates, you need to gather and develop data that will be helpful to your evaluations. Common sources and types of information are:

- Discussions with internal and external groups
- Discussions with vendors and/or service providers
- Independent (third-party) assessments of vendors, equipment, and/or service providers
- Modeling and simulation
- Information from risk assessments

Discussions with internal and external groups can provide information in addition to the candidate discussions. The intent of these discussions is to expand the amount of information by talking with more users and staff, and in greater detail. External groups may be able to relate their experiences with candidate vendors, vendor equipment, and service providers. This is especially useful when your design includes technologies or equipment that are new to you, but that others have experience with (e.g., implementation and operational experience with IPv6). External groups may be other organizations like yours, or groups with a common interest, such as user groups or technical forums.

Since we now have a list of candidates, discussions with vendors and/or service providers on that list are often useful to learn specific information about each

Font: Bembo & Helvetica Size:11/13pt Margins:Top:4pc Gutter:68pt Width:30pc Depth: 38 lines COP: Facing 1 Color

vendor, service provider, or piece of equipment. Such discussions can also be helpful in learning what each is capable of doing for your project: providing value-added services, access to labs, testing and validation, and competitive pricing are some examples. However, vendors and service providers should not be allowed to get too close to the evaluation. They are motivated to represent their products and services in the best light and may try to steer the evaluation in their direction. Therefore, it is best to keep such discussions isolated from the evaluations (i.e., don't invite them to any meetings where the evaluations are being discussed).

Along with the above discussions, at times it is useful to get independent (third-party) assessments of vendors, equipment, and/or service providers. Independent assessments provide a prospective that is often different from what internal groups, vendors, and service providers will give. Like external groups, independent assessors may be able to provide you with implementation and operational experience with technologies and equipment that are new to you. For independent assessments I tend to choose small companies or individuals (consultants) instead of large consulting firms.

Modeling and simulation of all or part of the network architecture can provide valuable information for your evaluations. For example, you can use computer models to make comparisons of how your network will perform using different service providers, based on their services offered and on their infrastructures. You may already have some of this information if you have done modeling and simulation to refine your network analysis and/or architecture.

In addition to the above, any information from risk assessments that is applicable to vendors and/or service providers can be useful in evaluations. Your risk assessment, performed early in the analysis process, may reveal certain risks for your project. One or more of these risk factors may be applicable to your evaluation. For example, your project may be adverse to the risk of applying a new protocol, technology, or service. If vendors or service providers apply one of these, it can be used as a part of your evaluation.

All of this information is usually compiled as a document in support of the evaluations and is carried forward into criteria refinement and ratings development. In addition, this information can be very helpful in developing formal documents needed in the procurement and implementation of your design. One example of this is development of a request for purchase (RFP) for products and services.

For example, using the set of initial evaluation criteria from Figure 10.7, data gathering would probably allow us to expand that set based on input from various organizations. Vendors would probably be able to add technology and standards-specific information, while we may be able to learn about risks from other organizations that have already completed similar projects. We may be able to learn specifics regarding performance, from these same organizations or from groups that do equipment and system testing.

10.4.4 Criteria Refinement and Ratings Development

Now that you have additional information regarding your criteria (from the data gathering exercise), you can use this information to refine the criteria. Often, in gathering and developing data, we learn that there are some new criteria that should be added; that some existing criteria are not as appropriate as first thought and perhaps should be removed; or that some criteria should be modified according to the new information. The result is a refined and better set of evaluation criteria.

In order to apply these evaluation criteria to your list of candidates, you should have some way to compare and contrast the candidates. This is commonly done with a system of ratings. Ratings show how the candidates compare to one another. Ratings are applied with criteria (which you already have) and weights that show the relative importance of each criterion. In this section we are concerned with developing weights for our criteria. Although this is one of the more subjective parts of the process, it is necessary in order to make selections from your list.

In our example we had the following initial criteria: costs, technology, performance, and risks. From our data gathering we learned that costs should be separated into initial costs and recurring costs, and that other criteria should be added: standards compliance, available services, operations, and scalability. If our seed set of two candidates expanded into five design candidates, our evaluation chart would look something like Figure 10.8. Note that we have a column for relative weights, and a row to total the ratings for each candidate, both of which we have yet to complete.

In this figure we have nine criteria. Each criterion will be given a weight based on the importance of that criterion to the evaluation. This is determined through group discussion, which may include voting on suggested weights. The range of weights that you apply is not as important as maintaining consistency throughout your evaluation. One common range for weights across the set of criteria is 0–1, where 0 means that that criterion has no importance to your evaluation, 1 means that it has the highest importance to your evaluation, and any value in between indicates that criterion's degree of importance. (Note that you could decide that all criteria are equal, in which case you need not assign weights, or you could give each criterion a weight of 1.)

For our example we would take the data gathered so far, along with our refined sets of criteria and candidates, and conduct a discussion regarding how to weight each criterion. Using a range of 0–1, Figure 10.9 shows how the weights might be applied.

	Evaluation Criteria	Relative Weight	Candidates				
			Candidate 1	Candidate 2	Candidate 3	Candidate 4	Candidate 5
1	Initial Costs						
2	Recurring Costs						
3	Technologies						
4	Standards Compliance						
5	Risks						
6	Performance						
7	Available Services						
8	Operations						
9	Scalability						
	Candidate Totals						

FIGURE 10.8 A Refined Set of Evaluation Criteria

	Evaluation Criteria	Relative Weight (0–1)	Candidates				
			Candidate 1	Candidate 2	Candidate 3	Candidate 4	Candidate 5
1	Initial Costs	**0.8**					
2	Recurring Costs	**1.0**					
3	Technologies	**0.5**					
4	Standards Compliance	**0.2**					
5	Risks	**0.9**					
6	Performance	**0.8**					
7	Available Services	**0.2**					
8	Operations	**0.5**					
9	Scalability	**0.1**					
	Candidate Totals						

FIGURE 10.9 A Set of Evaluation Criteria with Relative Weights Added

In this figure recurring costs are weighted highest, followed by risks, initial costs, and performance. Operations and technology are weighted in the middle of the scale, while standards compliance, available services, and scalability are weighted the least. By looking at this comparison chart, you can see the importance the evaluation team places on each criterion.

It is useful to write down how you arrived at each weight and keep this as part of your documentation. If you are ever asked why certain criteria were weighted higher than others, you will be happy to have it documented. I have found that memory does not serve well here: There have been times when I was certain I would remember the reasons for a particular weight (it seemed obvious at the time), only to forget when asked later.

There are additional ways to develop weights for criteria. As discussed during the analysis process, some characteristics that we can apply to criteria are urgency, importance, and relevance. *Urgency* is a measure of how time-critical the criterion is; *importance* is a measure of how significant the criterion is to this project; and *relevance* is a measure of the appropriateness of this problem to the project. The default characteristic is importance.

Each criterion can be evaluated in terms of urgency, importance, and relevance, and a weight assigned to the criterion that is based on all three characteristics. Then the candidates can be evaluated as in the previous example.

In the next section our weights are used with ratings based on how each candidate fares relative to one another for a given criterion.

10.4.5 Ratings and Prioritization

Armed with our criteria and the results of our data gathering exercise we now develop ratings for each candidate. As in developing the first set of weights, this should be a group effort that includes your best technical and executive decision makers. The size of such an evaluation group is important—you don't want such a large group that nothing gets accomplished, nor one so small that decisions made by the group will be protested by others in your organization.

Example 10.2.

My experience is that a group size somewhere between six and twelve is optimal for making these types of decisions. Oddly enough, this works from small to very large designs.

One important requirement of the evaluation group is that there be no vendors or service providers present. This may seem obvious, but it is surprising how often they are able to get involved. Clearly, in order to develop fair and balanced decisions it is best to leave them out. You should have already gotten all relevant information from vendors and service providers during the data gathering process, *before* you develop and apply ratings.

To continue our exercise, we have our evaluation group together along with our evaluation chart. We then need to agree on a scale for our ratings. One common range is the same range previously used, 0–1, where 0 means that the candidate is least relevant for that criterion and 1 means that it is most relevant. Other common ranges are 1–5 or 1–10, where a weight of 1 means that that candidate is the worst or least relevant for that criterion, and 5 (or 10) means that that candidate is the best or most relevant for that criterion. You could use these ranges to rank the candidates, giving the best candidate (according to that criterion) a 1, the next best a 2, and so on.

Expanding the scale, or reversing the numbers (so that 10 is the worst or least relevant) is entirely subjective. However, it is important that your evaluation group agrees on the scale. If you have problems with this, it indicates that you will have trouble during the evaluations.

Let's say that for this example we choose a scale of 1–5, 5 being the best and 1 the worst. We then take one of the evaluation criteria (e.g., costs) and discuss each design option based on this. This should be a democratic process, where each person gets a chance to express his or her opinion and vote on each candidate. The group leader or project manager would break ties if necessary.

For a discussion on costs we may have cost information (for each candidate design) that was provided to us by an independent assessment, from other organizations that are deploying a similar design, from our own past experience, or from the vendors and service providers themselves. You should be able to compile such cost information and use it to determine a relative rating for each design candidate. You would populate the evaluation chart with ratings for all of the candidates for the cost criterion and then do the same thing for the other criteria. Once you have given ratings to each of your candidates across all of the criteria, you can multiply the weight of each criterion with the rating given to each candidate. The result would look something like Figure 10.10.

In this example each candidate is rated from 1 (worst) to 5 (best) for each criterion. Those ratings are shown as the first number in each evaluation box. Then each rating is multiplied by that candidate's relative weight, resulting in a weighted rating. These weighted ratings are shown as the second number in each evaluation box. The ratings and weighted ratings for each candidate are added together, as reflected in the candidate totals at the bottom of the figure. Although totals for both weighted and unweighted ratings are shown (for illustration), only the weighted ratings would be totaled and applied to the evaluation. Having both weighted and unweighted ratings in an actual evaluation would be confusing.

	Evaluation Criteria	Relative Weight (0–1)	Candidates				
			Candidate 1	Candidate 2	Candidate 3	Candidate 4	Candidate 5
1	Initial Costs	0.8	3/2.4	5/4	4/3.2	2/1.6	1/0.8
2	Recurring Costs	1.0	4/4	5/5	3/3	1/1	2/2
3	Technologies	0.5	3/1.5	4/2.0	1/0.5	2/1.0	5/2.5
4	Standards Compliance	0.2	4/0.8	1/0.2	3/0.6	2/0.4	5/1.0
5	Risks	0.9	3/2.7	4/3.6	1/0.9	5/4.5	2/1.8
6	Performance	0.8	4/3.2	5/4.0	2/1.6	1/0.8	3/2.4
7	Available Services	0.2	5/1.0	1/0.2	3/0.6	2/0.4	4/0.8
8	Operations	0.5	3/1.5	5/2.5	4/2.0	2/1.0	1/0.5
9	Scalability	0.1	5/0.5	1/0.1	3/0.3	2/0.2	4/0.4
	Candidate Totals		34/17.6	31/21.6	24/12.7	19/10.9	27/12.2

FIGURE 10.10 A Set of Ratings for Candidates

Notice from this figure that, if we follow the unweighted ratings (the first numbers), Candidate 1 has the highest score. However, using the weighted ratings, Candidate 2 receives the highest score. This is because Candidate 2 has the highest ratings for those criteria that have the highest weights. Thus, applying weights allows you to focus the evaluation on those areas of importance to your project.

Finally, it is helpful to have a way to determine when the overall (summary) ratings are so close that you should declare a tie, and what to do when that occurs. Best practice is to declare a tie when candidates are within a few percentage points of each other.

When ratings development and application to the candidates are done well, it helps to take the politics out of the evaluation process.

Having rated our candidates, it is now time to refine the set of candidates, with the objective of selecting the optimal vendor, equipment, or service provider for our project.

10.4.6 Modifying the Set of Candidates

You now have an evaluation chart that shows how each network design candidate performed against various criteria. Weights were developed for the set of criteria. Each candidate received a rating for each criterion, which were then multiplied by that criterion's weight to form weighted ratings. The weighted ratings were combined to form a total for each candidate. The result of the evaluation is a set of overall or summary ratings that combines the individual ratings for each candidate.

Rank	Candidate	Deltas	
		Relative	Total
1	Candidate 2	0	0
2	Candidate 1	–4	–4
3	Candidate 3	–4.9	–8.9
4	Candidate 5	–0.5	–9.4
5	Candidate 4	–1.3	–10.7

FIGURE 10.11 Rankings of the Candidates after Evaluation

Summary ratings are used to prioritize the set of candidates. For example, from our previous evaluation chart we would get the prioritization shown in Figure 10.11. This figure shows the ranking of each candidate, along with the relative difference in ratings between that candidate and the one next higher in rank (relative delta), and the difference in ratings between that candidate and the top candidate (total delta). Either delta can be used to determine whether a candidate should be dropped from the set, or whether a tie should be declared. If, for this example, we had decided that a candidate had to be within one point of the next higher-ranking candidate (about 5% of the summary rating for Candidate 2) in order to declare a tie, we could see from the relative deltas that no candidates are close enough to be tied, and that Candidate 2 is the clear winner.

With a prioritized set you can eliminate candidates with bad ratings, with the goal of achieving a single optimal candidate. Usually a single candidate is chosen when it is clearly superior, that is, when its ratings are significantly better than all other candidates. Typically, on a first run of the evaluation there will not be one clearly superior candidate; instead, there may be few or several, depending in part on how many candidates you start with.

When the evaluation process yields more than one superior candidate, another iteration of the evaluation process should be performed on this reduced set. Successive iterations of the evaluation process begin with candidate discussions and continue through data gathering, criteria refinement and ratings development, ratings and prioritization, and another modification of the candidate set.

An advantage of doing another iteration of the evaluation process is that you already have most of the work done: the data sets and evaluation criteria. The work done in each successive iteration improves upon that information and should result in clearer and better ratings and prioritizations. Eventually, the evaluations should result in a single superior candidate.

Example 10.3.

Experience shows that often one or two iterations of the evaluation process will result in a single winner. If you do two evaluations and still have more than one candidate left (or worse, several candidates left), then you should reevaluate your evaluation data and how you are performing the evaluations.

10.4.7 Determining the Order of Evaluations

Remember that this process can be used to select vendors, their equipment, and service providers for your design, and that you will apply this process to only one type of evaluation at a time. However, there may be information that you gathered for one evaluation that can also be used for the other evaluations. For example, vendor information, gathered for vendor evaluations, may also be useful to evaluate specific items of equipment from that vendor. Likewise, service providers may be using network equipment similar to those you are evaluating. This may be particularly important from compatibility and end-to-end service perspectives.

This implies that there may be an optimal order for evaluations. In fact, it can be to your advantage to sequence your evaluations so that you can reuse and build upon information. A common order to evaluations (as indicated by the title to this section) is vendor, vendor equipment, and service provider. The logic for doing this is as follows:

1. You may want to choose which vendors you want for your design before making specific selections for network equipment. Vendor choices can strongly influence equipment selection.

   Of course, there is a counter argument to this: that you should not let vendor selections influence your choice of equipment. However, typically your selections for vendors and for equipment are at the very least loosely coupled.
2. Understanding which network equipment you plan to acquire will help you to make better decisions regarding service providers.

An alternative to this is to evaluate vendors and service providers concurrently. This does not mean you would combine your evaluations, but rather perform them separately during the same time frame.

10.5 Network Layout

Network layout takes topology and technology choices; architecture and design decisions; vendor, equipment, and service-provider choices; and strategic locations; and from these develops various views of your planned network design, including detailed logical diagrams, physical blueprints, and function-specific component plans. All products of this process should also show the parts of your existing network that will be incorporated into the new design.

10.5.1 Logical Diagrams

Logical diagrams show the connectivity and relationships among network devices. Relationships show how devices may interact with one another, how they may work together to provide service and support in the network, and what you might expect from them. For example, you could have a logical diagram showing the routers in your network, or showing just the border routers, or just the interfaces to all external networks. Such diagrams may also include security devices and how they will be connected to the routers, providing insight regarding how the routers and security devices would work together at an external interface.

Diagrams that focus on logical relationships do so at the expense of accuracy in physical descriptions (i.e., location accuracy). Such diagrams can provide approximate correlations between devices and their physical locations; however, they do not provide an accurate representation of physical space. I refer to such descriptions as logical diagrams and not blueprints, as they do not provide the traditional spatial accuracy and level of detail expected in blueprints. Diagrams showing logical relationships among devices are quite useful as companions to network blueprints, or as early drafts of blueprints. Figure 10.12 is an example of a network diagram.

This figure is an example of a communications closet. It shows the types of network devices planned for that closet, and how they are logically connected. For example, from the diagram you can tell that there are multiple firewalls and switches in the communications closet. You can also see the connectivity between devices, to the local networks, and to the Internet. At this stage, however, it does not describe the actual equipment or vendor selections, cable paths or types, or the physical arrangement of the devices in racks or shelves. Diagrams such as these are useful for planning purposes; however, they are not detailed enough to be considered blueprints.

Another example of a logical diagram is shown in Figure 10.13. Instead of describing a particular location, this diagram shows the logical interconnection of devices from across a network. This diagram is useful in that it describes the

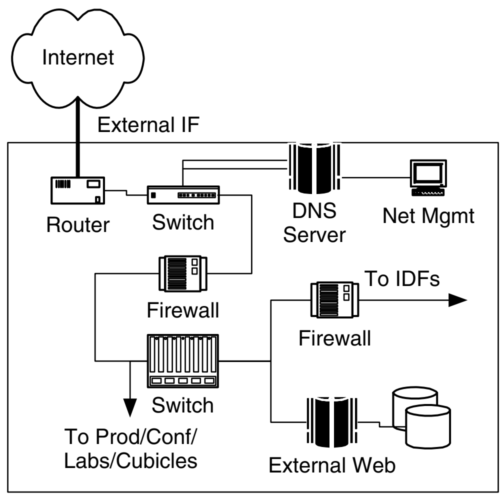

FIGURE 10.12 A Logical Diagram of a Communications Closet

hierarchy of connections, which can be easily mapped to traffic flows in the network.

10.5.2 Network Blueprints

As mentioned at the beginning of this chapter, network blueprints describe detailed physical aspects of your network design: locations of network devices, servers, cable plant, physical security, and secure locations; how devices are to be interconnected, their interface types and speeds; as well as device-specific and service-specific configuration information. There is significantly more detail in a blueprint than there is in a logical diagram.

Network blueprints can consist of a single diagram or sets of network diagrams, depending on network size. If your network design is large you may prefer to have high-level diagrams that show the entire network in some detail, along with more detailed diagrams that focus on specific network areas, such as geographical areas: a WAN, campus LANs, network backbones, individual buildings, even floors or rooms of a building. One focus of blueprints can be on strategic locations in your network.

Developing network blueprints consists of mapping strategic locations of your network onto network templates; applying topology and technology information; and adding your selections of network equipment and services.

Although these steps are described sequentially in this section, it is possible (and often desirable) to apply two or more of them concurrently. Depending on your

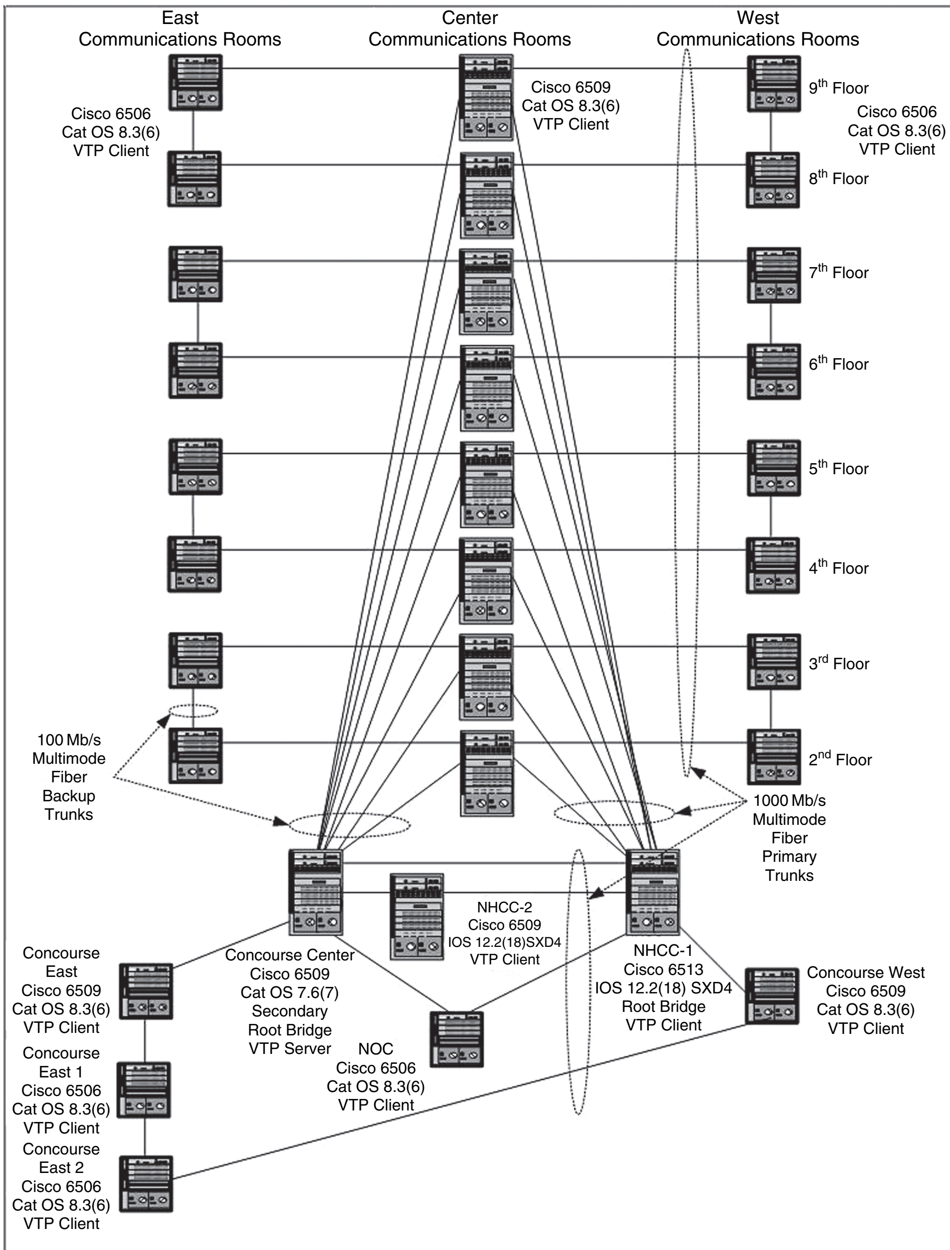

FIGURE 10.13 A Logical Diagram Showing the Interconnection of Devices across a Network

design and how far along you are with your selection process, your technology choices may be strongly coupled to one or more topologies. Also, if equipment choices have been made by this point, they are almost always strongly coupled to technology choices. Thus, you may find that you can map both equipment and technology at each strategic location as you step through the topology.

Mapping Strategic Locations

Developing network blueprints starts with collecting physical information regarding your network (diagrams of rooms, floors, buildings, and campuses; locations of physical plant, communication rooms/closets, and computing and server rooms). This information provides the templates upon which network specifics are described. Figure 10.14 shows an example of a diagram that has been used as the basis for a network blueprint. If such diagrams are not available, you will need to develop your own templates for the network blueprint. If available, you can use building or campus blueprints as the basis for your templates.

Using such templates we map the strategic locations for the network. We want to identify strategic locations as they are of particular importance to network design. From a financial perspective, these locations are likely to be where a number of network functions apply and thus where we may want multiple, high-performance (and more expensive) devices. As such, we will likely spend a higher proportion of our device budget on these locations. From a timing perspective, we typically plan to develop these locations first, as they are key to the development of the entire network. From a topological perspective, strategic locations are usually places where hierarchy is applied. Hierarchy occurs in the network through layering of the

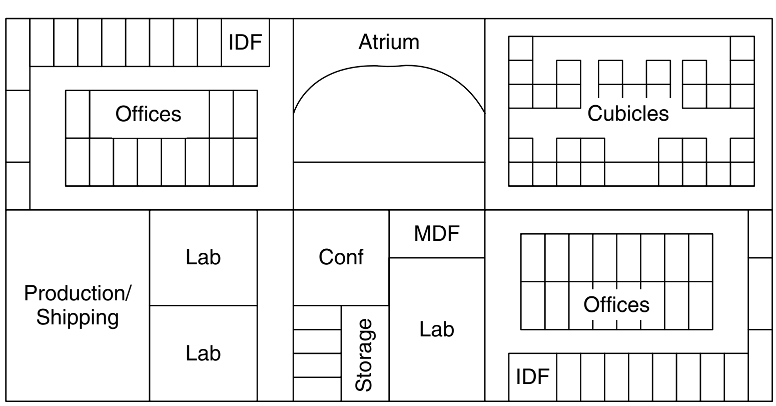

FIGURE 10.14 An Example of a Physical Diagram for a Small Enterprise

network in order to segment and isolate portions of the network, and aggregation of traffic flows as they move through the network. Since we usually impact more traffic flows (and thus more users of those flows) at strategic locations, we have the potential to impact more users at these sites.

You should already have identified likely strategic locations for your network during the architecture process and should be prepared to apply them to your physical diagrams. If not, at this point in the process you will need to determine whether or not you have strategic locations in your network and, if so, where they are.

Some examples of possible strategic locations based on network function are places where:

- Boundary points occur between security cells/zones.
- Major components of monitoring and management are located, particularly when this is for other than OAM functions.
- Multiple performance mechanisms, such as QoS and policies, are needed.
- Different classes or routers or switches interface. This indicates an aggregation of routes, networks, and/or traffic flows. This is also a strategic location based on hierarchy.
- Multiple routing protocols (EGPs and IGPs) coexist.
- Multiple network functions (e.g., security, network management, performance, routing) coexist.

The more that such examples are co-located, the more likely it is that that location is of strategic importance. Some examples of possible strategic locations based on hierarchy are at:

- LAN–MAN and LAN–WAN interfaces
- External interfaces of buildings.
- External interfaces of a campus (this may also be a LAN–MAN or LAN–WAN interface)
- Interfaces between access networks and a LAN backbone
- AS boundaries

Common strategic locations are found at external interfaces of the network, where your network connects to other networks, such as to other autonomous systems or to the Internet. In many small to medium-sized enterprise networks, this collapses to a single external interface (although multiple external interfaces may be necessary when external route diversity is required).

Such locations are strategic for several reasons: First, they are often the most significant boundary in the network (and may require different protocols for each interface, internal and external); second, they often require the presence of every major network function; third, hierarchy is often present, in that traffic flows are aggregated at these locations. This can result in co-locating relatively high-performance routers, security devices, network monitoring/management, and performance devices at this site.

Another set of common strategic locations is at the interfaces to the network backbone. Although such interfaces lie within your network, they can be places where security boundaries are located and where hierarchy occurs through traffic flow aggregation. Often different security, performance, and routing mechanisms are applied to either side of such interfaces. Recall from your flow analysis that network backbones usually provide transport only to transient traffic flows, whereas networks that connect to and are interconnected by a backbone originate and terminate traffic flows. Using this definition we can apply different mechanisms to transient and non-transient traffic flows. For example, we may choose to apply differentiated services as the IP QoS mechanism in the backbone, where we can map traffic classes to aggregates of transient traffic flows. However, as traffic flows get closer to their originations and destinations, we want to be able to apply QoS on a more granular basis, possibly per IP address, per port, or per traffic flow. Interfaces at the network backbone are logical locations for such duality of function, and this also makes them strategic from a topological perspective.

From a planning perspective, locations that incorporate several of the above examples are strategic in that they are key places for future growth of your system. Since strategic locations have some or all of the functions of security, performance, monitoring, routing, and others, they are excellent places for services that can take advantage of co-location with these functions. VoIP, IP video, messaging, and a variety of distributed applications are examples.

Because strategic locations house a number of important network devices, you want to ensure that their physical environments provide appropriate support. Such locations should provide a relatively high degree of physical security, electrical conditioning, and HVAC. You may need to improve the physical environment of wiring closets, communications rooms, computing rooms, and the like as part of your network deployment when they are identified as strategic locations in your network.

Example 10.4.

On visiting a customer to perform a site review of its network infrastructure, I was amazed to find a strategic network location where critical equipment (routers and security devices) were sitting on tables, in the middle of an open room, in 2 to 3 feet of standing water. There was no restricted access (the doors were wide open), no wiring closet (in fact, the wiring was hanging from the ceiling), and no special HVAC considerations. Needless to say, the renovation of this location was of primary importance in the new network design.

It is possible that there are no obvious strategic locations in your network. If, during the requirements and architecture processes, you cannot identify any of the possible strategic locations listed earlier, you may be developing a structurally and topologically homogeneous network. This is rare, however, and care must be taken to ensure that you have not missed potential locations.

Applying Topology Selections

Your choices of strategic locations should map to your selected topology for the network. Topologies imply hierarchical interfaces, which indicate traffic flow aggregation and the potential to identify strategic locations.

Your topology choice describes the high-level structure of your network and consists of the major locations that support this structure and the interconnectivity between these locations. The term *location* is used here, as there can be much flexibility in describing what is connected in a topology. Each location in a topology may be as large as a network or as small as a single device. Typically, a topology consists of network devices (routers and/or switches), possibly co-located with security, monitoring, and/or performance devices and servers. The sets of (network, support, server) devices at each site that are interconnected by the topology are what I refer to as locations, and are shown in Figure 10.15.

Interconnectivity describes how the locations are connected. This may be as basic as describing the physical paths, or may include the cable plant as well as the technologies used at one or more layers in the network. Depending on your design, you may be able to fully describe the connections at this point or may need to wait until after the topology is applied to the design.

Consider, for example, a ring topology. This topology consists of a ring of locations, where each location is interconnected to two neighboring locations. As shown in Figure 10.16, these are often the two neighbors closest in distance.

Each location implies a hierarchical connection to another portion of the network. For example, if a ring topology is applied to a WAN, then each location

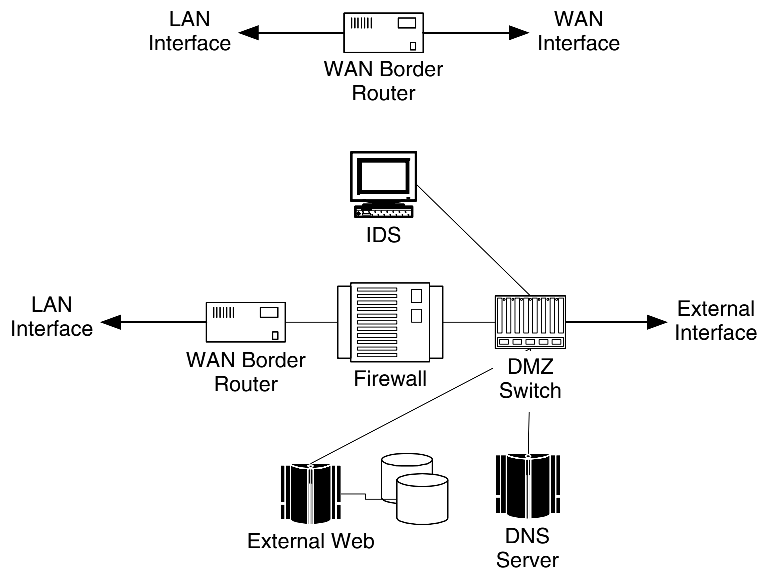

FIGURE 10.15 An Example of Locations

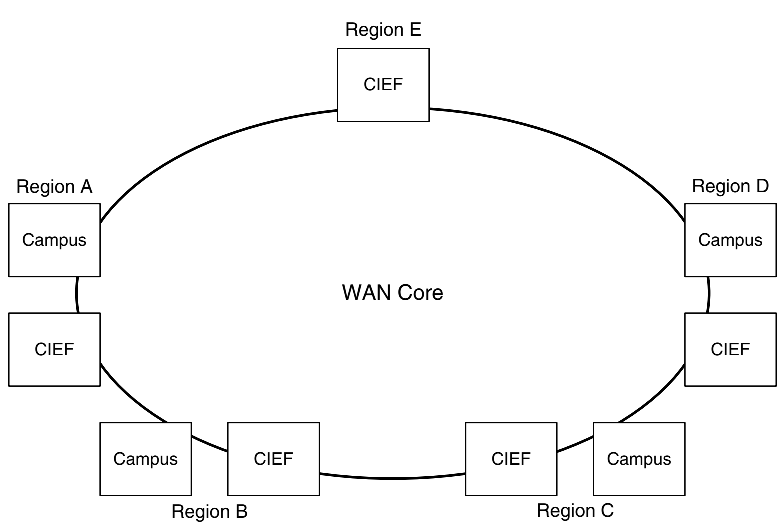

FIGURE 10.16 A Ring Topology

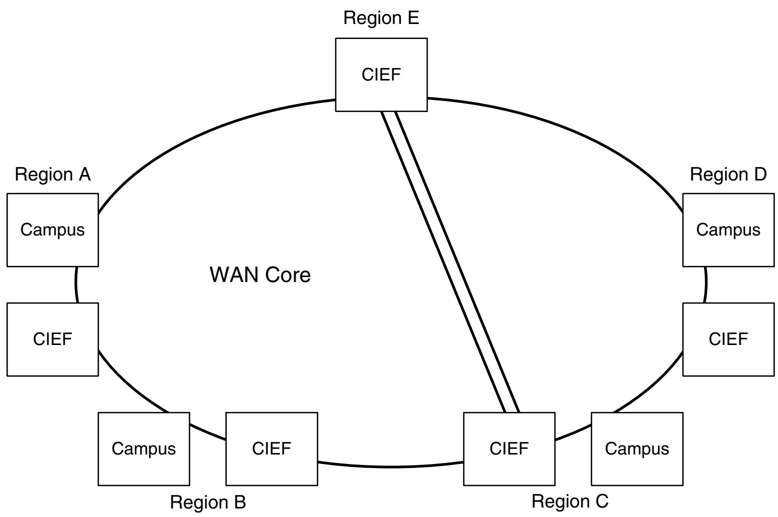

FIGURE 10.17 A Dual- or Split-Ring Topology

can be an interface to one or more regional MANs or LANs connecting to this WAN. If a ring is applied to a LAN, then this implies a LAN backbone, where each location can be an interface to access (edge) or distribution networks in different parts of a building or campus.

A ring topology can be modified into a dual-ring topology (also known as a split-ring or bifurcated-ring topology) by adding a connection between selected locations. Doing so not only modifies the topology, but also changes the traffic flow patterns for the network, adding route diversity. This increases the importance of those locations selected for the additional connection.

In Figure 10.17 additional connections are added between the CIEFs in Regions C and E, splitting the ring and adding diversity across the WAN core.

Other examples are mesh and partial-mesh topologies (a partial-mesh is shown in Figure 10.18). Like the ring or dual-ring topologies, these topologies also show locations and their interconnectivity. The relative importance of each location in these topologies is dependent on the requirements and traffic flows of the network, both of which are expressed to some degree by the choice of technology used for each connection.

The importance of the locations in a topology indicates that they may be strategic locations in the network. Having identified such locations, your selected topology is applied to them. Since strategic locations are closely coupled to topology, typically the full set of locations is used when applied to your topology choice. However, you may apply all or part of your set of strategic locations to your topology.

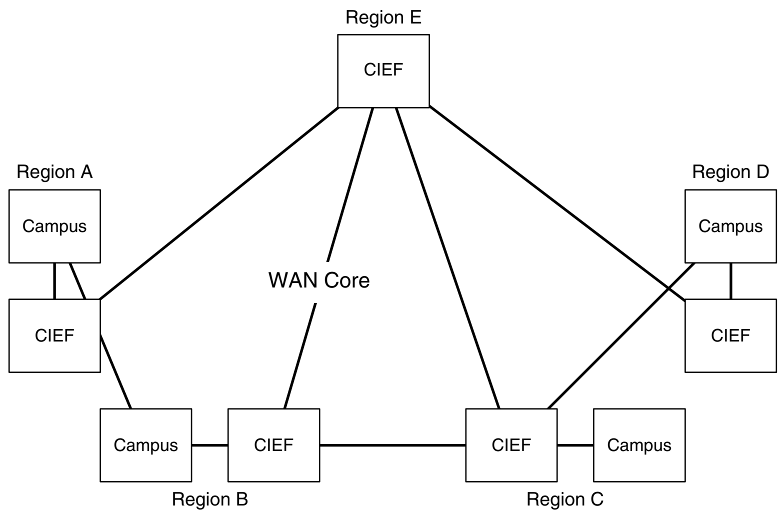

FIGURE 10.18 A Partial-Mesh Topology

A topology provides a start toward developing the design; however, it will not describe every location in your network. It is usually best to begin at the top of your network hierarchy (WAN or LAN backbones) and work your way down. You may apply a topology iteratively, several times throughout your network, or change topologies as you move down the hierarchy.

Applying Technology Selections

At this point you should be able to apply your technology choices from the architecture process. This consists of fully describing the technologies used for your network, and how devices will be interconnected.

In the previous section we described how locations in the topology are interconnected, in terms of physical paths and cable plant. Now we want to apply technology choices across the topology, working from the highest level of hierarchy down to the edges of the network.

For our WAN backbone with a dual-ring topology, we would start at a strategic location (probably one of the locations where the dual rings intersect), describe the technology choices within that location (between devices at that location), as well as technology choices between that location and the other three that it connects with. We would then describe the other location at the intersection of the rings, and then walk the ring, describing each of the locations making up the rings. The result would look something like Figure 10.19. This may be done in conjunction with applying equipment selections, discussed next.

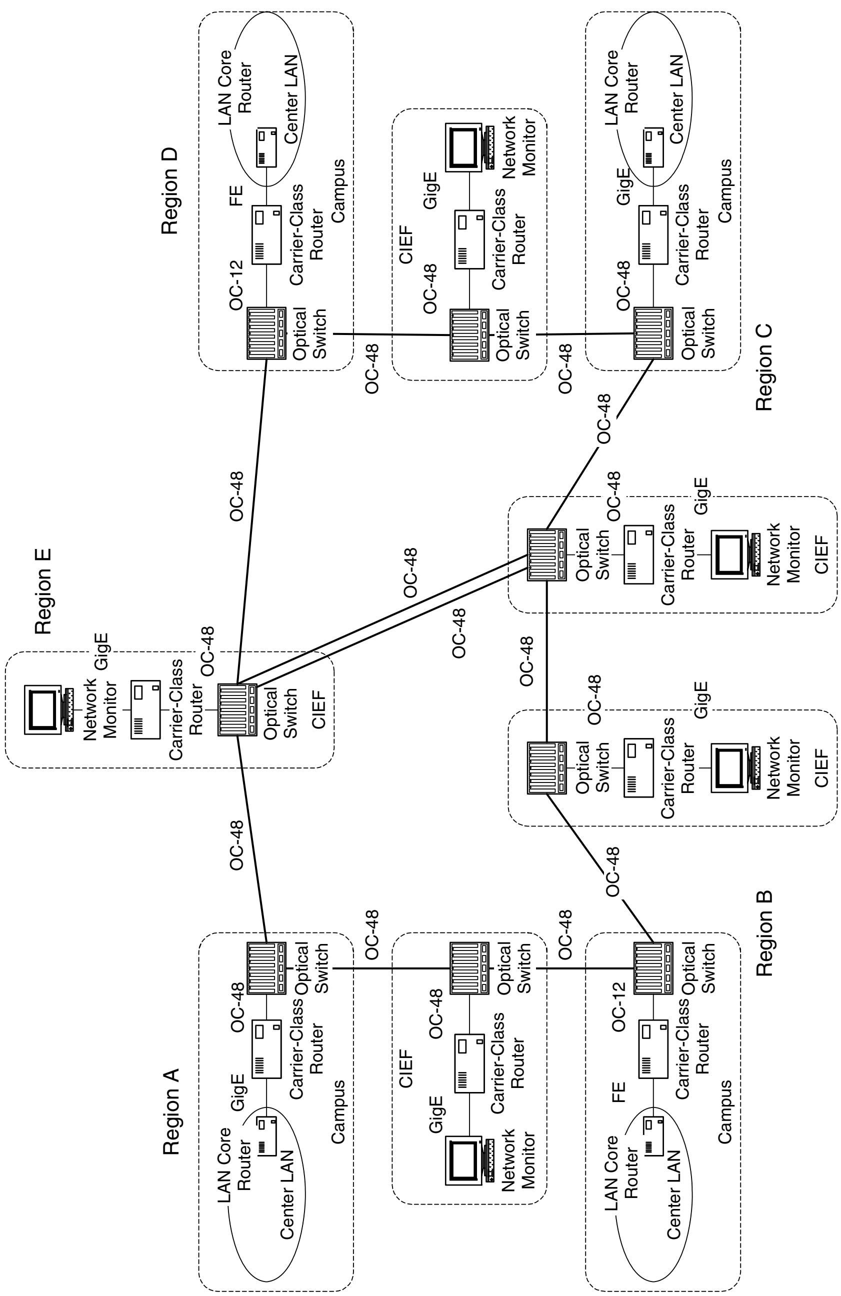

FIGURE 10.19 A Dual-Ring Topology with Technology Choices Added

Applying Equipment Selections

If you have made your vendor and equipment choices by this point, you can apply them to the design. If not, you can use the equipment type and class choices made during the architecture process. Here you should specify some or all of the following for each device:

- Equipment vendor
- Equipment type/class
- Device ID
- Interface types and rates
- Device hardware configuration
- Device OS level/revision
- Degrees of connection, power, internal (card, board, engine) diversity
- Any appropriate vendor-specific information
- Routing protocols used

In order to save space and make the drawings cleaner you can group devices with the same information into specific classifications and provide the information for each classification once on the diagram or in a companion diagram or document. This adds the final detail to the description of your blueprints.

Each location in your network now provides topology information, technology information, and details regarding the equipment to be placed at that location. This can take up a lot of space in your blueprints, and it may be at this point that you decide to develop a separate set of blueprints that separate the network and allow you to provide more information per location. Figure 10.20 below shows a portion of our WAN ring with some equipment information added.

10.5.3 Component Plans

During the network architecture process you should have decided whether or not you were going to provide detailed component plans for specific functions of your network. Component plans build upon information you gathered in developing the network architecture (i.e., mechanisms of each function, interactions among these mechanisms, and interactions among functions) and include diagrams for each function.

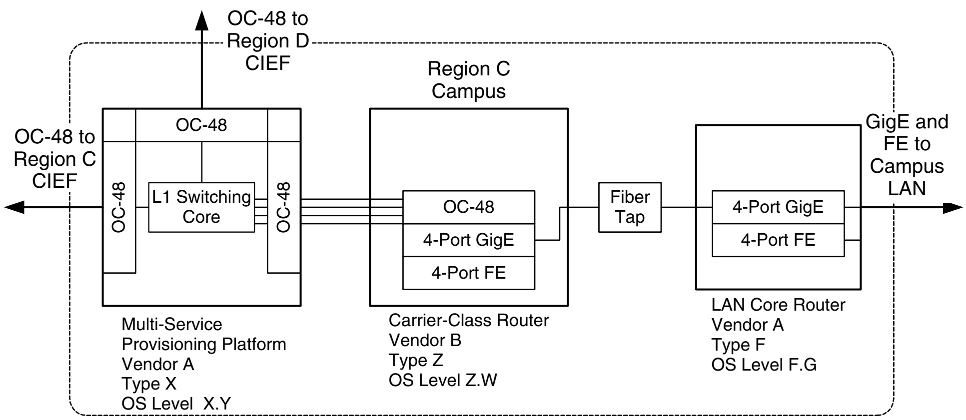

FIGURE 10.20 Region C Campus with Equipment Information Added

Each component plan is similar to a network diagram or blueprint but focuses on a specific function. You need not have a component plan for each function of your network, but rather only for those functions that you choose to focus on.

A common component plan focuses on security, showing which security mechanisms you will deploy and where they will be located. Such a plan can even focus on specific mechanisms. An example of this is describing a VPN architecture as part of a security component plan. A VPN architecture can be complex enough to warrant its own diagrams, as a subset of the overall security diagrams.

Another common component plan is for routing and addressing. This component plan shows where each type of routing protocol is applied, route-specific information, AS information, and where peering and route aggregation occur.

Such plans are especially useful when developed as overlays, so that you can see how each function contributes to the overall design. For example, the component plan for security can overlap with plans for routing, addressing, network management, performance, and others, showing where security devices are co-located and interconnected with devices from other plans.

The purpose of having diagrams, blueprints, and component plans is to provide both a complete set of design data and sets that focus on network specifics. Thus, while blueprints are useful to quickly get the big picture of the network, diagrams are helpful in understanding logical connectivity, while component plans let you focus on routing, security, or other network functions.

In addition, having additional sets of diagrams and plans can actually simplify your views of the network. Network blueprints can become unwieldy, packing so much information into the drawings that they become difficult to read, understand, and use, particularly when trying to trace interactions among various devices, whereas diagrams and component plans are tailored to tracing interactions.

Blueprints and component plans are meant to be complementary. When both blueprints and component plans are developed, the network blueprint describes physical infrastructure, including the cable plant, physical security, and HVAC; while component plans provide detail in each of the network functions. Typically, routing, switching, and addressing are included in the blueprint as well as in a component plan, as they are the most basic network functions.

It is useful to have component plans developed as overlays, where you can place one diagram on top of another, in order to see how and where the functions are applied within the network, and where they overlap (Figure 10.21). As such,

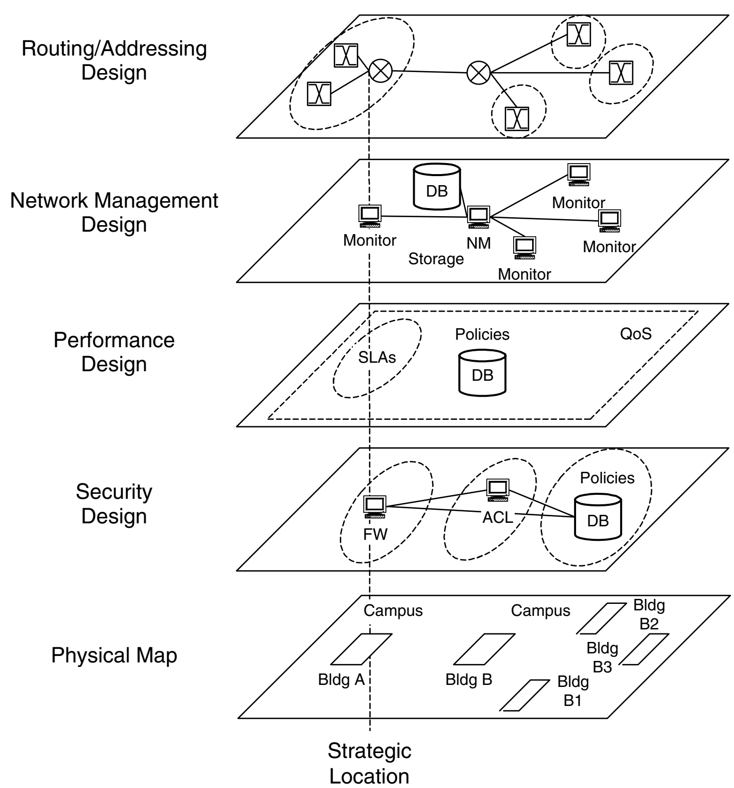

FIGURE 10.21 Component Plan Overlays Line Up at Strategic Locations

Elsevier US Job code: NAAD Ch10-P370480 3-5-2007 10:26a.m. Page:422 Trim:7.5in×9.25in TS: Integra, India

they need to line up structurally, so that each location in the network is shown at the same place in each component plan.

However, in developing your network design you may decide to have the traditional set of blueprints that describes all aspects of your design. For small-scale networks you may be able to describe all appropriate information in a single blueprint or set of blueprints. If you choose to do this, information in this section on developing component plans would simply be applied to developing a set of blueprints.

10.6 Design Traceability

Each decision made regarding your network design should be traceable to one or more architecture decisions and thus to requirements and problem statements. This completes the traceability of your decisions, a critical part of the analysis, architecture, and design processes.

The ability to trace each design decision all the way back to how it addresses the project's architecture, requirements, and problem statements is a powerful capability. Figure 10.22 illustrates this traceability. Each design decision maps to one or more architecture decisions, which then map to requirements, which map to problem statements. This allows you to demonstrate how each and every design decision addresses your project's requirements and problems.

Without traceability you may not know that some of your project's problems and requirements are not being met, that there are architecture and design decisions that do not address project requirements, or that your design may be heavily weighted toward particular requirements and problems. Traceability demonstrates how well your network design addresses project requirements and problem statements.

An example of traceability is shown in Figure 10.23. This example is for a project to define and implement a network security perimeter as part of a defense-in-depth strategy. The design decisions regarding vendor and equipment selections, and standards development, trace to the architecture decisions each one addresses, which map to their respective requirements, and finally to the problem statement. With this technique you can clearly see how decisions made during the analysis, architecture, and design processes apply to solving the project's stated problems.

I have found traceability useful in a variety of situations, including those described here.

Font: Bembo & Helvetica Size:11/13pt Margins:Top:4pc Gutter:68pt Width:30pc Depth: 38 lines COP: Facing 1 Color

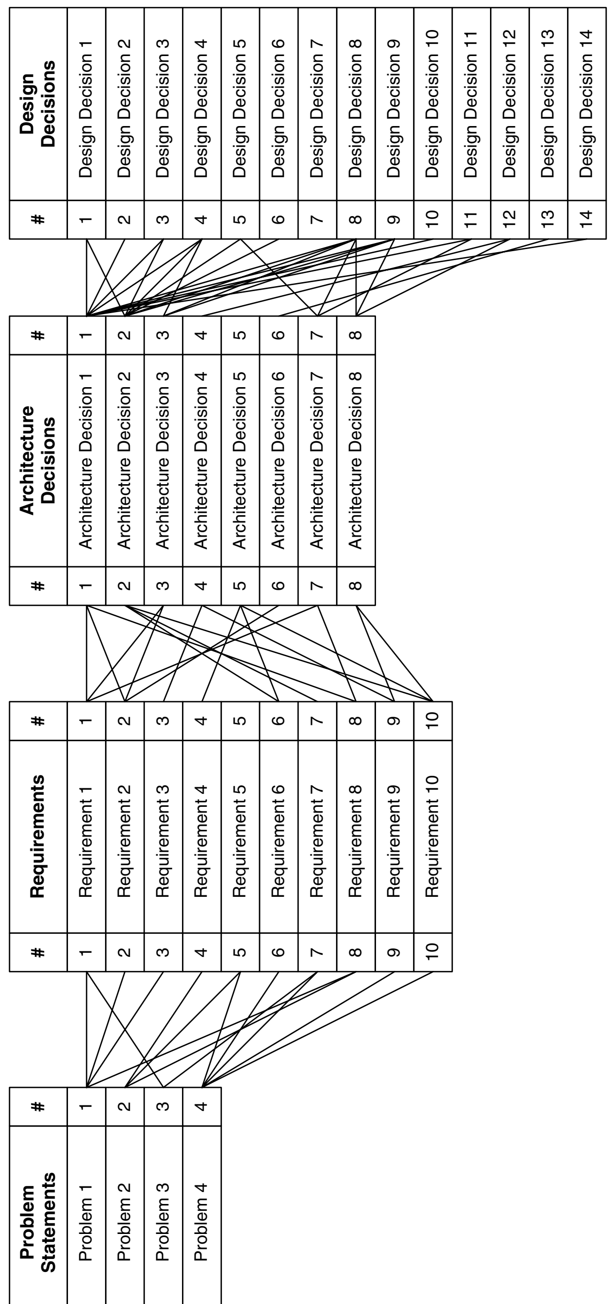

FIGURE 10.22 Design Decisions Trace to Architecture Decisions, Requirements, and Problem Statements

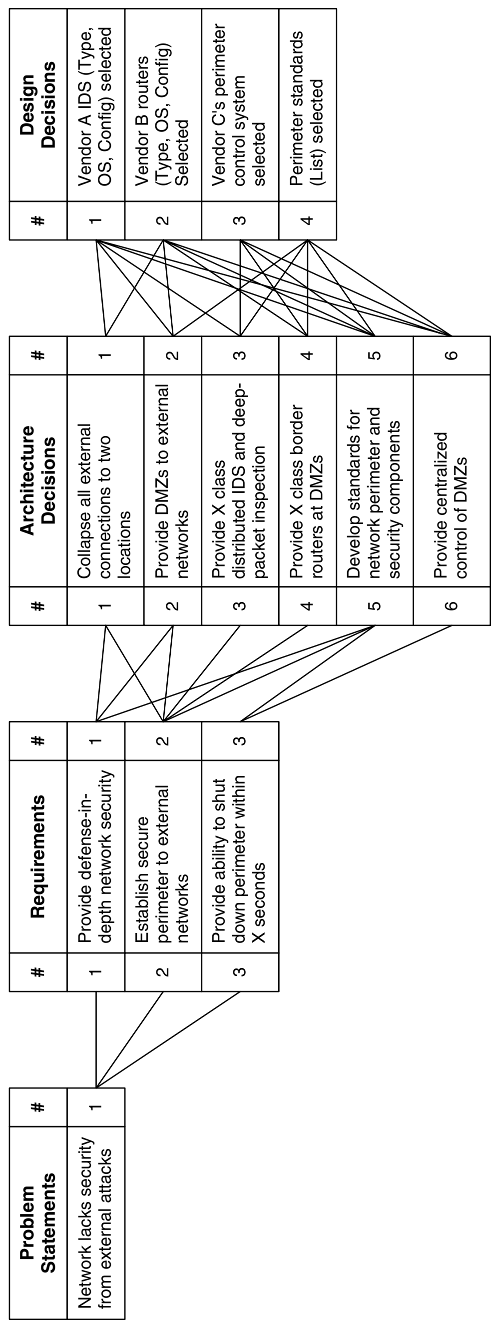

FIGURE 10.23 An Example of Traceability for a Network Security Perimeter Project

Addressing challenges to the network design. For a variety of reasons someone may challenge all or part of your network design. Such challenges may be reasonable and expected, as when new network engineers are added to the design team, or when engineers outside of the project see the design for the first time. With this process such challenges should be rare, however, as network engineers within your organization should have early and frequent access to project information as it evolves through the analysis, architecture, and design processes.

At times these challenges may be more political than technical. This is an unfortunate aspect of projects, one that we must be prepared to address. Fortunately, there will be a wealth of information by the time you get to the design, all of which can be used to address challenges, regardless of whether they are genuinely technical or merely political in nature.

Before applying the discipline described in this book, my experiences were often that network designs would get challenged late in the design process, sometimes in competition for funding, other times over differing opinions regarding design choices. When such challenges were made, there inevitably followed long periods of discussion, wrangling, arguing over competing designs. What I discovered was that, without a well-documented set of problem statements, requirements, architecture and design decisions, it was difficult to address competing design choices (technologies, vendors, equipment). Arguments often came down to differences in personal choices or desires, or at times simply what felt most comfortable—not a particularly strong position from which to drive a project to completion.

However, after applying the analysis, architecture, and design processes, my experiences have been surprisingly (or maybe not so surprisingly) consistent. Armed with well-documented problem statements, requirements, and architecture and design decisions, I have been (and continue to be) consistently able to address any and all challenges to my projects. In fact, when such challenges arise, all arguments quickly dry up as soon as I describe the traceability from decisions to their requirements and problem statements. As you will see for yourself, this is a testament to the power of this process.

Addressing budget, schedule, and resource questions. Another type of challenge to the network design is to justify your budget, schedule, and resource expenditures. Questions such as "Why are we spending X$ for this project?," "What are we getting for this money?," "Why will it take this project so long to complete?" may need to be addressed. In addition, you may need to request more time or funding to complete your project and will have to make the necessary

Elsevier US Job code: NAAD Ch10-P370480 3-5-2007 10:26a.m. Page:426 Trim:7.5in×9.25in TS: Integra, India

arguments. The background information provided by the processes in this book can be extremely helpful in arguing your case.

Bringing newcomers up to date on how the design evolved. The documentation that you have developed throughout the project can be quite useful in helping others follow the evolution of your design. For example, I have found that those unfamiliar with the design (e.g., newly hired engineers) often have questions regarding why and how particular choices were made. Having to explain the design to such folks can be avoided by providing the project documentation. In this regard, it can be handy to develop a project folder that contains the relevant analysis, architecture, and design information. Keeping the proper documentation available is useful, not only to bring newcomers up to date, but also to bring to meetings and discussions when you might be called upon to provide such background information.

You can describe traceability either textually or via diagrams, or with a combination of these. I prefer to use diagrams, because they make it easier to see the connections between problem statements, requirements, architecture decisions, and design decisions. However, I also use text (spreadsheets) to describe traceability, usually as part of the documentation I deliver as products of the project.

From the designer's perspective, showing how design decisions trace back to architecture decisions provides two important sanity checks. First, if you have design decisions that do not map to any architecture decisions (and thus do not map to any requirements or problem statements), then either your design decisions are not appropriate for the project or your set of architecture decisions is not complete. Second, if the mapping between design decisions and architecture decisions is heavily weighted toward a few architecture decisions, then those architecture decisions may be too broad in scope. While the second case is not as critical as the first, it still may indicate a need for a reexamination of the architecture decisions.

In Figure 10.24 there are design decisions that do not map to any architecture decisions and thus do not map to any requirements or problem statements. This demonstrates either that these design decisions (numbers 2, 6, 7, and 14) are important to the project, but we are lacking the proper architecture decisions (and possibly requirements and problem statements); or that these design decisions are not important for our project. Using traceability you can clearly see when this occurs. Without it, you could end up with design decisions that are inappropriate and costly for your project.

In the same figure we also have the opposite problem: There is an architecture decision that has no design decisions mapping to it (number 5). This demonstrates

Font: Bembo & Helvetica Size:11/13pt Margins:Top:4pc Gutter:68pt Width:30pc Depth: 38 lines COP: Facing 1 Color

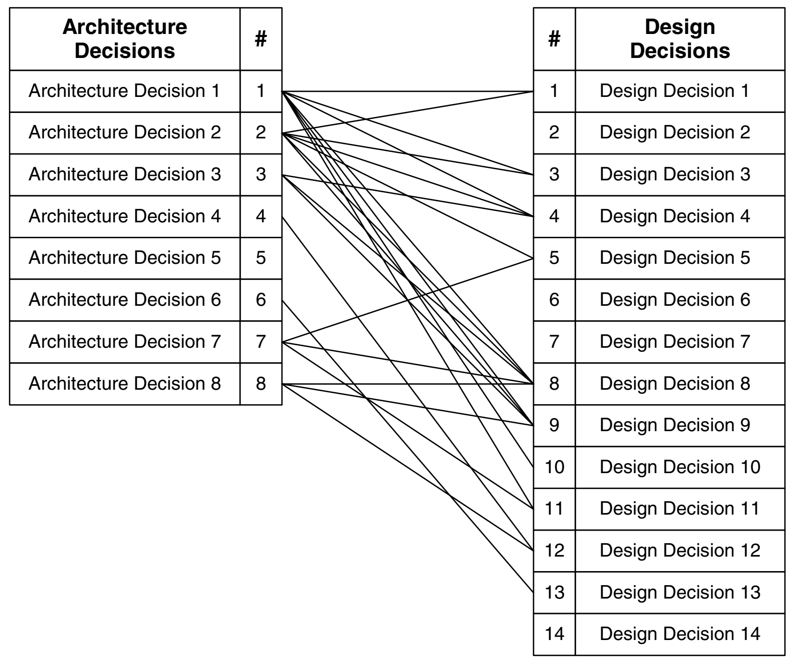

FIGURE 10.24 Traceability Shows Incomplete Mapping between Design and Architecture Decisions

either that we have not done a complete job of developing our design decisions, so that there are some that address Architecture Decision 5; or that this architecture decision is not important to our project. If it is not important, it should be removed and its mapping to requirements reexamined.

In Figure 10.25 we have a complete mapping between architecture and design decisions; however, the mapping is heavily skewed toward a few architecture decisions (numbers 1 and 2). This may not be a problem; in fact, it may demonstrate that these decisions are more significant than the others. Your development of problem statements and requirements may have even indicated that these architecture decisions would be more important than the others. However, it may also indicate that these architecture decisions are too broad, that they could be separated into more focused decisions. When you see such skewing in traceability, it indicates that you should reexamine your architecture decisions to make sure they are what you really want.

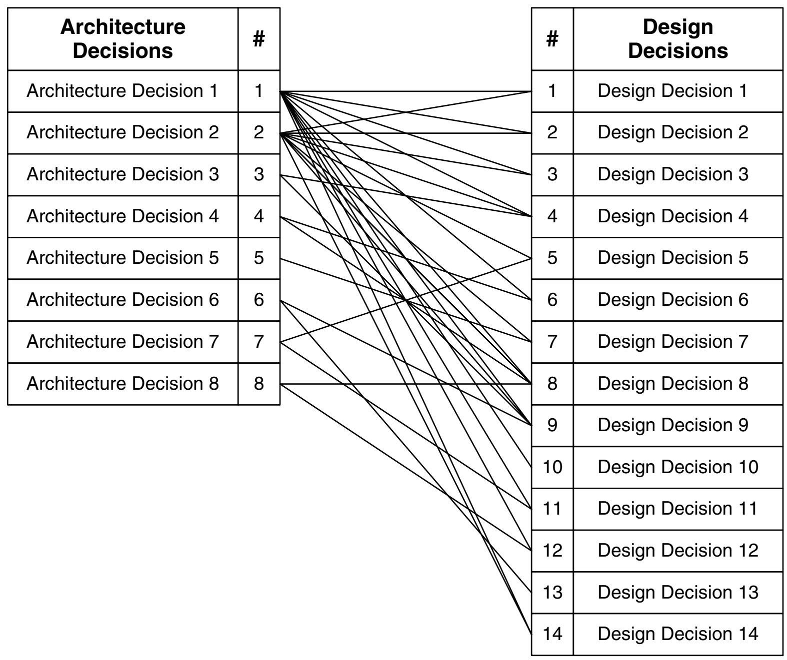

FIGURE 10.25 Traceability Shows Mapping between Design and Architecture Decisions Heavily Skewed toward Architecture Decisions 1 and 2

10.7 Design Metrics

Recall that a necessary part of each validated requirement is that it have one or more metrics associated with it, and that these metrics are used to determine your success in meeting each requirement. Coupling metrics to network requirements is an important commitment that you make to ensure project success. The metrics associated with requirements should also be coupled to metrics you have developed for the design.

Like metrics associated with requirements, design metrics describe how you measure the success of your design decisions. Common examples of design metrics include the ability to achieve a desired diversity level in the network, via routing, topology, equipment redundancy, provider redundancy, and the like, often associated with metrics for performance requirements such as network availability and reliability, or service delivery requirements such as SLAs; the ability to bound end-to-end or round-trip delay in the network, often in support of real-time or interactive applications, again associated with metrics for performance and service delivery requirements; or the ability to deliver desired capacity or throughput levels (either throughout the network or to specified applications and/or devices), also associated with metrics for performance and service delivery requirements.

As you might expect, an additional component of traceability is the ability to trace metrics from the design back to its requirements. It is desirable to do this whenever possible, as it provides direct evidence for satisfying requirements with the delivered design. In the previous examples where design metrics are associated with (and thus traceable to) requirements metrics, demonstrating performance and service delivery in the network satisfies both requirements metrics and design metrics.

In this sense design metrics are useful in validating your design. Figure 10.26 shows an example of this. Using the security perimeter traceability matrix shown earlier, a metric is provided for Requirement 3: Demonstrate that the security perimeter can be shut down within X seconds, where X is to be determined. This metric is supported by two of the architecture decisions—to develop standards for network perimeter and security components, and to provide centralized control of DMZs—by adding the shut-down time (X) as one of the standards to be developed, and by making X one of the criteria by which vendors and products will be evaluated.

Design decisions regarding vendor and service-provider evaluations are influenced by X, which is determined during standards development. One metric for Design Decision 4 is the successful development of X, which can then be used to successfully demonstrate shut down of the perimeter within X seconds, and satisfy the metrics for Design Decision 3 and Requirement 3.

Note that the value of X is critical to the resulting network design and success of the perimeter. Thirty minutes versus thirty seconds for X makes the difference between having a fully automated perimeter and being able to have humans in the loop. It also can make the difference in keeping a security attack from reaching beyond your perimeter.

10.8 Conclusions

The network design is the ultimate product of the analysis, architecture, and design processes. The design is where your architectural decisions regarding network topology and technologies come together with selections for equipment, vendors, and service providers to provide detailed diagrams, blueprints, and component plans for your network.

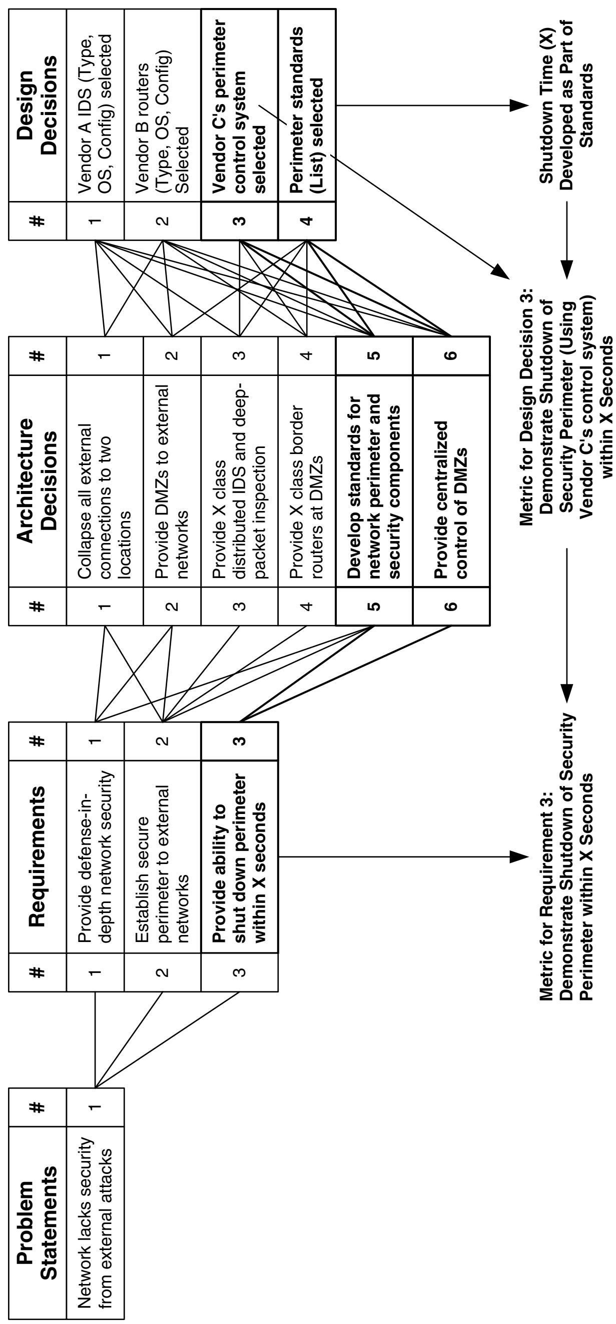

FIGURE 10.26 Metrics for Design Can Be Coupled with Metrics for Requirements

Elsevier US Job code: NAAD Ch10-P370480 3-5-2007 10:26a.m. Page:431 Trim:7.5in×9.25in TS: Integra, India

By using the processes described in this chapter, as well as throughout this book, you will be able to develop better network designs.

10.9 Exercises

1. What are ad hoc design decisions? How do such decisions reduce the quality of the resulting design? Give an example of an ad hoc design decision.
2. If we have followed the analysis and architecture processes, what products from those processes are used as input to the design?
3. What are the primary differences between first-order, second-order, and third-order design products?
4. What are network blueprints, network diagrams, and component plans? Why would a network design have sets of each of these?
5. What are the major components of the evaluation process for vendors, service providers, and equipment?
6. In evaluating vendors, service providers, or equipment for a network design, why would you seed the evaluation process?
7. Which of the following evaluation criteria most likely apply to equipment evaluations, which ones apply to service-provider evaluations, and which apply to both?
   - Available service-level agreements (SLAs)
   - Standards compliance (IETF)
   - Mean time between failure (MTBF)
   - Mean time between service outage (MTBSO)
   - Hardware scalability
8. In the evaluations matrix in Figure 10.10, how would the ratings have changed if:
   a. The relative weights were all changed to be 1 (i.e., there were no relative weights)?
   b. The last three criteria (available services, operations, scalability) were each given a weight of 1.0, while the others remained unchanged?
   c. Each relative weight was multiplied by 10 (i.e., the scale changes from 0–1 to 0–10)?
9. The results of a service-provider evaluation are as follows: Service Provider A: 28.7 points; Service Provider B: 28.5 points; Service Provider C: 29.0 points; Service Provider D: 27.5 points; Service Provider E: 21.9 points; Service Provider F: 27.6 points. How would you modify this set of service-provider candidates?

Font: Bembo & Helvetica Size:11/13pt Margins:Top:4pc Gutter:68pt Width:30pc Depth: 38 lines COP: Facing 1 Color