

Internetwork Planning and Design

Version 1.02

14. January 1998

Markus A. Boeing

This document provides a methodology to deliver the service of planning and designing an internetwork. The design method introduced will enable a network engineer to analyze a customer's current data network and to develop a design that meets the requirements from the analysis.

Designing internetworks requires broad and in-depth theoretical and practical experience. This document does not intend to teach network design. It presents a method to structure the service of network analysis and design and gives some best practices.

Table of Contents

1	PREFACE.....	1
1.1	OBJECTIVES.....	1
1.2	INTENDED AUDIENCE.....	1
1.3	OVERVIEW OF NETWORK DESIGN.....	1
1.4	ACKNOWLEDGEMENTS.....	2
2	ORGANIZATION OF AN INTERNETWORK PLANNING AND DESIGN PROJECT	3
2.1	PROJECT OBJECTIVES STATEMENT.....	3
2.2	FLEXIBILITY MATRIX.....	3
2.3	BENEFITS	3
2.4	POTENTIAL RISKS	3
2.5	WORK BREAKDOWN STRUCTURE.....	4
2.5.1	<i>Deliverables</i>	4
2.5.2	<i>Requirement Analysis</i>	4
2.5.2.1	Determine The Information Required To Do The Analysis	5
2.5.2.1.1	Network Design Questions.....	6
2.5.2.1.1.1	Determine The Service Levels Of The Network	6
2.5.2.1.1.2	Determine Where The Information Flows In The Company	7
2.5.2.1.1.3	Determine The Customers Address Strategy.....	8
2.5.2.1.1.4	Determine The Customers Network Naming Architecture And Design	9
2.5.2.1.1.5	Determine All Aspects Of Network Managing And Support.....	10
2.5.2.1.1.6	Determine WAN Connectivity And Routing Information	11
2.5.2.1.1.7	Determine LAN Connectivity Information.....	12
2.5.2.1.1.8	Determine All Aspects Of External, Non-corporate Connectivity.....	13
2.5.2.1.1.9	Miscellaneous Topics	14
2.5.2.2	Gather The Required Information.....	16
2.5.2.2.1	Gather Preliminary High-level Network Information	17
2.5.2.2.2	Gather Detailed Network Information	18
2.5.2.3	Analyze The Collected Data And Information.....	19
2.5.2.4	Write Requirement Analysis Document.....	20
2.5.2.5	Deliver Requirement Analysis Document And Get Acceptance.....	21
2.5.3	<i>Network Design</i>	22
2.5.3.1	Miscellaneous Network Design Tasks.....	22
2.5.3.2	Develop The Network Design Document.....	23
2.5.3.2.1	Miscellaneous Document Development Tasks	23
2.5.3.2.2	Define External, Non-corporate Connectivity	24
2.5.3.2.3	Develop A Routing Strategy.....	25
2.5.3.2.4	Develop An Address Concept.....	25
2.5.3.2.5	Develop A Naming Concept.....	26
2.5.3.2.6	Develop A Test Concept	27
2.5.3.2.7	Develop A Disaster Recovery/Prevention Concept.....	27
2.5.3.2.8	Develop A Network Design	28
2.5.3.2.9	Finish The Network Design Document.....	29
2.5.4	<i>Implementation</i>	30
2.6	PRECEDENCE NETWORK & RESOURCE REQUIREMENTS.....	31
2.6.1	<i>Skills And Experience Required</i>	32
3	BEST PRACTICES.....	34
3.1	TOPOLOGY	34
3.1.1	<i>Flat Model</i>	34
3.1.2	<i>One-layer Model</i>	34
3.1.3	<i>Two-layer Model</i>	35
3.1.4	<i>Three-layer Model</i>	36
3.2	ROUTING.....	39
3.2.1	<i>OSPF</i>	40
3.3	ADDRESSING	41

3.4	GRAPHING DATA	44
3.4.1	<i>Graphing Utilization Of LAN Segments</i>	44
3.4.2	<i>Graphing Error Condition Of Links</i>	45
3.4.3	<i>Annotated Network Map</i>	46
3.4.4	<i>Traffic Matrix</i>	47
3.5	TESTING AND MODELING	48
3.6	APPLICATION DESIGN BASICS	50
4	SUPPORT MATERIAL	51
4.1	TOOLS.....	51
4.1.1	<i>Free Tools</i>	51
4.1.1.1	CAIDA Measurement Tool Taxonomy	51
4.1.1.2	Multi Router Traffic Grapher	51
4.1.1.3	NetSCARF	52
4.1.2	<i>Commercial Tools</i>	52
4.1.2.1	CACI.....	52
4.1.2.2	Cisco	52
4.1.2.3	Optimal Networks	52
4.1.2.4	Network General	53
4.2	STANDARDS BODIES AND ORGANIZATIONS ON THE INTERNET	53
4.2.1	<i>ACM Special Interest Group On Communications</i>	53
4.2.2	<i>ANSI</i>	53
4.2.3	<i>IEEE</i>	53
4.2.4	<i>Internet</i>	54
4.2.4.1	Internet Ad Hoc Committee	54
4.2.4.2	Internet Architecture Board.....	54
4.2.4.3	Internet Assigned Numbers Authority	54
4.2.4.4	Internet Engineering Task Force	55
4.2.4.5	Internet Research Task Force	55
4.2.4.6	Internet Society.....	56
4.2.4.7	RFC Archive Indexed By Topic.....	57
4.2.4.8	RFC Editor	57
4.2.4.9	Réseaux IP Européens.....	57
4.2.5	<i>International Telecommunication Union</i>	57
4.2.6	<i>Network Device Test Lab</i>	57
4.3	OTHER USEFUL INTERNET SITES.....	58
4.3.1	<i>Searchable Archives Of Mailing Lists</i>	58
4.3.2	<i>Cisco Mailing List</i>	58
4.3.3	<i>Local Area Networks</i>	58
4.3.3.1	The BIG-LAN Mailing List.....	58
4.3.3.2	Charles Spurgeon's Ethernet Site.....	58
4.3.3.3	Ethernet Type Codes.....	58
4.3.3.4	Gigabit Ethernet Alliance	58
4.3.3.5	Measured Capacity Of An Ethernet.....	59
4.3.3.6	ASTRAL	59
4.3.3.7	Token Ring Consortium.....	59
4.3.4	<i>Wide Area Networks</i>	59
4.3.4.1	Dan Kegel' s ISDN Page.....	59
4.3.4.2	SMDS	60
4.3.4.3	SONET & SDH	60
4.3.5	<i>Security</i>	60
4.3.5.1	Computer Emergency Response Team.....	60
4.3.5.2	Electronic Privacy Information Center	60
4.3.5.3	Firewalls	61
4.3.5.3.1	Internet Firewalls Frequently Asked Questions By Marcus J. Ranum	61
4.3.5.3.2	Introduction To Firewalls By NIST	61
4.3.5.3.3	Firewall Mailing List.....	61
4.3.5.3.4	Firewall Products.....	61
4.3.5.3.4.1	Cisco's Centri Firewall.....	61

4.3.5.3.4.2	Cisco's 1605-R Router.....	61
4.3.5.3.5	Nomad Mobile Research Center.....	61
4.3.5.3.6	Phrack Magazine.....	62
4.4	TEMPLATE OF A NETWORK DESIGN DOCUMENT.....	62
4.5	RECOMMENDED READING.....	62

Table of Figures

CPN Requirement Analysis.....	31
CPN Network Design and Implementation.....	31
Flat model.....	34
One-layer model.....	35
Two-layer model with WAN backbone.....	36
Two-layer model with LAN backbone.....	36
Network addresses are allocated bit-wise from the left to the right, host addresses are allocated bit-wise from the right to the left.....	42
Network Utilization and Broadcast Information.....	44
Under normal conditions you would expect the Ethernet error condition graph to form a sail between Utilization and Collisions.....	45
Network Map Annotated with Traffic Volume.....	46
Traffic Source/Destination Matrix.....	47
Measurement vs. Modeling.....	49

Document History

Version	Date	Status	Author	Remarks
X0.1	29. Oct 1997	Draft	MAB	Applied changes to RA suggested by John Jupp
X1.0	13. Nov 1997	Draft	MAB	Released for first review
X1.1	18. Dec 1997	Draft	MAB	Included changes suggested by Rainer Raupach
V1.0	18. Dec 1997	released	MAB	
V1.01	6. Jan 1998		MAB	w/o Pref, Word95 format
V1.02	14. Jan 1998		MAB	

1 Preface

1.1 Objectives

This document provides a methodology to deliver the service of planning and designing an internetwork.¹ The design method introduced will enable a network engineer to analyze a customer's current data network and to develop a design that meets the requirements from the analysis.

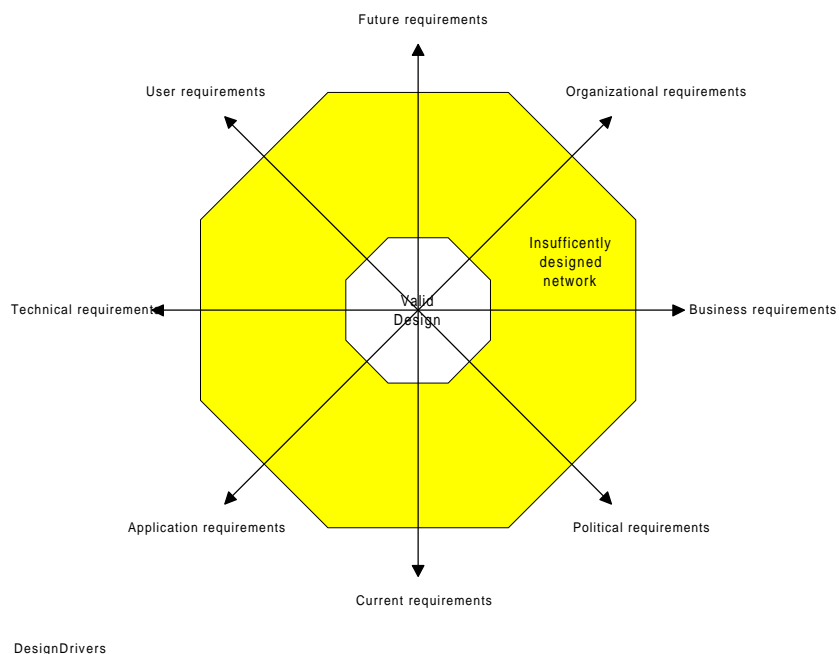
Designing internetworks requires broad and in-depth theoretical and practical experience. This document does not intend to teach network design. It presents a method to structure the service of network analysis and design and gives some best practices.

1.2 Intended Audience

This document has been written for experienced network engineers involved in planning and designing of internetworks. It is assumed that the reader is familiar with project management methods like work breakdown structures and alike.

1.3 Overview Of Network Design

Network design is not limited to the transport of data. It is a broad field that covers analysis, planning and implementation of user applications as well as network architecture, infrastructure and topology, data transport techniques, standards, services and protocols. Focus can not be split between the local and the wide area network, both must be treated as a single system. One of the key features of a good design is that it is built to support both, the present and the future needs. An efficiently designed data network will become vital for a business organization. It can create many options that are not available under the old style of communications. It can be extremely harmful to an organization if it is designed incorrectly.



The factors that influence network design (business, technical, economical and political) are often conflicting. The objective of a good design is reaching a balance point between each end of the spectrum.

¹ An internetwork is a collection of networks interconnected by routers and other devices that functions as a single large network.

Network design is a series of tasks from determining the basic requirements to managing the running network.

The first step in network design should be to gather data such as the corporate structure, the current topology, the applications in use, the information flows within the company, the budget and the resources available.

Analyze the customers business and technical goals. Determine what new applications and services must be supported by the new network design. Discuss the criteria for the success of the internetwork with your customer.

Develop the structure of the new internetwork. Design a scalable and simple topology, determine the hard- and software to be deployed and select the network management strategy.

Estimate the expected performance of the new internetwork using simulation and modeling tools.

Assess the costs of your design.

Implement your design.

Monitor your new internetwork. Is it working like you expected? Are there bottlenecks? Did any applications stop working? Is the utilization on all links like you expected?

1.4 Acknowledgements

Large quantities of the material presented here are derived from Digital Equipment's „TCP/IP Planning and Design Service SDK” and Cisco Systems documents on network design.

2 Organization Of An Internetwork Planning And Design Project

We assume that a feasibility study and project setup have been done already and start right away with the project planning phase.

2.1 Project Objectives Statement

The project objective statement defines briefly the scope, schedule and resources of a project. It includes success and failure criteria.

„Build a network that connects all sites of the XYZ Company within the next six month. The network installation cost may not exceed a total of x DM and the monthly communication cost may not exceed y DM.“

2.2 Flexibility Matrix

The flexibility matrix defines the most and least flexible aspects of a project. You may place just one mark per line. It is a useful tool to take guided decisions on the design direction during the project.

	Least Flexible	Optimize	Most Flexible	Why?
Scope	X			
Schedule		X		
Resources			X	

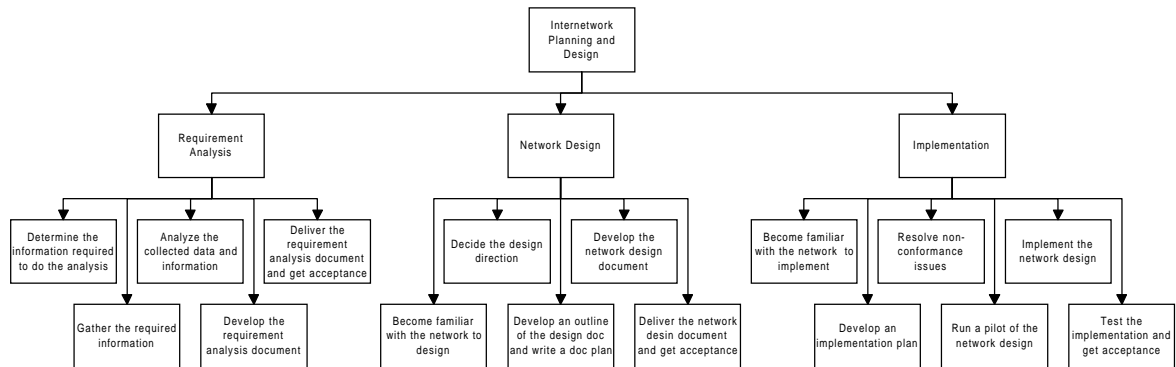
2.3 Benefits

- The customer will develop a better understanding of his organization and how the network fits into it.
- The customer will develop a better understanding of how the network can solve his business problem.
- The customer will get an unbiased expert analysis of his networking environment.
- The customer will get an objective analysis of his present and future networking needs.
- The customer will get a standardized and consistently designed network architecture with potential for growth

2.4 Potential Risks

- The customer does not have correct information and documentation about his networking environment. ⇒ Validate all information using measurement! (Analyze routing tables, debug spanning tree protocol, run your own topology discovery, look for inconsistencies in the customer's documentation, etc.)
- Undocumented connections (leased line or dial-up) may exist.
- Routing backdoors may exist. (Or even worse: They exist and use dial-up lines.)

2.5 Work Breakdown Structure



WBS

2.5.1 Deliverables

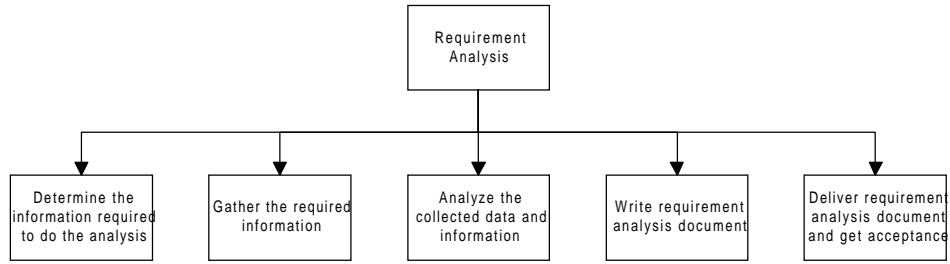
- A requirement analysis document for the customer
- A document for internal use that identifies opportunities and risks for both The consultant and the customer
- A presentation of the requirement analysis document to the customer
- A network design document for the customer
- A presentation of the network design document to the customer
- An implementation plan
- An implementation and acceptance test report

2.5.2 Requirement Analysis

During the requirement analysis you analyze the customers existing and future networking environment. A requirement analysis document will be developed along with a set of recommendations to enhance the customer's network based on the analysis of future needs.

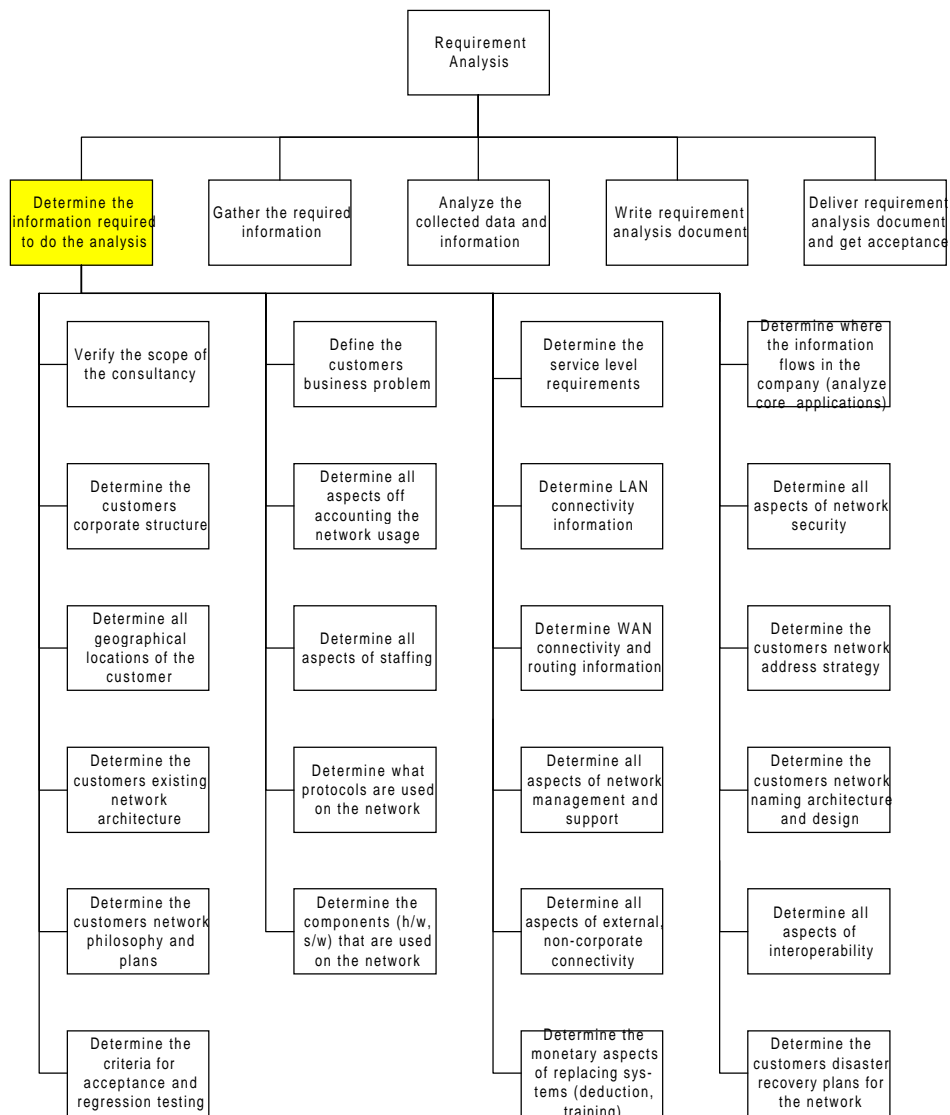
Deliverables

- A requirement analysis document for the customer
- A document for internal use that identifies opportunities and risks for both The consultant and the customer
- A presentation to the customer



WBS RA

2.5.2.1 Determine The Information Required To Do The Analysis



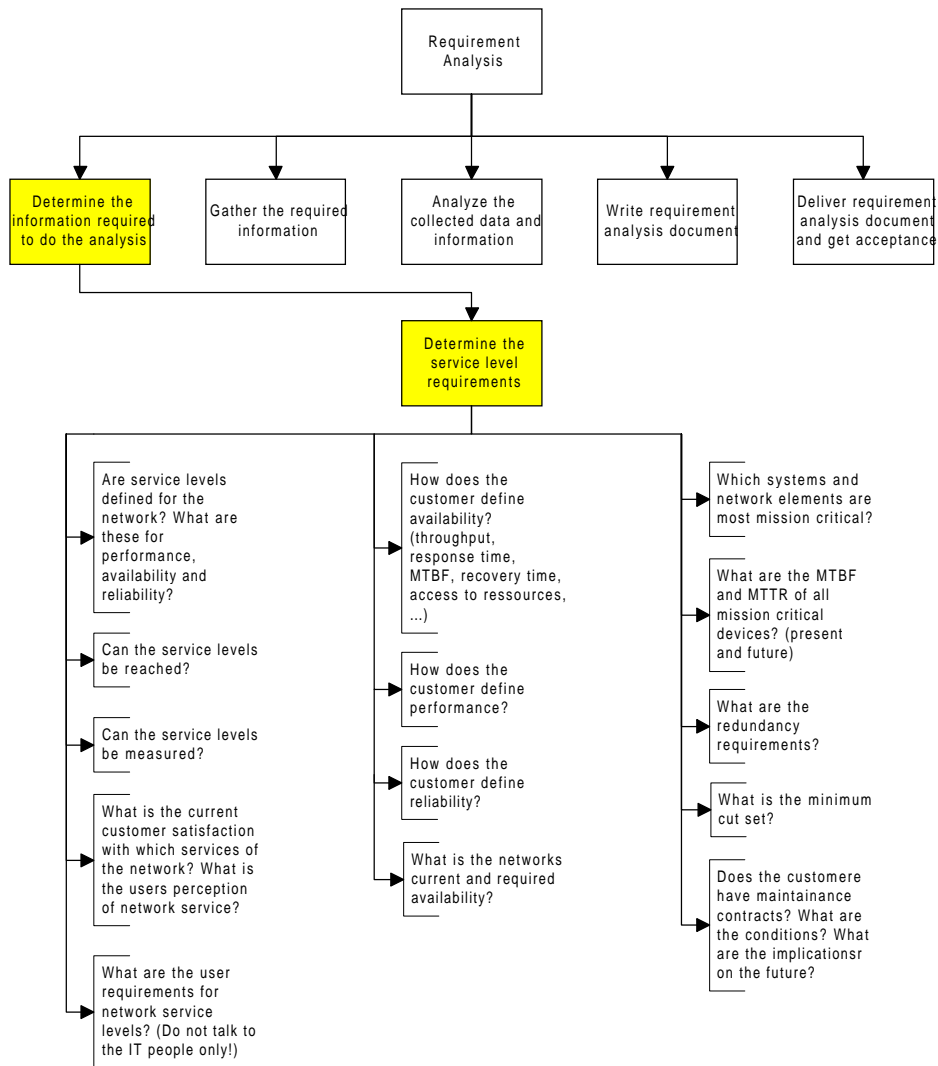
WBS RA.DetReqInf

The items above will be referred to as „Network Design Questions“. They are not listed in any particular order.

The rank of an item within the list does not imply a priority. Please refer to the section „Network Design Questions“ on pages 6 to 14 for more detail.

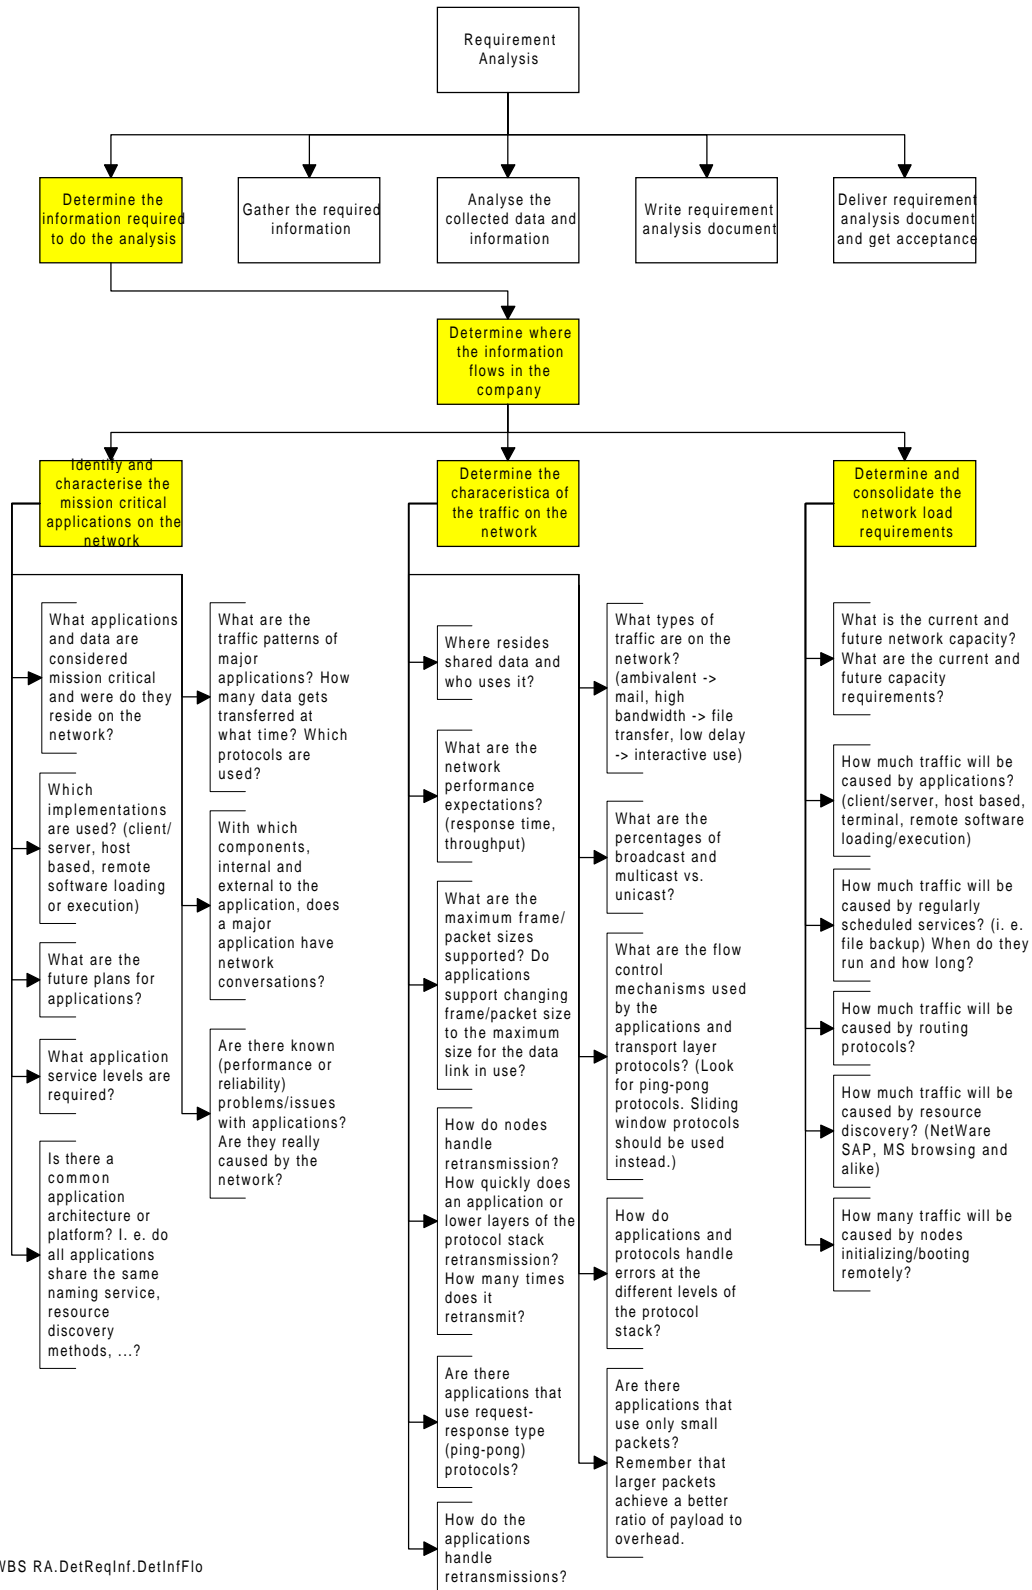
2.5.2.1.1 Network Design Questions

2.5.2.1.1.1 Determine The Service Levels Of The Network

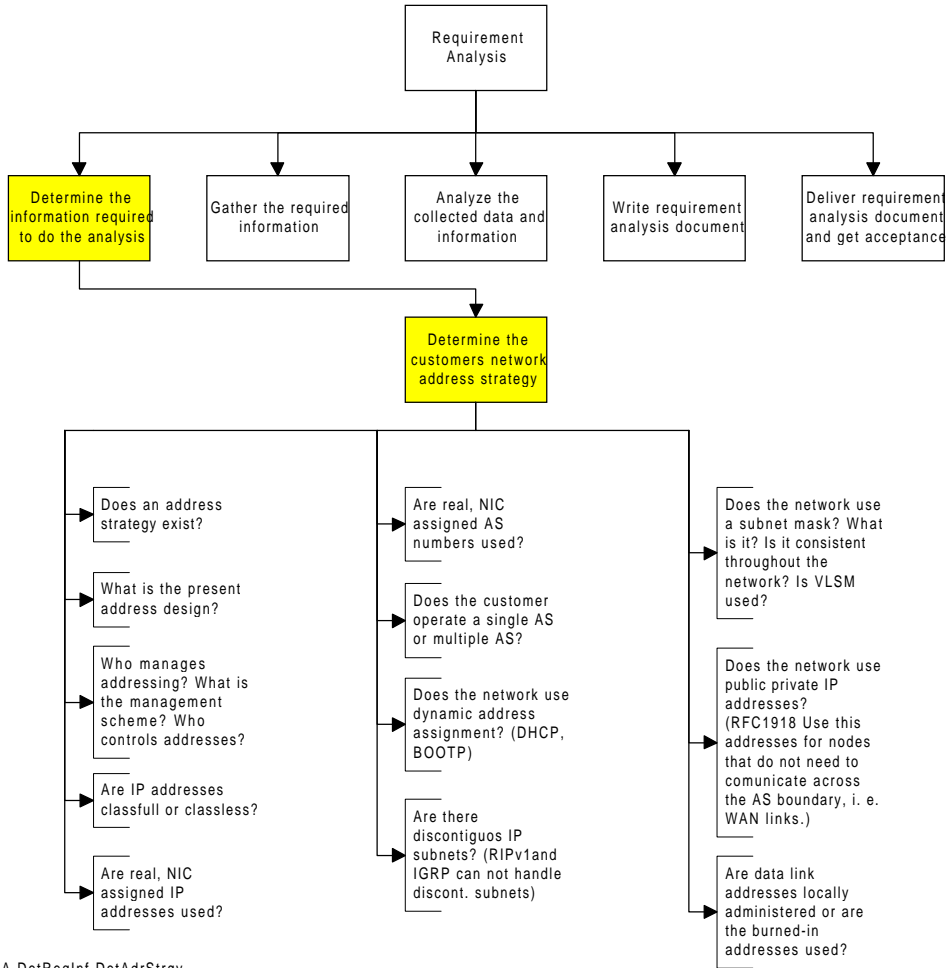


WBS RA.DetReqInf.DetSrvLev

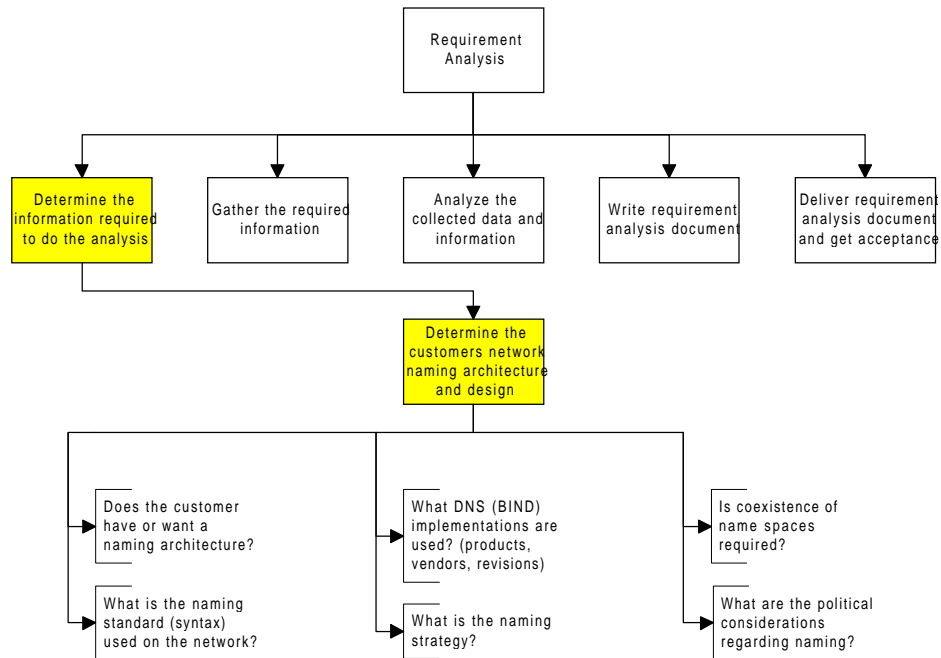
2.5.2.1.1.2 Determine Where The Information Flows In The Company



2.5.2.1.1.3 Determine The Customers Address Strategy

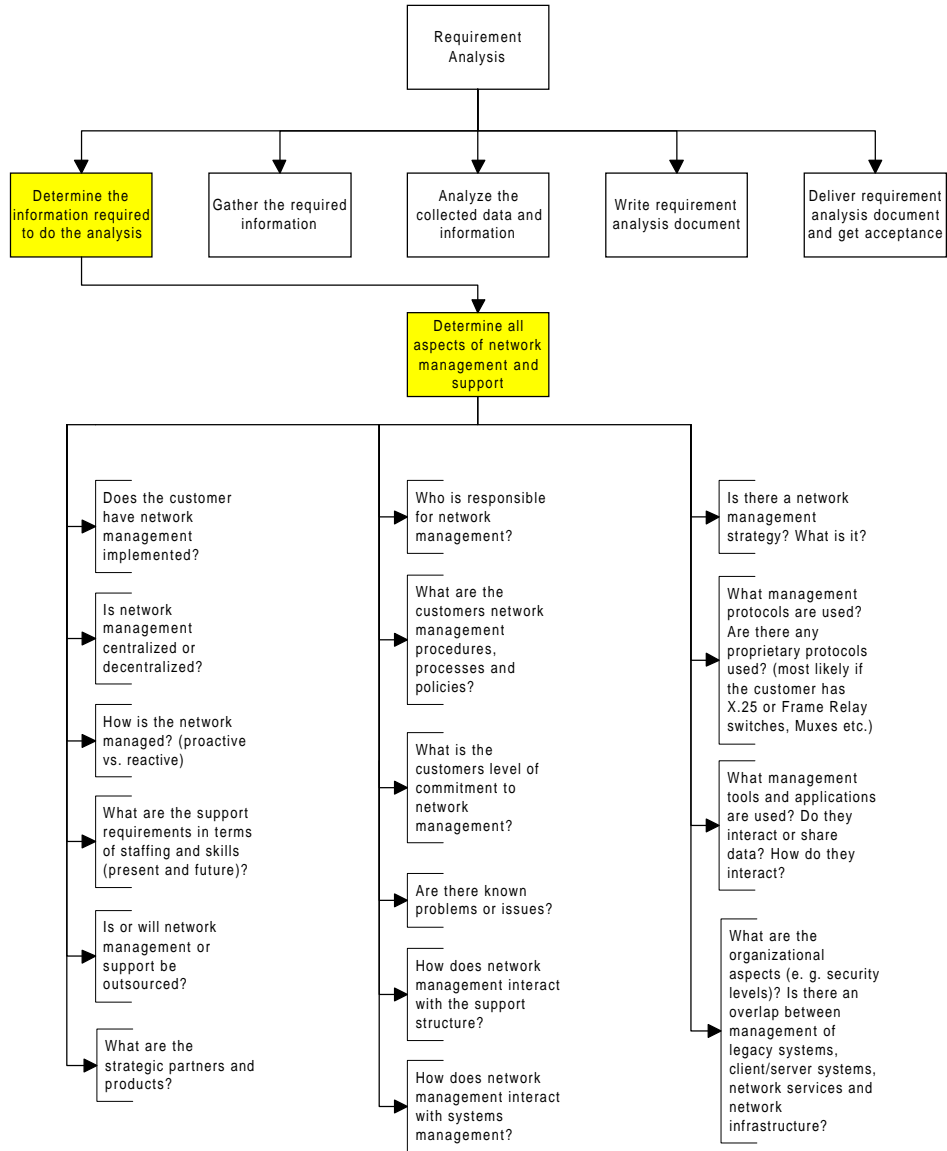


2.5.2.1.1.4 Determine The Customers Network Naming Architecture And Design



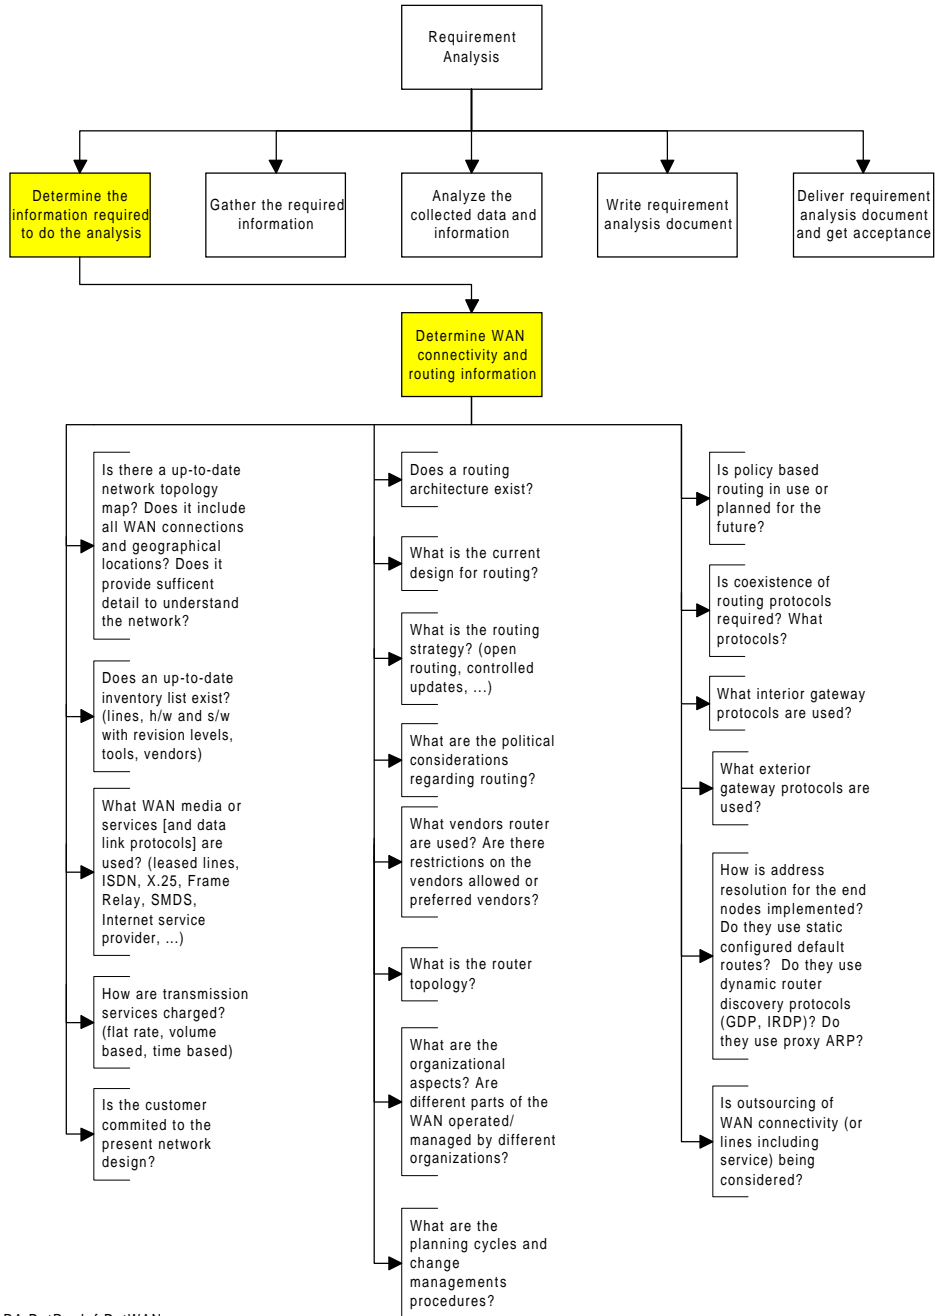
WBS RA.DetReqInf.DetNamArch

2.5.2.1.1.5 Determine All Aspects Of Network Management And Support



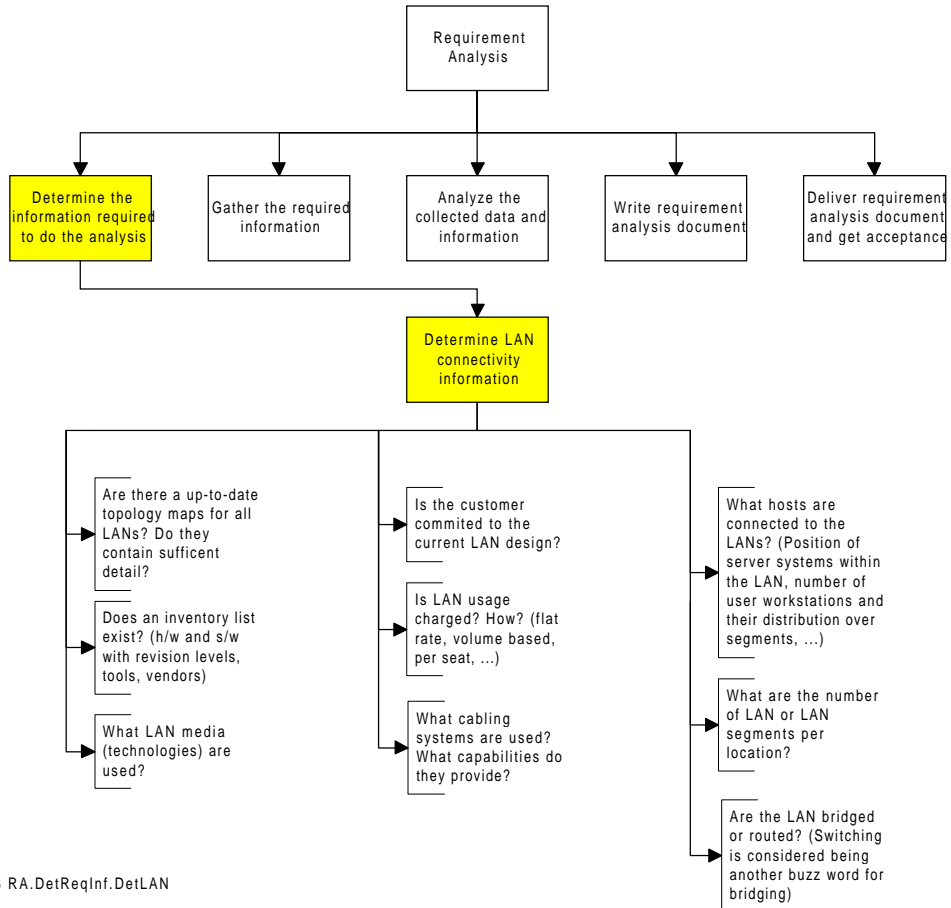
WBS RA.DetReqInf.DetNetMan

2.5.2.1.1.6 Determine WAN Connectivity And Routing Information

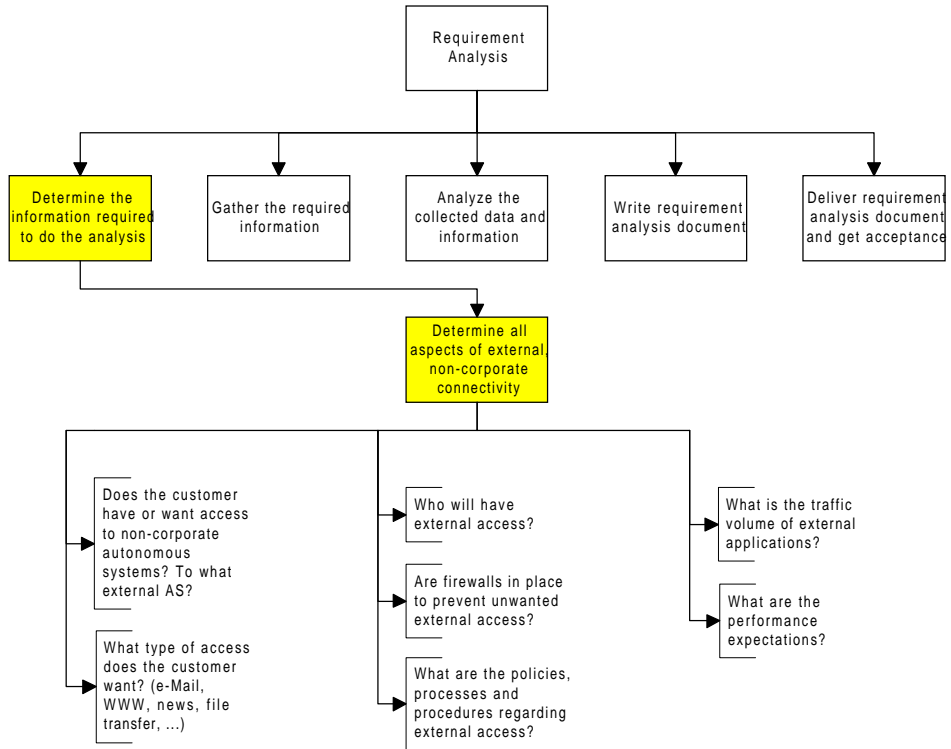


WBS RA.DetReqInf.DetWAN

2.5.2.1.1.7 Determine LAN Connectivity Information

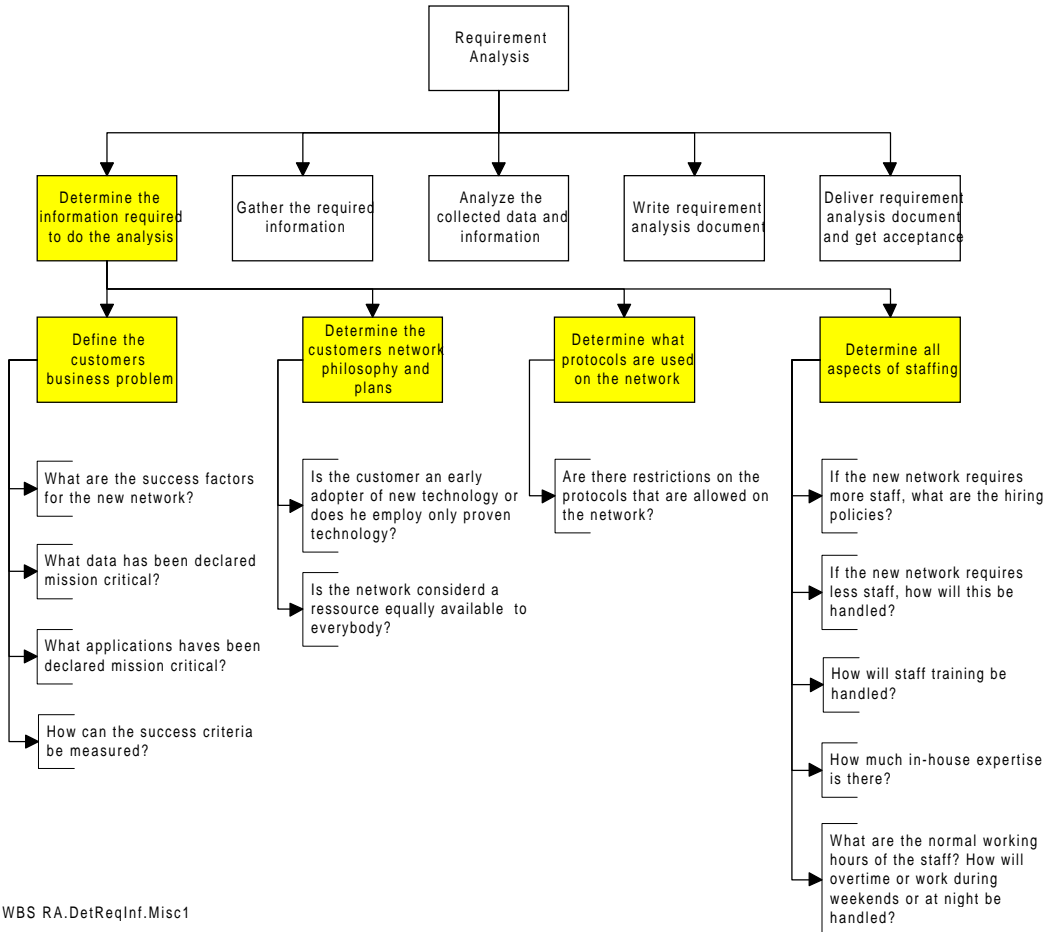


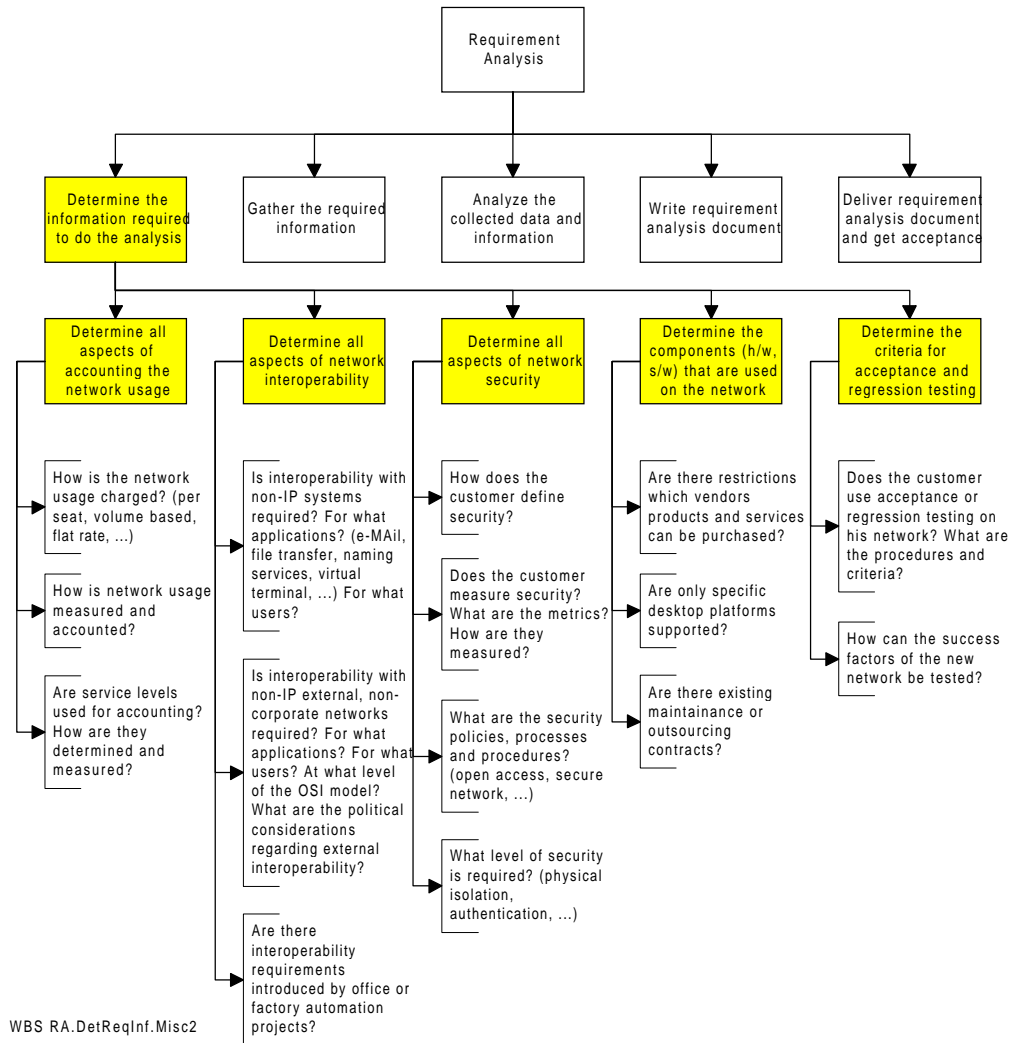
2.5.2.1.1.8 Determine All Aspects Of External, Non-corporate Connectivity



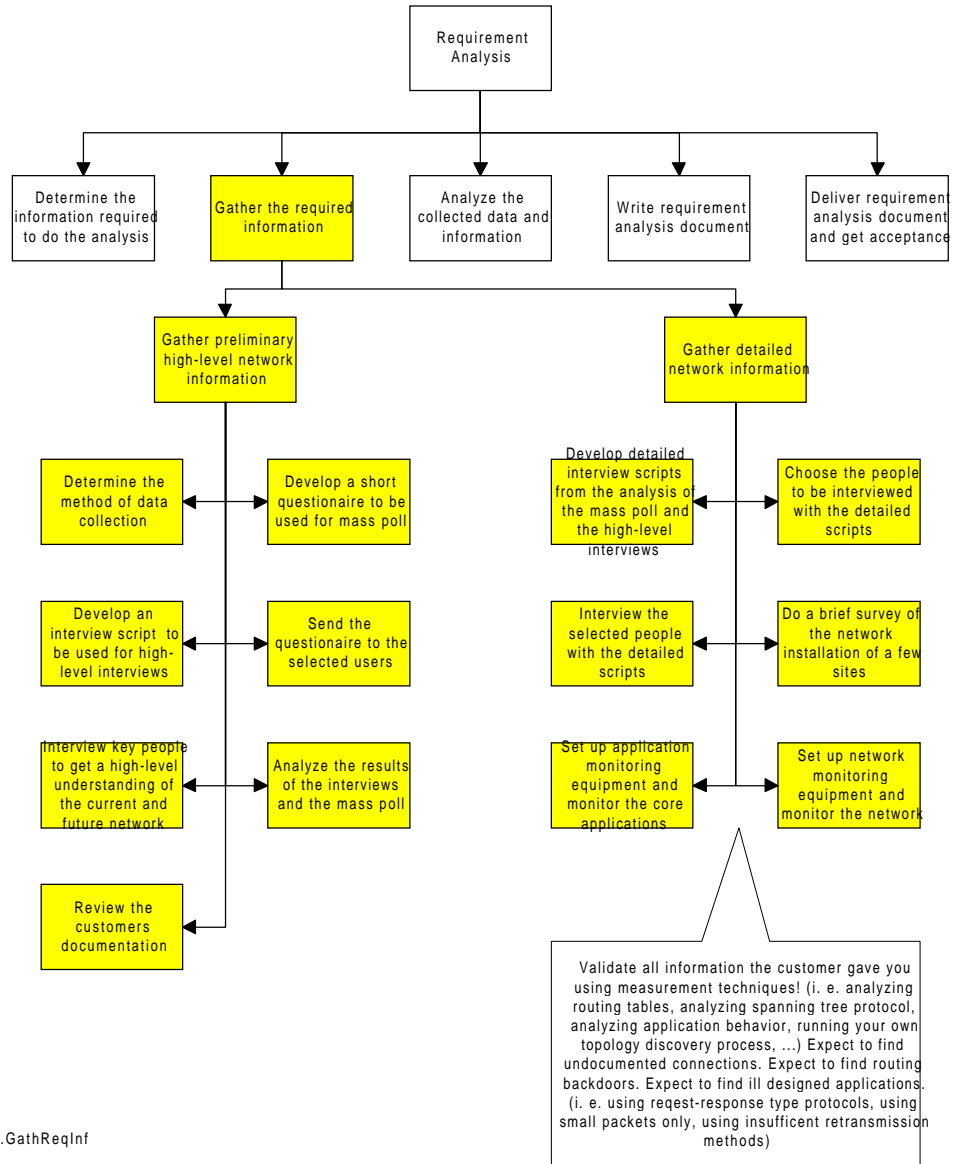
WBS RA.DetReqInf.DetExtCon

2.5.2.1.1.9 Miscellaneous Topics



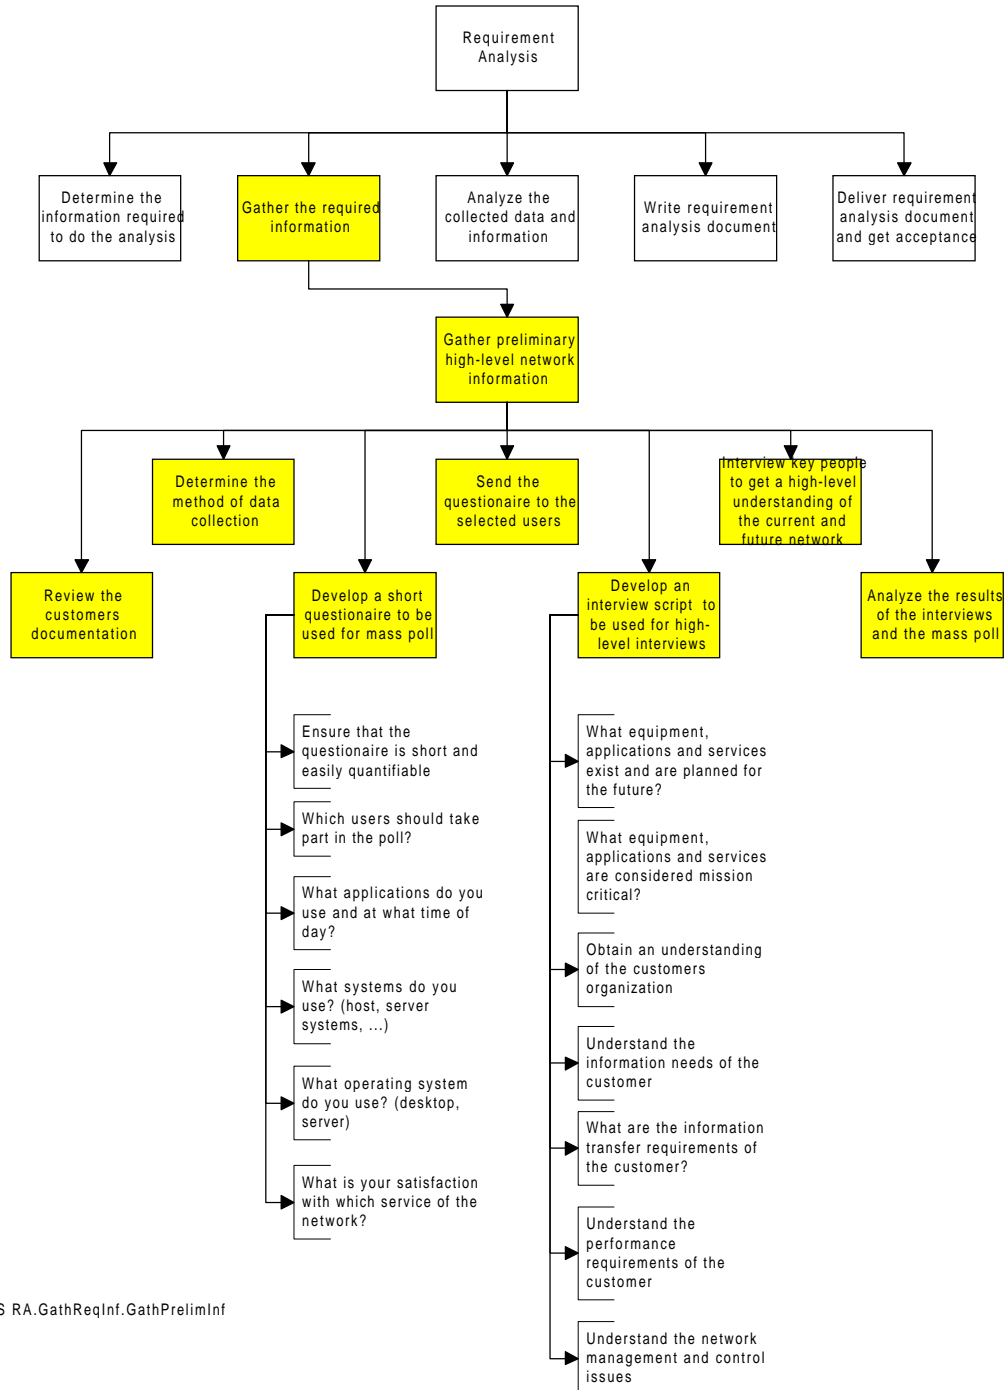


2.5.2.2 Gather The Required Information



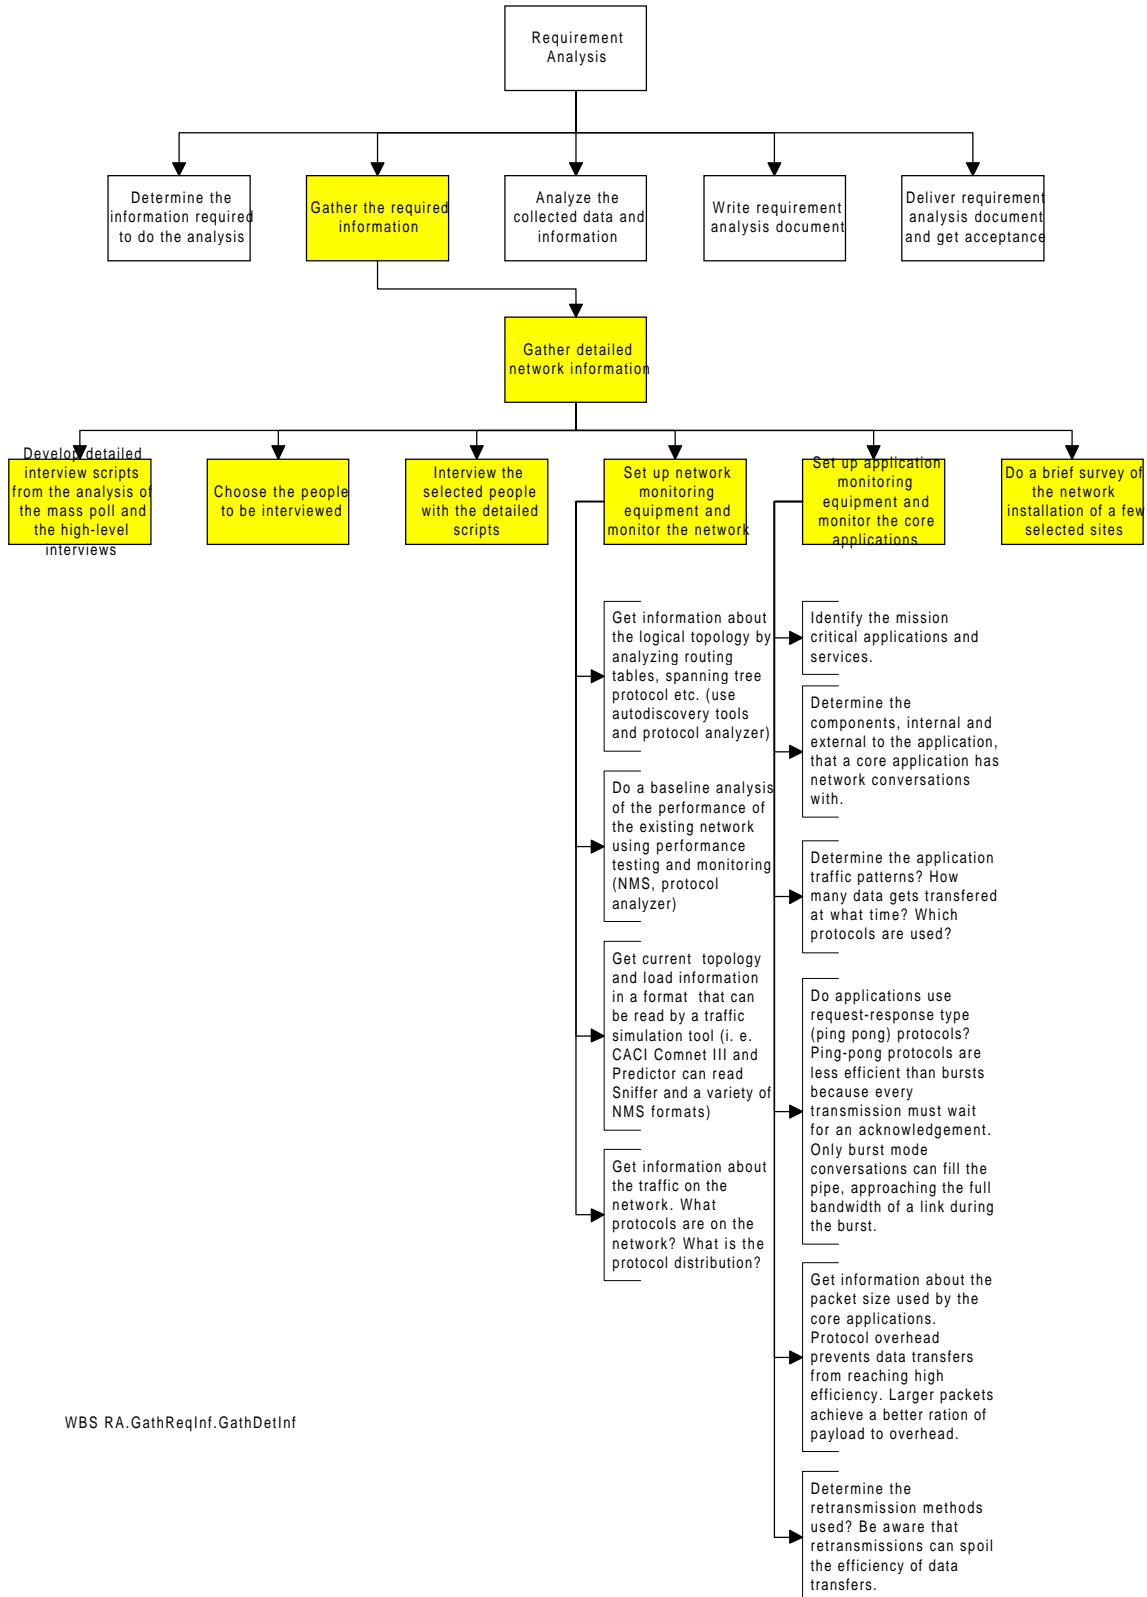
WBS RA.GathReqInf

2.5.2.2.1 Gather Preliminary High-level Network Information



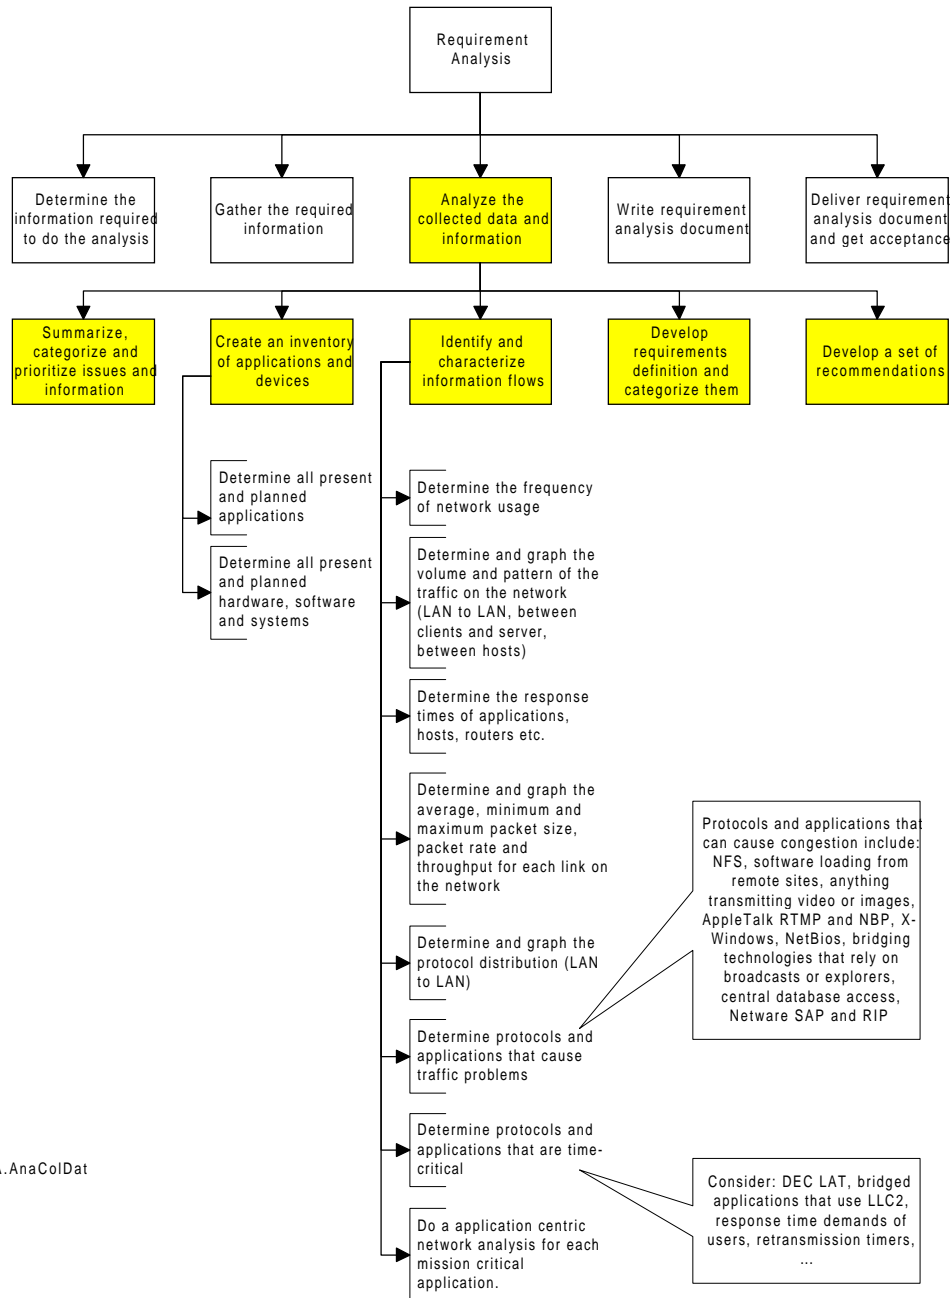
WBS RA.GathReqInf.GathPrelimInf

2.5.2.2.2 Gather Detailed Network Information



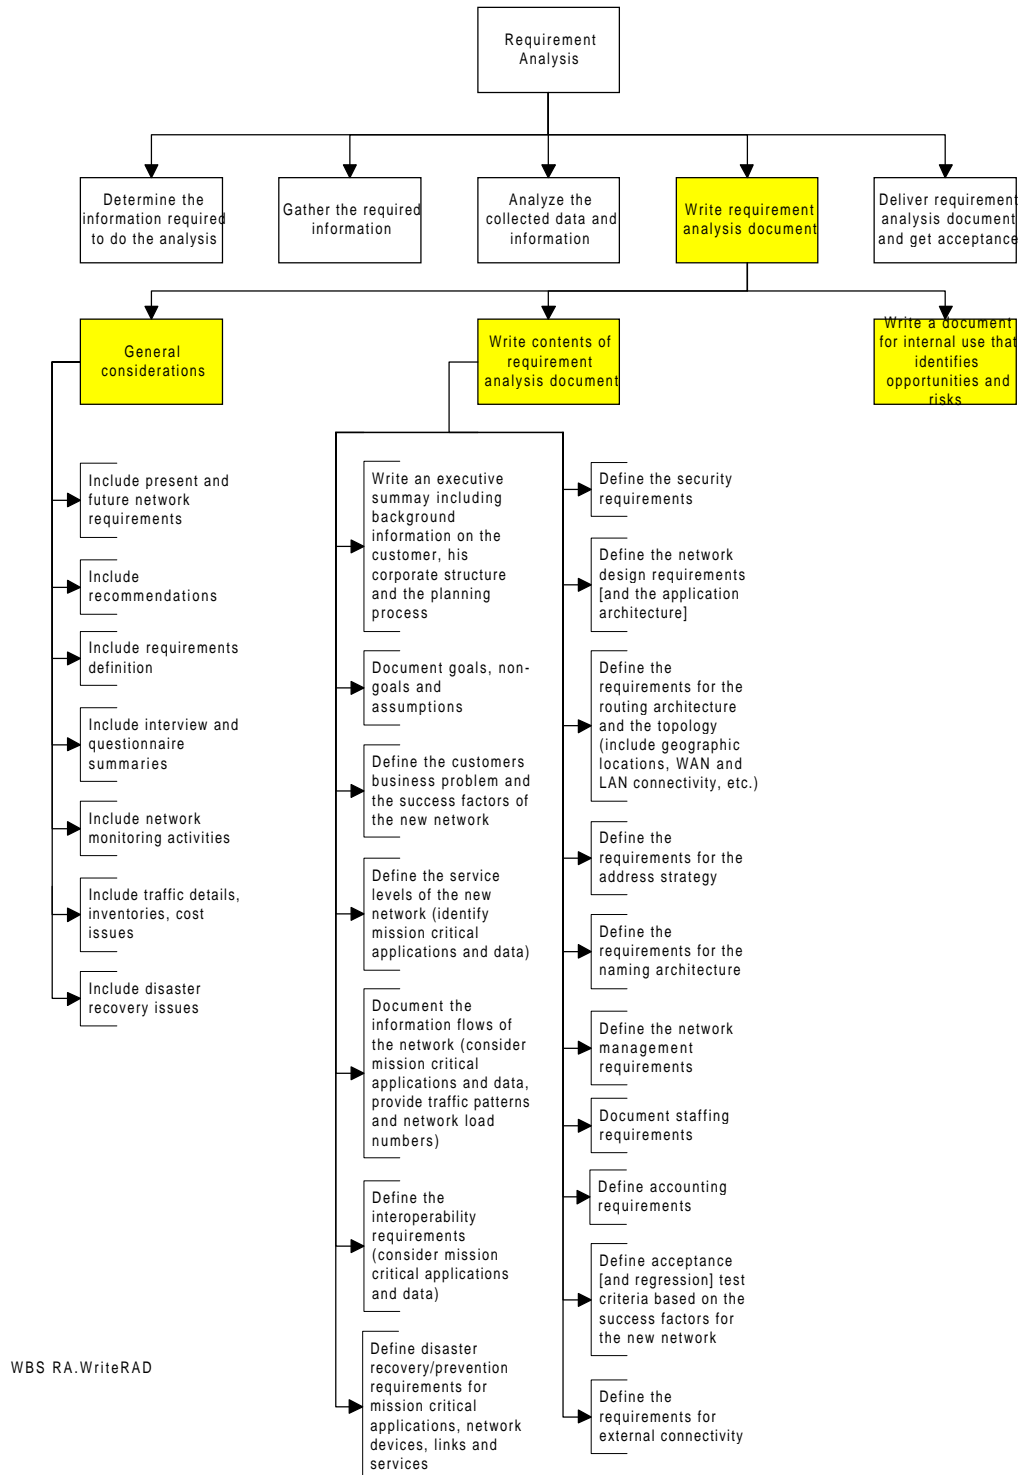
WBS RA.GathReqInf.GathDetInf

2.5.2.3 Analyze The Collected Data And Information

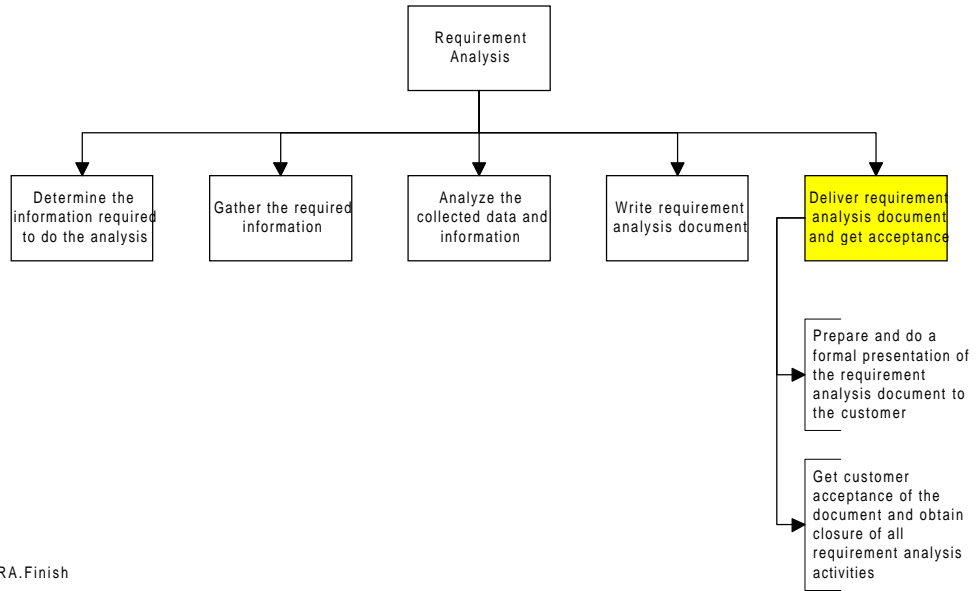


WBS RA.AnaColDat

2.5.2.4 Write Requirement Analysis Document



2.5.2.5 Deliver Requirement Analysis Document And Get Acceptance



WBS RA.Finish

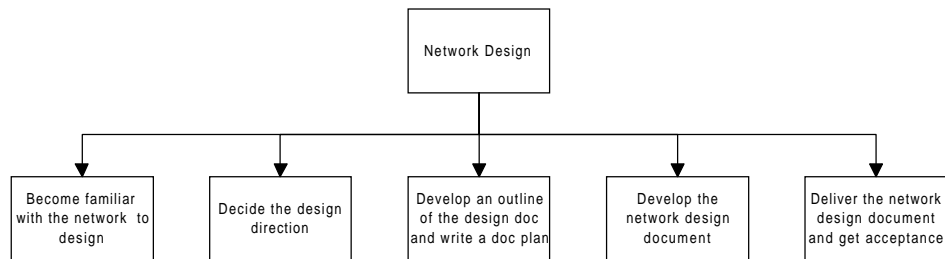
2.5.3 Network Design

„Morte certa, ora incerta“

During the network design you develop a design document according to the requirement analysis, and obtain customer approval of the design.

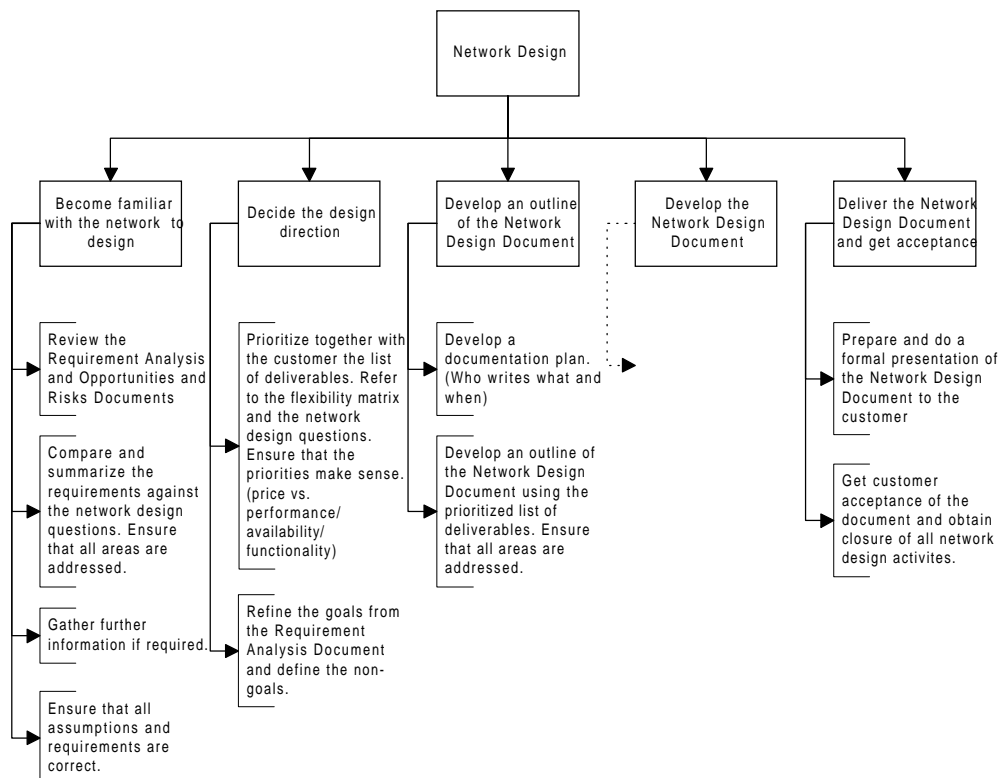
Deliverables

- A network design document for the customer
- A presentation to the customer



WBS ND

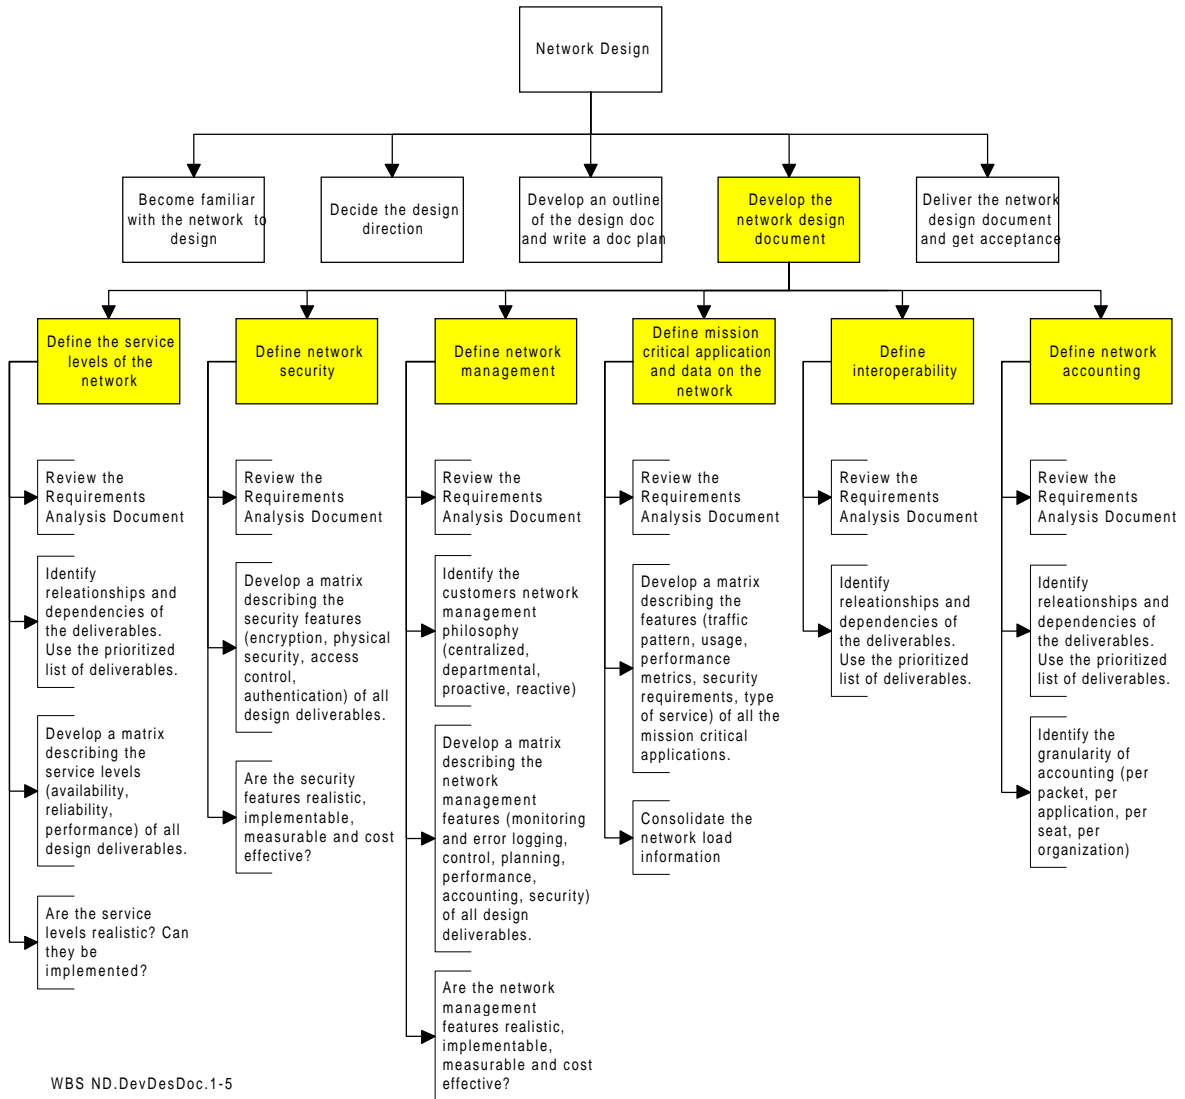
2.5.3.1 Miscellaneous Network Design Tasks



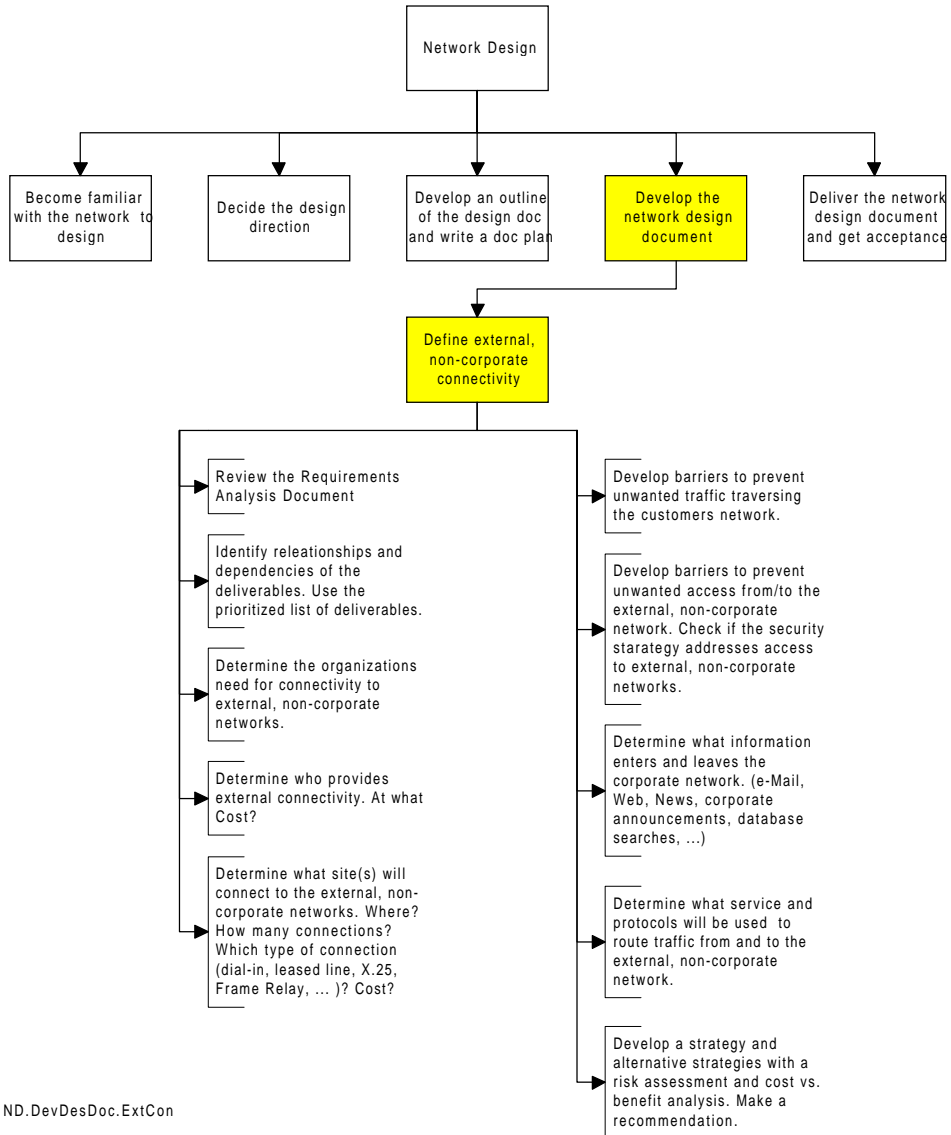
WBS ND.Init+DecDesDir+Outline+Deliver

2.5.3.2 Develop The Network Design Document

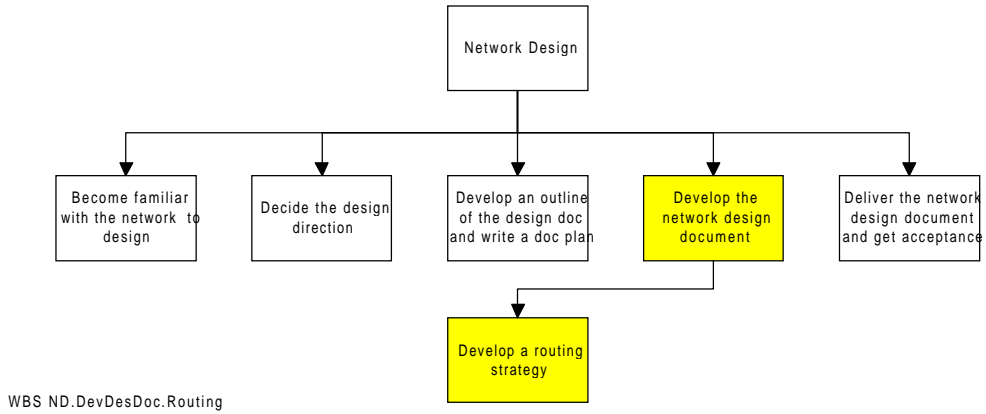
2.5.3.2.1 Miscellaneous Document Development Tasks



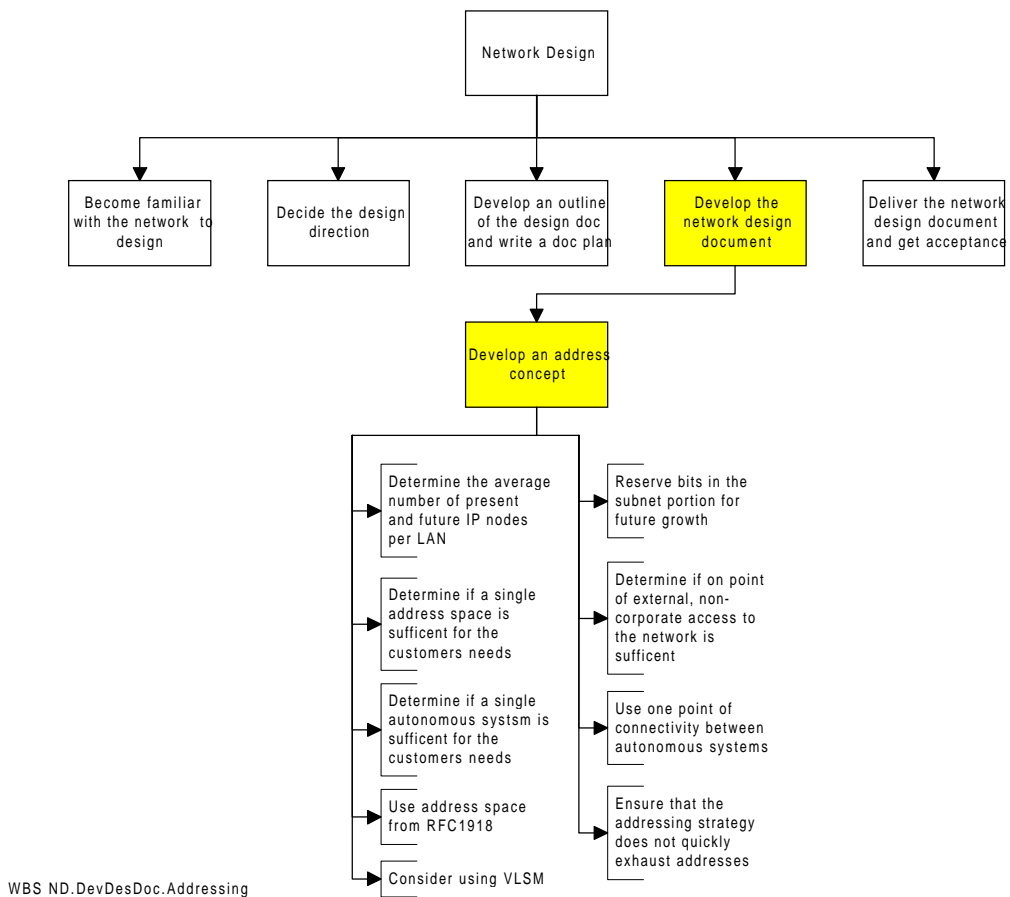
2.5.3.2.2 Define External, Non-corporate Connectivity



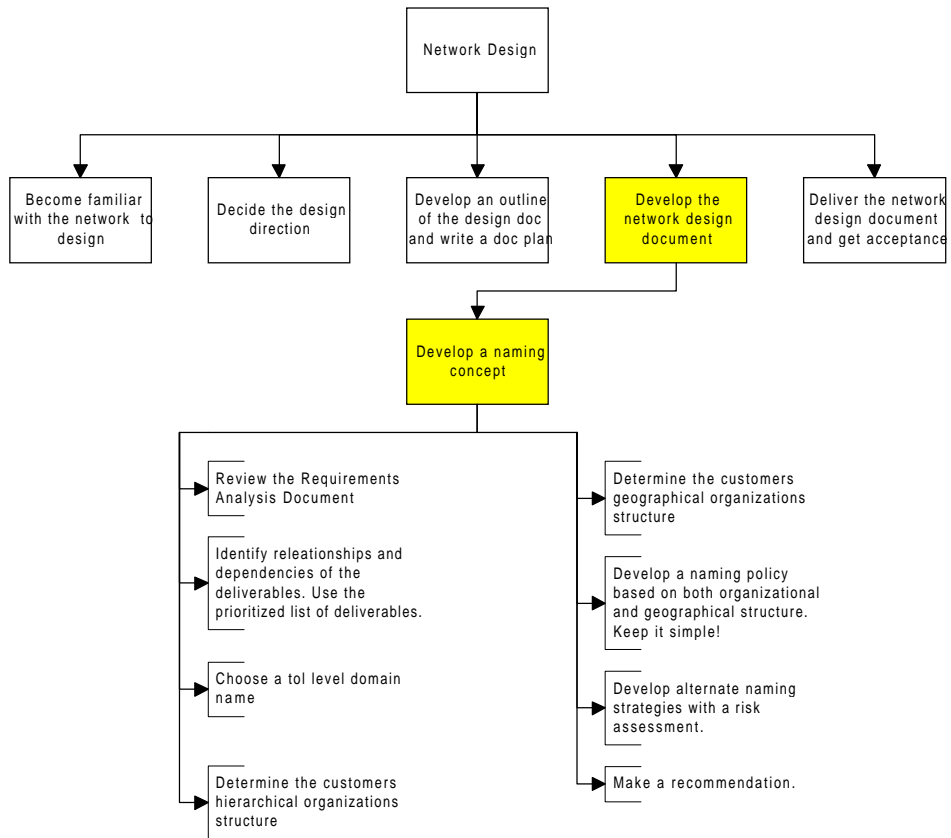
2.5.3.2.3 Develop A Routing Strategy



2.5.3.2.4 Develop An Address Concept

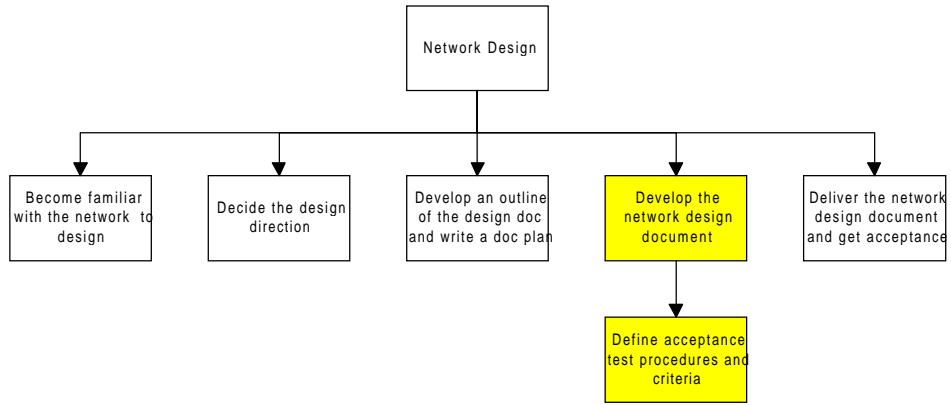


2.5.3.2.5 Develop A Naming Concept



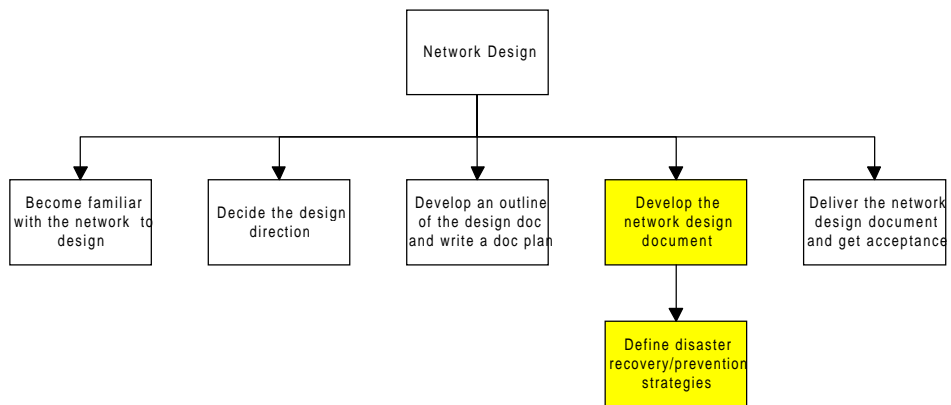
WBS ND.DevDesDoc.Naming

2.5.3.2.6 Develop A Test Concept



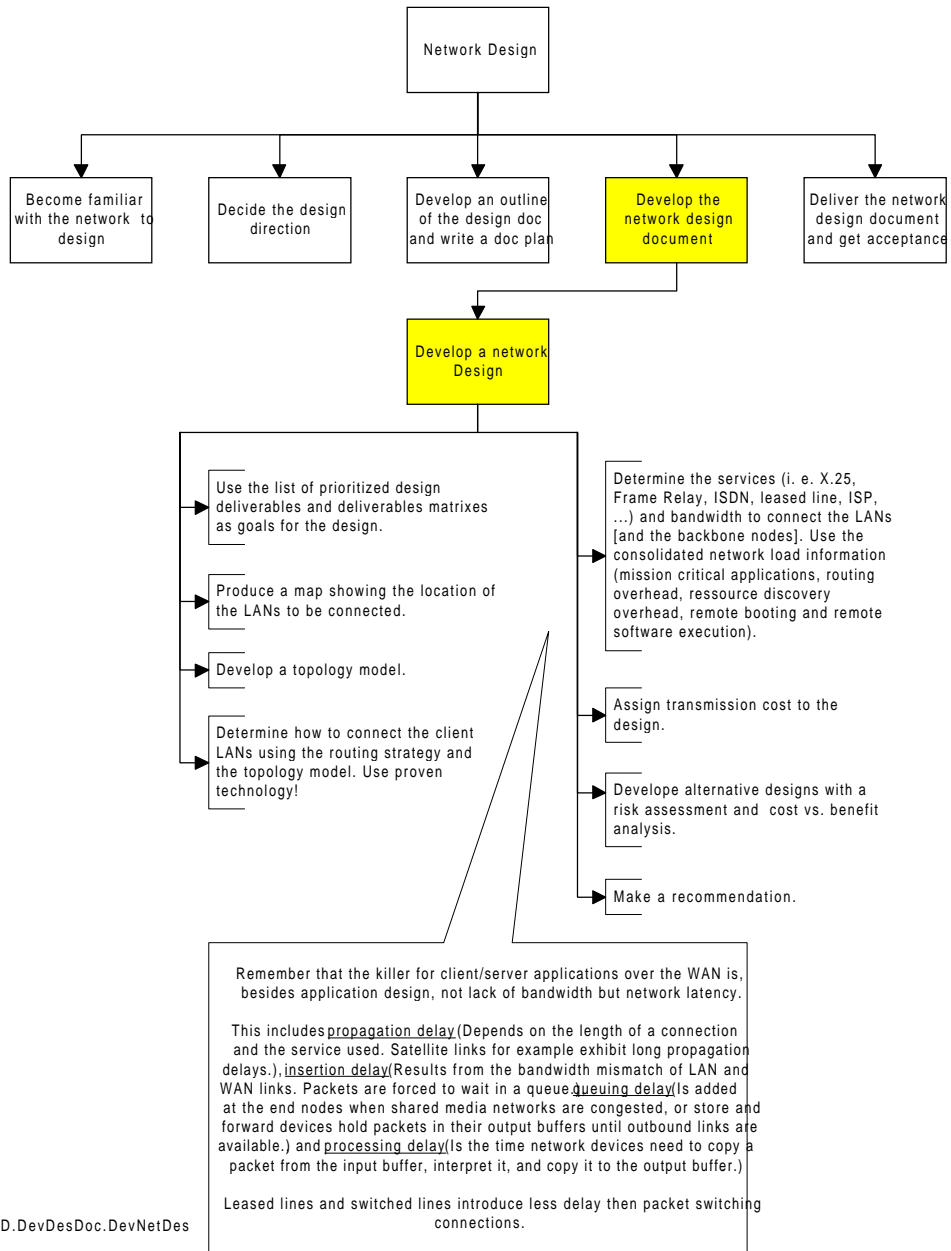
WBS ND.DevDesDoc.Test

2.5.3.2.7 Develop A Disaster Recovery/Prevention Concept



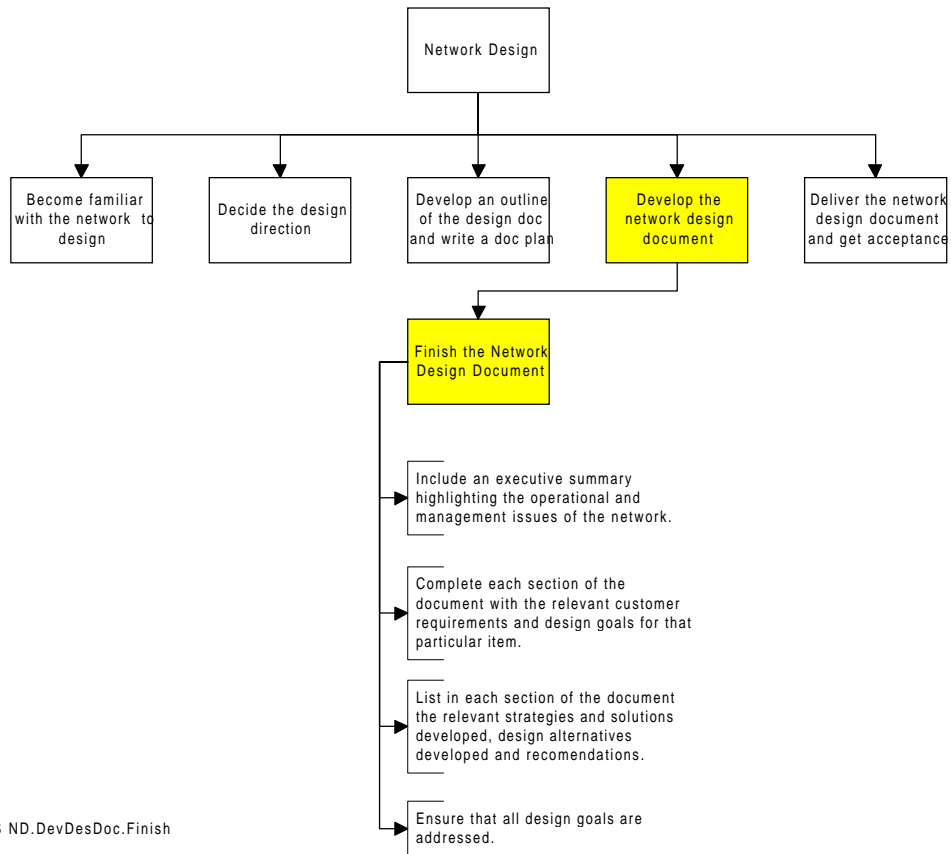
WBS ND.DevDesDoc.Disaster

2.5.3.2.8 Develop A Network Design



WBS ND.DevDesDoc.DevNetDes

2.5.3.2.9 Finish The Network Design Document



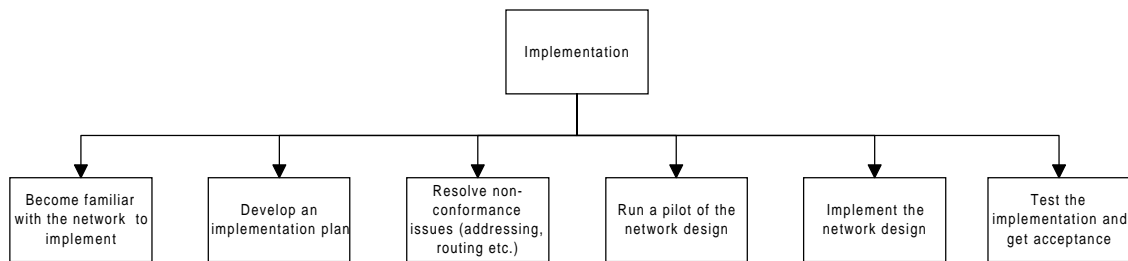
WBS ND.DevDesDoc.Finish

2.5.4 Implementation

During the implementation phase you develop an implementation plan for the network design, and implement it.

Deliverables

- An implementation plan
- An acceptance test report



WBS Imp

2.6 Precedence Network & Resource Requirements

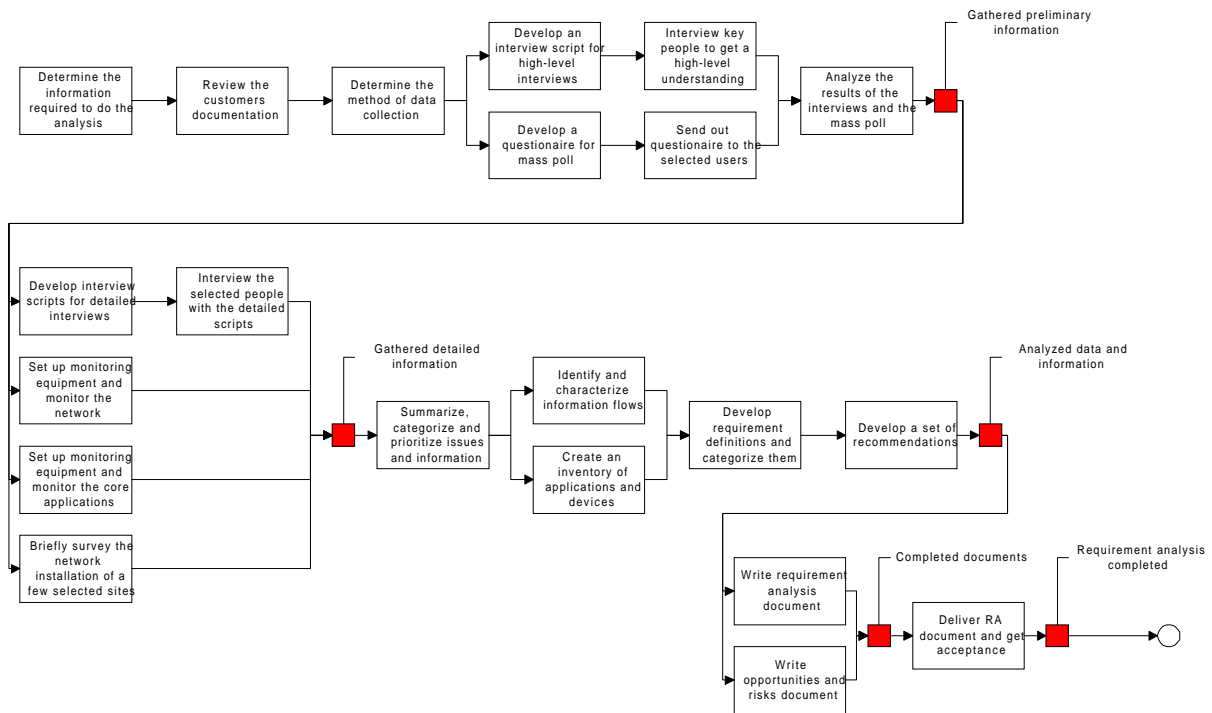


Figure 1, CPN Requirement Analysis

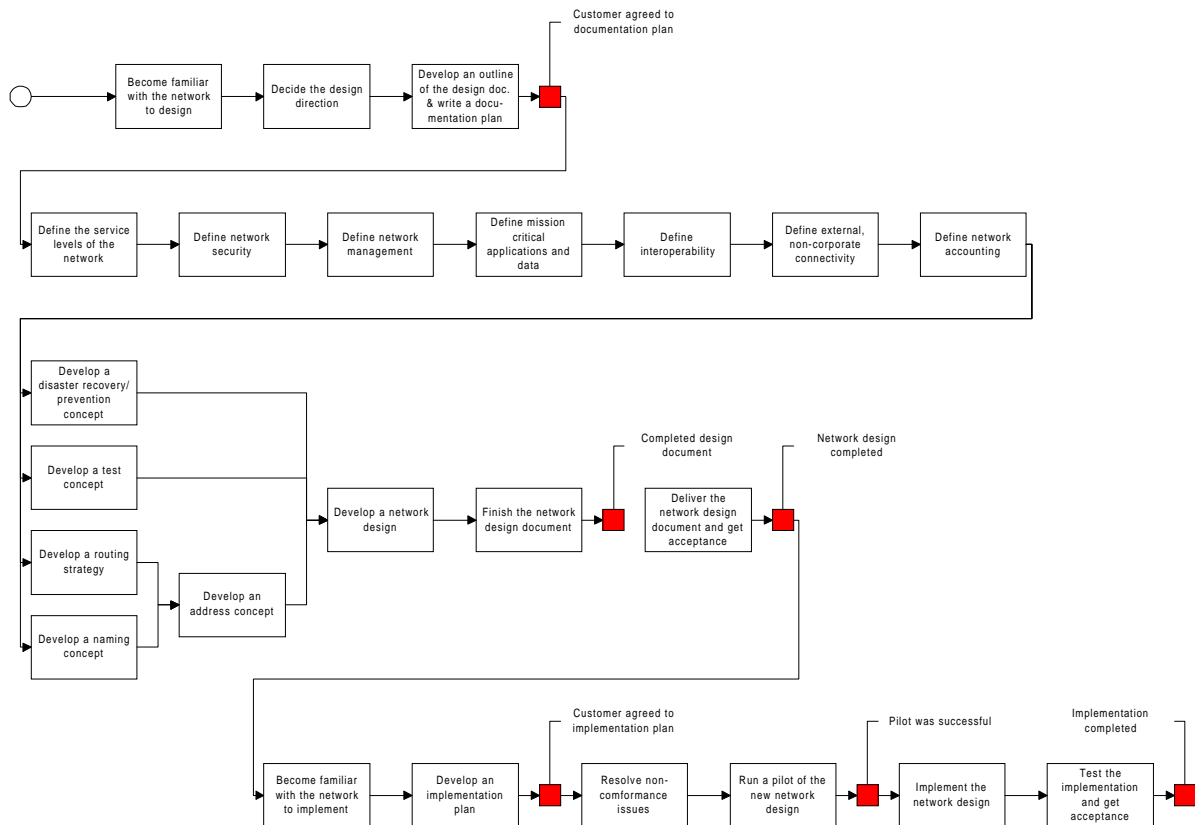


Figure 2, CPN Network Design and Implementation

2.6.1 Skills And Experience Required

Analyst

- >5 years networking experience
- Broad theoretical and practical knowledge of networks and technology (TCP/IP, desktop protocols, LAN and WAN technologies, operating systems and network operating systems, PCs, ...)
- Proven project management skills
- Writing and presentations skills
- Good interviewing skills
- Ability to communicate at all levels
- Analytical skills
- Ability to stay objective and unbiased
- Ability to understand the customers business environment („see the big picture“)
- Experience in the delivery of requirement analysis
- Ability to deliver innovative but feasible concepts and solutions
- Ability to understand network traffic and use traffic analysis tools

Designer

- >5 years networking experience
- Good theoretical and practical knowledge of networks and technology (TCP/IP, desktop protocols, LAN and WAN technologies, operating systems and network operating systems, PCs, ...)
- Good knowledge of network routing and addressing
- Good knowledge of naming
- Good knowledge of network performance
- Good knowledge of network security
- Knowledge of traffic analysis and use traffic analysis and simulation tools
- Knowledge of network management tools and protocols
- Knowledge of network products
- Knowledge of carrier services
- Knowledge of the IP/Internet culture
- Previous experience of managing and implementing networks and networked systems
- Writing and presentations skills
- Ability to communicate at all levels

Implementers

- >3 years networking experience
- Broad theoretical and practical knowledge of networks and technology (TCP/IP, desktop protocols, LAN and WAN technologies, operating systems and network operating systems, PCs, ...)
- Good knowledge of network routing and addressing
- Good knowledge of naming

- Knowledge of network performance
- Knowledge of network security
- Writing and communication skills

3 Best Practices

3.1 Topology

Design rules within the section Topology are given within subsections. This does not imply that the rules apply only to that particular subsection. Unless noted otherwise design rules apply globally. The rank of a rule within the sequence does not imply a priority.

At this stage of the design you lay out the overall structure of the network and choose the internetworking technology.

Use modular and hierarchical models to simplify complex design problems.

In order to design a scalable network it is important to choose a simple, hierarchical structure. Rules and models simplify complex design problems to a collection of smaller, manageable parts. Use the hierarchical model consistently for better predictability of latency, performance and routing, and easy troubleshooting.

The hierarchical model provides a physical topology of the network. Specialization is a feature of hierarchical models. Different router features are implemented at specific layers.

The simpler the topology, the simpler the implementation.

It is virtually impossible to build large networks without imposing a hierarchy. Hierarchical networks scale well!

3.1.1 Flat Model

Using only bridges and switches is internetworking without hierarchy. This is known as a flat network or a flat address space.

Bridges and switches are transparent to routers. They do not organize address space.

The scalability of a flat network is limited by the amount of broadcast traffic. A flat network, VLAN is another example, forms a single broadcast domain.

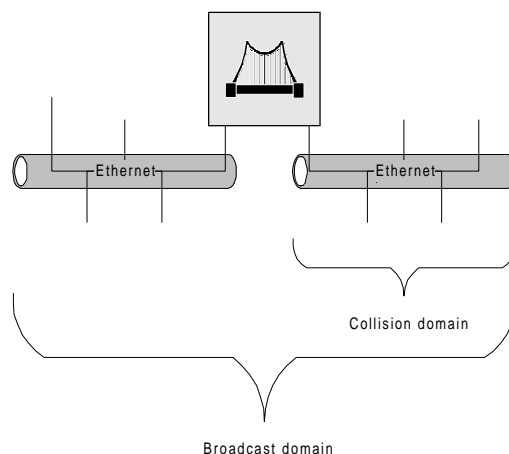


Figure 3, Flat model

Flat models can be used to build workgroup networks.

3.1.2 One-layer Model

Use routers to terminate broadcast domains.

In a one-layer model routers can be added to without increasing the number of logical layers. All routers are considered peers if they are connected to a common backbone. This is typical when the backbone is in the riser of a building with routers in the satellite equipment rooms of each floor. Client LANs connect the computers on the floor to the router in the wiring closet.

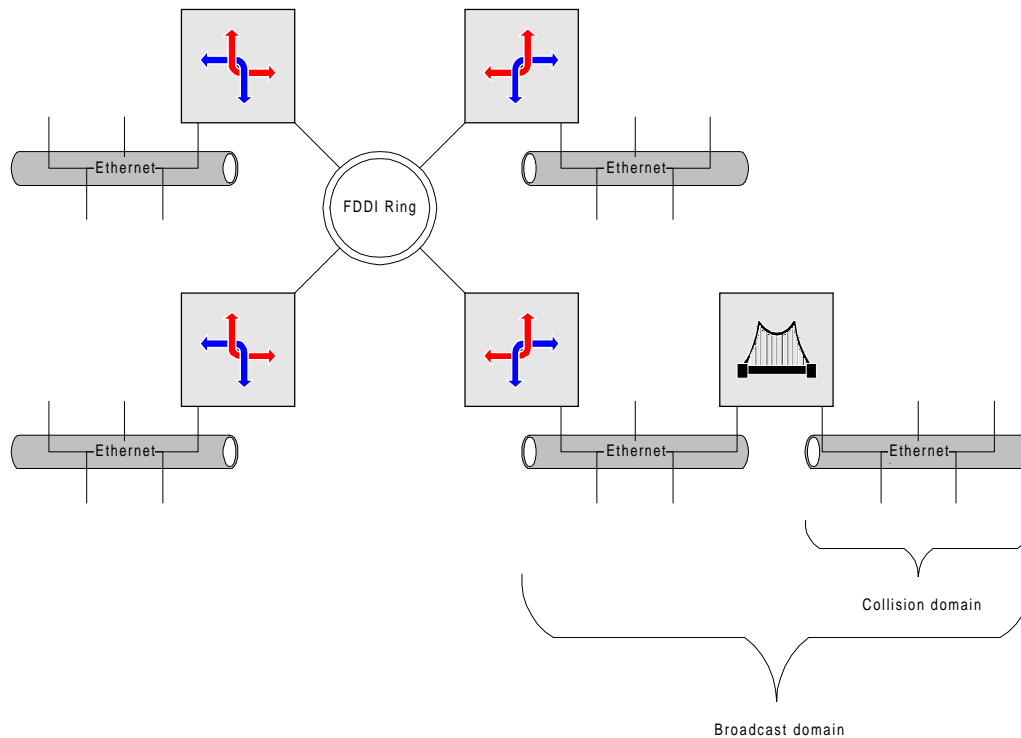


Figure 4, One-layer model

Use router features to make efficient use of WAN links.

Another one-layer model has a WAN backbone instead of the LAN backbone. Here routers are used to provide features to improve WAN performance and efficiency. These features include compression, bandwidth reservation and proxy services. Apply these features to the WAN ports as this has less impact on performance than it would have on the LAN ports.

Carefully consider the performance impact of value-added features on the routers core function, path-determination and packet forwarding.

3.1.3 Two-layer Model

In the two-layer model, core routers provide access to the core WAN, and access routers provide client LANs access to the corporate network.

Core routers use value-added features to improve WAN performance and efficiency. Apply these features to the WAN ports as this has less impact on performance than it would have on the LAN ports.

Again: carefully consider the performance impact of value-added features on the routers core function, path-determination and packet forwarding.

The network diameter should be as small as possible and consistent throughout the network.

It is important to restrict the diameter of the network to achieve low delay and fast routing convergence. The

diameter must be consistent for predictability of latency, performance and routing, and easy troubleshooting.

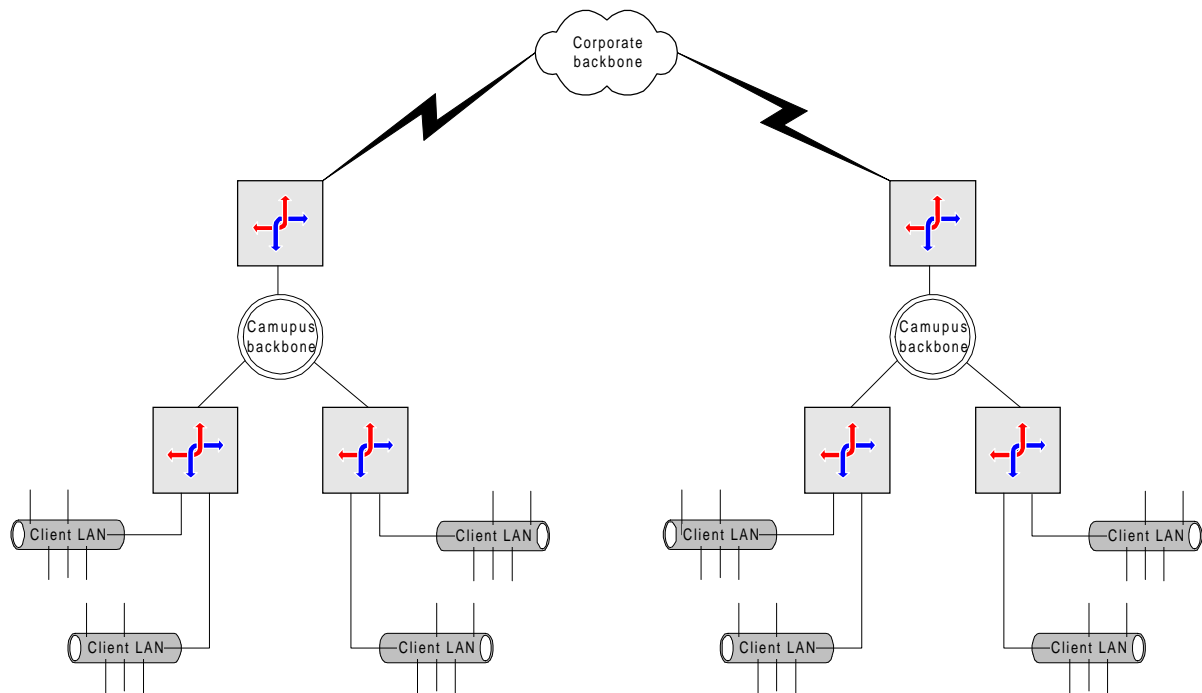


Figure 5, Two-layer model with WAN backbone

Core LAN redundancy is less expensive than core WAN redundancy.

Another two-layer model has a LAN backbone instead of the WAN backbone. The core LAN backbone is duplicated with total redundancy. The WAN links to the remote sites are duplicated to different core routers. This configuration provides cost-effective redundancy.

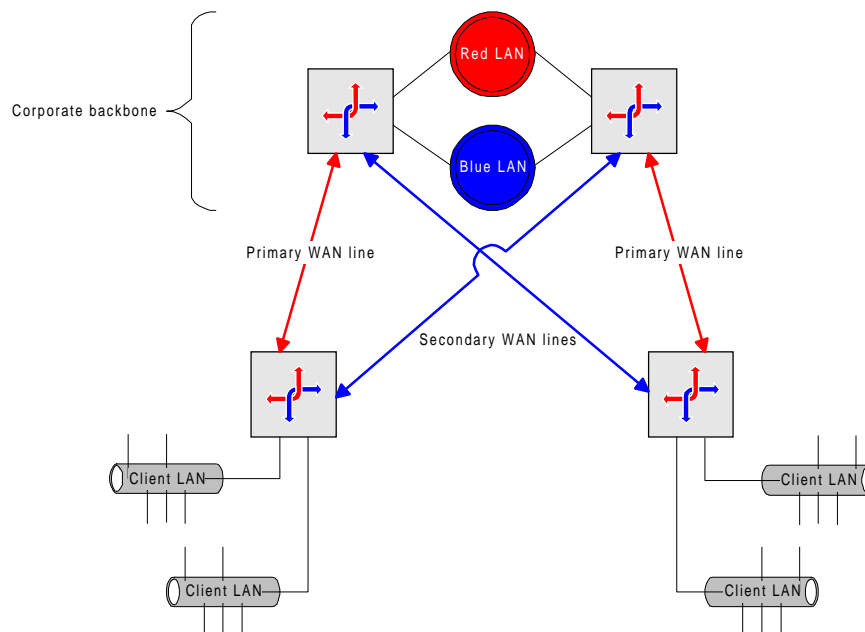


Figure 6, Two-layer model with LAN backbone

3.1.4 Three-layer Model

This model uses three distinct layers: the core layer, the distribution layer and the access layer. This model is used to build very large-scale networks. The model permits traffic aggregation and filtering at three successive routing levels.

The function of the core layer is to provide an optimized transport structure. Including dynamic load sharing across the core structure, efficient and controlled use of bandwidth, optimized connectivity between distribution networks, a defined and consistent network diameter, predictable traffic patterns, and robust connectivity. The core layer is a high-speed switching backbone and should be designed to switch packets as fast as possible. This layer should not perform any packet manipulation such as packet filtering.

At the distribution layer are policies implemented to control access to resources. This is the layer where packet manipulation should take place. It provides policy-based connectivity and address or area aggregation, serves as redistribution point between routing domains and the demarcation between static and dynamic routing. It can be the point at which remote sites access the corporate network.

Access routers provide access for workgroup LANs, and control traffic by localizing broadcasts and service requests. Functions include segmentation into networks or subnetworks, proxy service, router discovery, filtering and connectivity on demand. It can be the point where remote sites access the corporate network via wide-area technology.

The three layers do not have to exist in clear and distinct physical entities. The layers are defined to aid network design and to represent functionality that must exist in a network. The instantiation of each layer can be in distinct routers or switches can be represented by a physical media or can be combined in a single device.

Provide sufficient capacity where traffic aggregates.

Concentration of traffic from lower layers must be taken into consideration when planning core bandwidth.

Design redundancy at the core layer.

Router distribution varies from a few in the core to many at the access layer. Core routers are mission critical. If a core router fails many client LANs are affected. Design the core for maximum availability. It should tolerate link failures and still provide connectivity.

Use an even number of routers for fast convergence and load balancing.

Design connectivity with an even number of meshed or partial meshed routers. The odd number of routers will prevent equal cost load balancing. If you have parallel equal cost paths both are installed in the routing tables. In case of failure of one path the alternate path is available at instance because it is already in the routing tables. If you have unequal cost parallel paths the routing protocol must converge before the longer path can be used.²

It is important to restrict the diameter of the network to achieve low delay and fast routing convergence. The diameter must be consistent for predictability of latency, performance and routing, and easy troubleshooting.

Keep bandwidth consistent within a given layer of the hierarchy.

Keep bandwidth consistent within a layer of the hierarchical model to provide load balancing and improve routing performance. Routing protocols converge much faster if multiple equal-cost paths to a destination network exist.

Implement routing policies at the distribution layer.

² All the time consuming mechanisms that prevent routing loops, for example counting to infinity or path hold-down, are executed before the alternate path becomes available.

Policies control access to resources. This includes access to parts of the internetwork, and access to services of the internetwork. Policies can be implemented using routing metrics and packet filters.

Limit upstream traffic at the access layer.

Use static routes to save bandwidth when a part of the internetwork can be reached only via one path. If you can not use static routing, use routing protocols that advertise only changes. Consider using value-added features like Cisco's snapshot routing.

Use static service advertisement on the access router if possible.

Use local proxy service on the access router instead of forwarding broadcasts if possible.

Move workgroup servers down to the workgroup if possible.

Collect related workgroups at the access layer.

Cluster similar LAN technology at the access layer.

Do not design backdoor routes or chains of access routers.

Do not violate the hierarchical model at the access layer. Do not attach workgroups by adding routers below the access layer. Do not allow backdoor routes between workgroups. Backdoor routes cause unpredictable, non-deterministic routing and make troubleshooting difficult.

3.2 Routing

Instead of writing a brief and incomplete introduction to routing protocols I would rather like to point you to some excellent books: „Interconnections“ by Radia Perlman, „Routing in the Internet“ by Christian Huitema, and „Internet Routing Architectures“ by Bassam Halabi.

Please note that the rank of a design rule within the sequence of rules does not imply a priority.

Use a gateway discovery protocol to support redundancy.

Hosts are often configured with a default gateway. This is a static configuration that fails if the router fails.

Alternatively hosts can run the Gateway Discovery Protocol (GDP) or ICMP Router Discovery Protocol (IRDP) to dynamically determine a gateway router.

Another method is to configure hosts so that they send an ARP request for every destination. Configure the router to respond to any ARP if it has a matching route in its table. This is called proxy ARP. Proxy ARP is not recommended because the router is spoofing when it responds on the behalf of a host that may or may not be active.

Running silent RIP (routed -q) on the hosts is not recommended either. Routers would have to inject RIP broadcasts into the client LANs. If a routing protocol other than RIP is used, route redistribution into RIP is required. Being a classfull routing protocol RIPv1 will cause you pain in environments that use different subnet masks or VLSM.

Routing is an overhead activity. It consumes bandwidth, CPU and memory resources. The routing protocol, its timers, and the use of static routes determine the bandwidth overhead.

Use static routing to stub networks.

A part of the internetwork that can be reached by only one path is called a stub. Static routes are preferred to connect stub networks because they reduce routing overhead.

Hey, do you know what „routing by propaganda“ and „routing by rumors“ is all about. Check out the recommended books if you are not sure.

Hierarchical network design dictates a diameter of up to six hops.

With distance vector protocols convergence time is a function of the network diameter and its complexity. Routing information must propagate from one edge of the internetwork to the other. The network diameter should be kept small and consistent.

Plan well for scalability! Use summarization!

Carefully plan link-state internetworks. Scalability of link-state networks is a complex problem. It is affected by the number of routing nodes in an area, the number of networks in an area, the number of areas, the design of the address space, the effective use of summarization, and the stability of the links. Keep your areas small! As a rule of thumb do not plan to have more than 50 to 100 routers within an area and not more than 5 to 10 neighbor routers.

Routing scales if routing information consolidates. Route summarization is important for all routing protocols. Without summarization there is a flat address space with a specific route to every single subnet. Access routers summarize host routes into subnetworks. Subnets are grouped into major networks. Subnets and networks can be collected in areas. Networks can be grouped into autonomous systems.

Carefully consider this while developing the addressing scheme!

Consider convergence time.

Convergence time is the time it takes for routers to get a consistent understanding of the network topology after a change. Packets may not be reliably routed to all destinations until convergence takes place. Convergence is a critical design constraint for time-critical applications and protocols, for example SNA over IP (DLSw).

Convergence has two components. The time it takes to detect the link failure and the time it takes to determine and install a new route.

Use multiple equal-cost paths.

The fastest convergence takes place with load balancing. If the routing table has multiple paths to a destination, all traffic will be immediately routed over the remaining path.

3.2.1 OSPF

Do not use OSPF virtual links.

Virtual links are paths that connect the backbone area through another area. They act like tunnels, which maintain the contiguity of the backbone when failures occur. Virtual links add complexity to the design. Therefore the backbone should be designed with redundancy and multiple-path routing. Virtual links can be seen as a way to repair existing OSPF networks. Any network that is designed to use virtual links is broken by design!

Maximize summarization.

Consider the requirements of OSPF for summarization when you develop the IP addressing scheme!

Configure route summarization between areas at the area border routers. OSPF areas are useless without summarization! Without summarization, every router in the autonomous system recomputes the Dijkstra algorithm when any link changes. This is as if every router were in the same area. This does not scale very well!

Keep the areas small.

Do not use more than 50 to 100 routers within an area.

Keep the backbone diameter small.

Keep the backbone topology simple.

Avoid complex meshed topologies. A LAN backbone design is ideal for OSPF.

Isolate the backbone from intra-area traffic. Use it strictly as a transit area. Put workgroups in areas other than the backbone.

Use stub areas.

3.3 Addressing

Instead of writing just another introduction to IP addresses I would rather like to point you to an excellent document: „Understanding IP Addressing: Everything You Ever Wanted To Know“ by Chuck Semeria. It covers classfull and classless IP addresses, subnetting, Variable Length Subnet Masks (VLSM) and Classless Inter-Domain Routing (CIDR). I think it is the most comprehensive single source on IP addressing and it is free. You can download it from the URL: <http://www.3com.com/nsc/501302.html>.

Please note that the rank of a design rule within the sequence of rules does not imply a priority.

Addressing will be used to encode topological information only!

Addressing will be used to encode topological information only. Names will be used to encode organizational information. Do not encode organizational information in addresses! Use carefully structured addressing and naming conventions to make filtering, growth and faultfinding easier.

Use hierarchical addressing schemes.

It is impossible to build large networks without imposing a hierarchical structure.³ The structure of IP addressing suggests a hierarchy⁴. Routers operate in hierarchical address spaces. Only hierarchical networks scale well!

Plan your addressing strategy in a way that supports route aggregation. Summarization requires that multiple IP network addresses share the same high-order bits. The addresses can then be aggregated into a smaller number of routing table entries. Route summarization is important for all routing protocols. Hierarchical topologies and routing architectures are useless without address aggregation.

Use addresses from the private network address space (RFC 1918) for hosts that do not need to access the Internet.

The Internet Assigned Numbers Authority (IANA) has reserved blocks of IP addresses for private networks. A Company that uses IP addresses from that address space can do so without coordination with IANA. Many companies can use the address space. Addresses are only unique within one company.

Use the private network address space to address all your WAN connections. In most cases it is preferred to use the private address space for all nodes and provide Internet access via application level gateways.

Use Variable-Length Subnet Masks to make efficient use of address space.

VLSM makes more efficient use of address space by allowing big and small subnets. Remember that the routing protocol must be able to support VLSM. It must be able to base routing decisions on the IP address and a prefix. RIPv1 and Cisco's IGRP are examples for protocols that can not handle VLSM. RIPv2, OSPF and Cisco's EIGRP can handle VLSM.

The best way to use VLSM is to first subnet the entire address space by using a fixed subnet mask. Then take one of these big subnets and further subnet it with an extended subnet mask. If the small subnets are grouped, routing information can be aggregated.

Remember that IP addresses are binary entities! While you assign subnets reserve bits between the subnet and the host part for future growth. Allocate network numbers bit wise from the left to the right. Allocate host numbers bit wise from the right to the left.

³ Telephony networks are good examples for that.

⁴ Hosts aggregate into subnets, subnets aggregate into networks and networks aggregate into autonomous systems.




Figure 7, Network addresses are allocated bit-wise from the left to the right, host addresses are allocated bit-wise from the right to the left.

Avoid using more than two subnet masks for a single internetwork.

Use VLSM to save subnets in the WAN when allocating meshed serial lines.⁵ Serial lines each need a subnet number. Each serial line has only two host addresses. Use a regular subnet that is further subnetted to number all serial lines of the core network or a single distribution network. This allows route summarization⁶ for all subnets in the core or distribution network.

Limit the number of end nodes in one flat address space.⁷

The maximum number depends on the protocols used. A good rule of thumb is to have not more than 500 nodes within one flat address space. If you use Appletalk or NetBIOS you would probably want to limit the number of nodes to 200. Flat Novell IPX address spaces can probably be as big as 500 nodes. Flat IP address spaces can probably be as big as 1000 nodes as long as they are well behaved.⁸ Remember that less is better if you design for stability!

Use dynamic host address assignment (DHCP) to simplify configuration tasks.

To reduce configuration tasks, use a dynamic host address assignment protocol, for example the Bootstrap Protocol (BootP) or the newer Dynamic Host Configuration Protocol (DHCP). DHCP is the preferred protocol.

Renumbering IP hosts may be necessary to make use of the new IP address scheme. To prepare for the migration start using DHCP.

Derive addresses of other protocols from your IP addresses if possible.

You can easily derive the addresses for other protocols, at the data link or the network layer, from a valid IP addressing scheme.

A Novell IPX address is comprised from a four-byte network part and a six-byte host part. The adapter MAC address will be used for the host part. The network part, which has to be unique within the internetwork, can

⁵ In this context the term „serial line“ covers X.25 or Frame Relay point-to-point connections as well.

⁶ Route summarization is also called aggregation or supernetting. It refers to allocating multiple IP addresses in a way that allows summarization into a smaller number of routing table entries. This reduces memory consumption and routing traffic. Route summarization requires that multiple IP addresses share the same high-order bits and that the routing protocol carries the prefix length with the IP address.

⁷ „Flat address space“ means a router does not separate that part of the network. A switched network hanging off a router is a flat address space. Bridges and switches operate in flat address space.

⁸ The Unix boxes do not run rwho deamons and the route daemon is silent, i. E. the node listens to, but does not send RIP packets.

easily be derived from the IP network address. I. E. the decimal IP network address 10.47.11.0/24 will produce the hexadecimal IPX network address 0A2F0B00.

If you really want to use locally administered MAC addresses⁹, simply derive it from the adapters IP address. A MAC address, it is six bytes in length, has to be unique only within its own, flat address space. Using a valid IP addressing scheme guarantees that it is even unique within the whole internetwork. I. E. the decimal IP address 10.47.11.100 will produce the hexadecimal MAC address 0A:2F:0B:64:00:00 with two bytes of padding.

⁹ In my opinion there is no good reason to have locally administered MAC addresses.

3.4 Graphing Data

I assume you will not have the super duper tool that gathers all required data and turns it into meaningful graphs. There is a good probability that you will end up with a lot of data from various sources that you will need to correlate¹⁰ and turn into a meaningful, easily interpretable graph.

I would like to suggest some ways to graph data that I find useful.

3.4.1 Graphing Utilization Of LAN Segments

The network utilization and broadcast graph shows you the health of a LAN. The percentages are percent of total media bandwidth. For example a broadcast percentage of 10% means that broadcasts consumed 10 % of the total bandwidth not 10 % of the measured utilization. The maximum values are individuals. They do not belong to a set.




Figure 8, Network Utilization and Broadcast Information

¹⁰ Awk, grep, perl and companions are nice tools for that. ☺

3.4.2 Graphing Error Condition Of Links

The error condition graph shows the link utilization mapped with error conditions. The reference for this graph is the number of transmitted frames. As the number of frames is significant to determine the severity of an error it is important to annotate this graph with the sample size. For example a collision rate of 50 % is not a problem if the sample is only 10 frames long.

The example shows the error condition graph for an Ethernet segment. Use similar graphs to show the error conditions for other LAN or WAN technologies.




Figure 9, Under normal conditions you would expect the Ethernet error condition graph to form a sail between Utilization and Collisions.

3.4.3 Annotated Network Map

Use the utilization or packet rates obtained from your measurement to annotate a network map and give a graphical interpretation of the traffic intensity. The map shows clearly which links are heavily loaded, as well as which leaf networks have heavy traffic to the core.

Use multiple maps to graph large networks: One graph showing the core with the distribution networks feeding in. One graph for each distribution network with the access networks feeding in. One graph for each access network with the leaf networks feeding in.

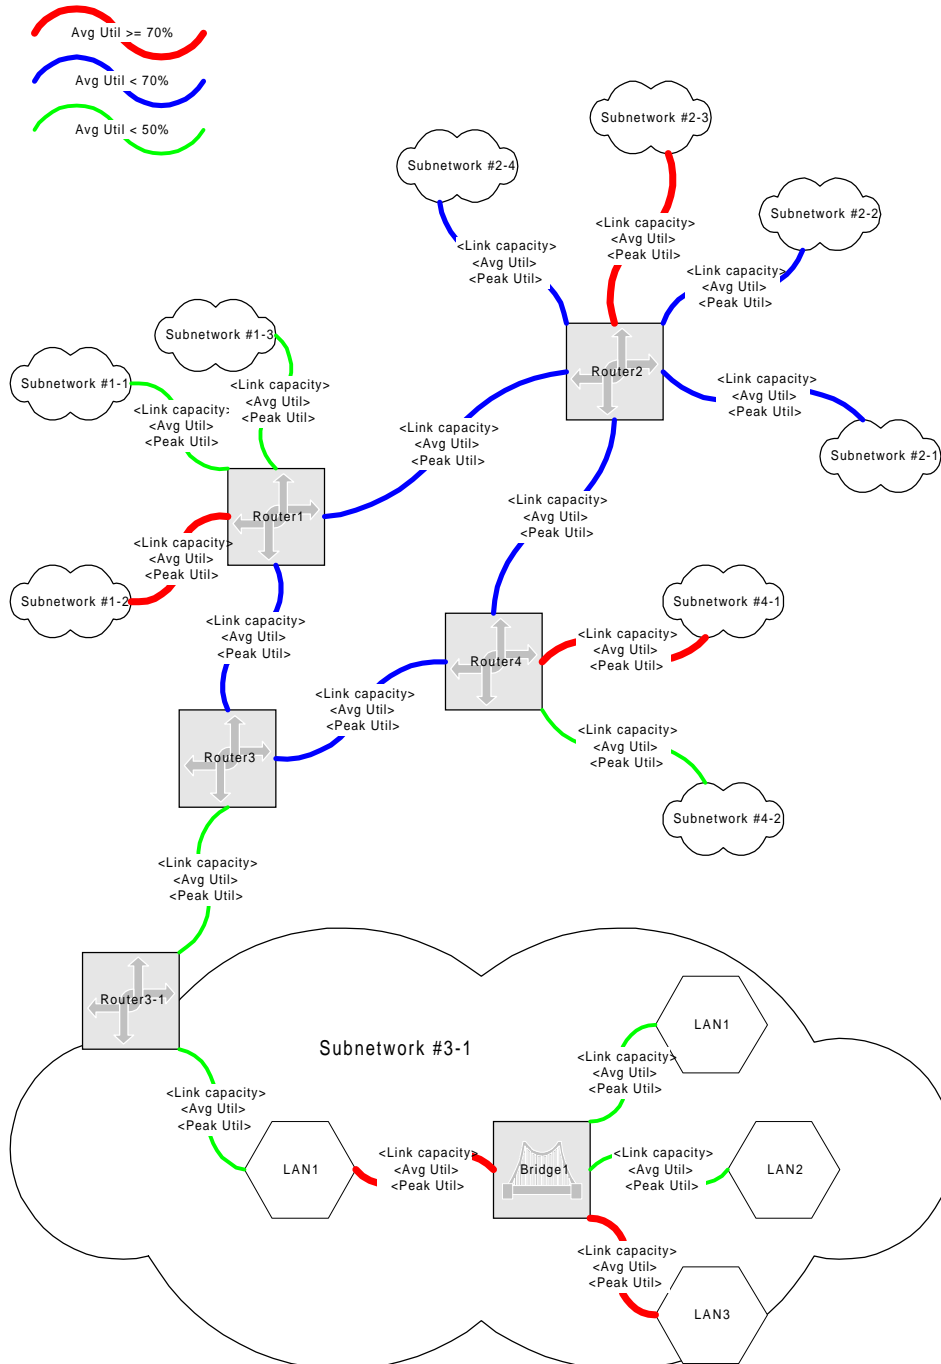


Figure 10, Network Map Annotated with Traffic Volume

3.4.4 Traffic Matrix

Although the annotated network map is informative and easy to interpret it does not show the correspondence between sources and destinations. A better picture of who is talking to whom can be obtained by a matrix of source/destination pairs, with the traffic volume of the subnetworks being represented by the square at the intersection. The rows and columns are sorted by total destination traffic.

Both the annotated network map and the source/destination matrix provide valuable information about the traffic flows in a network. The annotated network map shows how the network is utilized. The traffic matrix explains why the traffic flows as it does in the map.

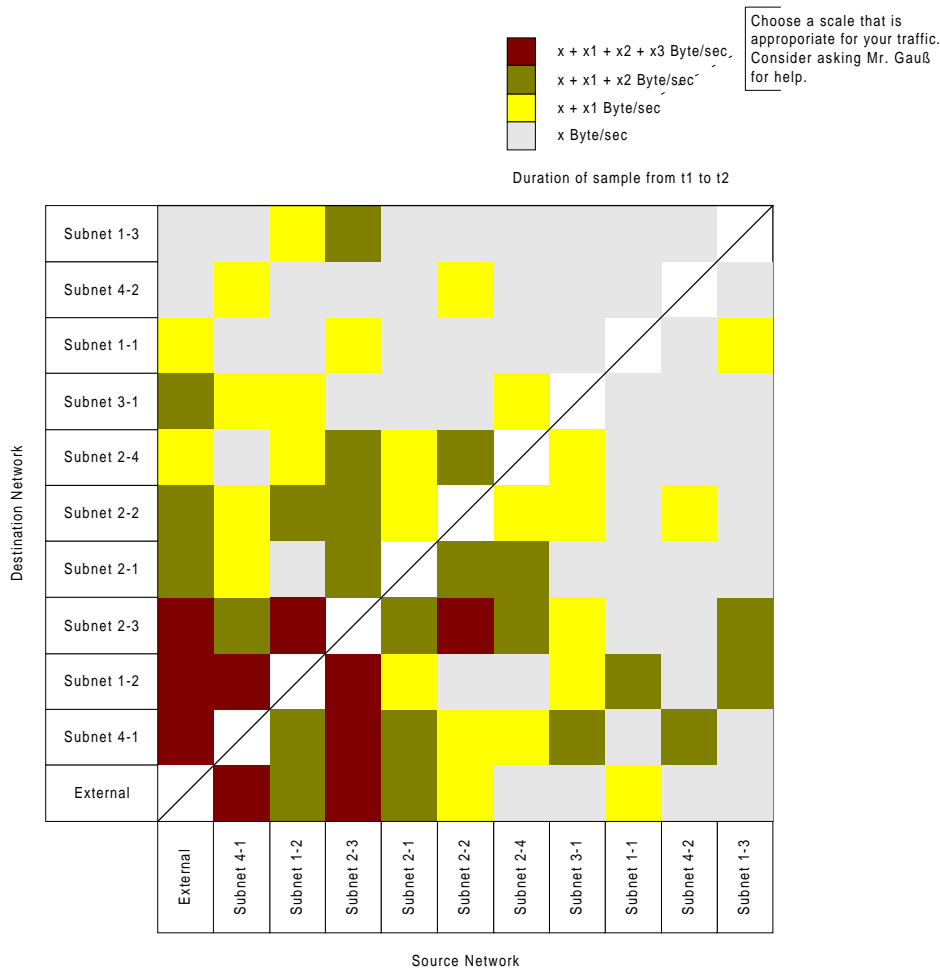


Figure 11, Traffic Source/Destination Matrix

3.5 Testing And Modeling

The techniques for performance evaluation are measurement, analytical modeling and simulation.

Measurements are possible only if something similar to the existing system already exists. Analytical modeling and simulation can be used for situations where measurement is not possible. In general it would be preferred to base analytical modeling and simulation on previous measurement.

Measurements may not give valid results because many of the environment parameters, such as workload or time of the measurement may be unique to the experiment. The accuracy of results varies from very high to none.

Analytical modeling requires so many simplifications and assumptions that if the results turn out to be accurate even the analysts are surprised.

Simulation incorporates more detail and requires less assumption than analytical modeling. The results are usually closer to reality.

Analytical models provide the best insight into the effects of various parameters and their interactions.

Simulation allows searching the space of possible parameters for the optimal combination but often it is not clear what the trade-off is among different parameters.

With measurement it is often difficult to tell if the improved performance is a result of some random change or due to the particular parameter setting.

It is good practice to use two or more techniques simultaneously. For example you may use measurement and simulation together to verify and validate the results of each one. *Until validated, all evaluation results are suspect!*

Table 1, Criteria for Selecting an Evaluation Technique

Criterion	Analytical Modeling	Simulation	Measurement
Stage	Any	Any	Post prototype
Time required	Small	Medium	Varies
Tools	Analysts ^{MAB}	Simulation languages ^{MAB}	Instrumentation
Accuracy¹	Low	Moderate	Varies
Trade-off evaluation	Easy	Moderate	Difficult
Cost	Small	Medium	High
Saleability	Low	Medium	High

¹ In all cases, results may be misleading or wrong.

Copied without permission from Raj Jain, The Art of Computer Systems Performance Analysis.

^{MAB} Powerful object-oriented, GUI-driven tools for analytical modeling and simulation of networks and applications are available today, for example CACIs Comnet III, Predictor and Profiler.

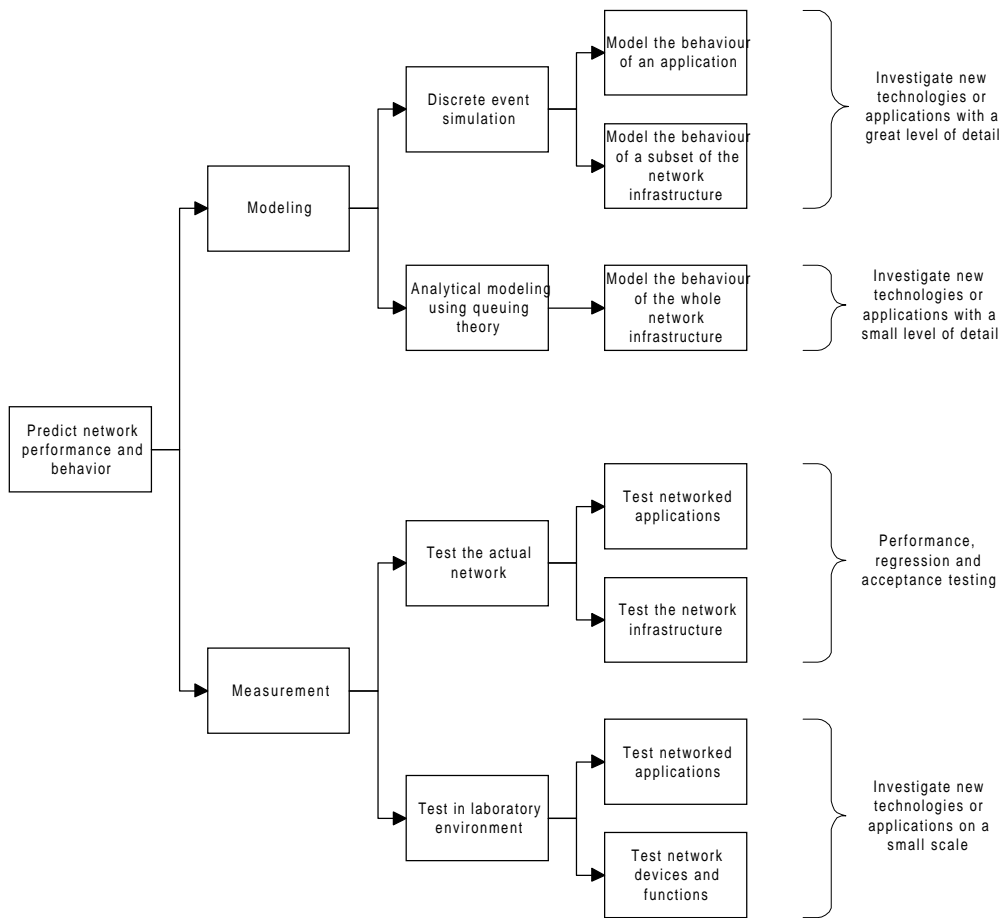


Figure 12, Measurement vs. Modeling

Use measurement techniques to profile existing applications and workloads.

Use modeling techniques to consolidate load information and optimize the topology.

Always validate your results!

Use measurement techniques for acceptance testing.

3.6 Application Design Basics

Use sliding window protocols.

Do not implement request-response type protocols for your application. They are less efficient than bursts because every transmission must wait for an acknowledgement. Only conversations that send bursts can approach the full capacity of a link during the burst.

Use large packets.

Code your application to use large packets. Avoid conversations that use small packets. Larger packets yield a higher ratio of payload to overhead.

Use common services.

Ensure that all applications share the same resource discovery methods, naming services etc.

Use efficient protocols.

Avoid using LLC2, NetBios and any bridging technology that relies on broadcasts or explorers.

Use low latency WAN technologies.

Remember that the killer for client server applications over the WAN is network latency (propagation delay, insertion delay, queuing delay, processing delay). Leased lines and switched lines introduce less delay than packet switching connections.

Locate servers and clients as close as possible.

4 Support Material

The *descriptive text* was copied from the home pages.

4.1 Tools

4.1.1 Free Tools

4.1.1.1 CAIDA Measurement Tool Taxonomy

URL: <http://www.caida.org/Tools/taxonomy.html>

This tool taxonomy provides a preliminary overview of Internet and TCP/IP performance measurement tools and efforts. We categorize tools used to measure the Internet as well as general TCP/IP performance tools. We will update the information here as we receive/review updates on tools and activities. If you are aware of any tools that are missing from this listing or would like to share your experiences with any of these tools, please contact us at info@nlanr.net.

4.1.1.2 Multi Router Traffic Grapher

URL: <http://ee-staff.ethz.ch/~oetiker/webtools/mrtg/mrtg.html>

The Multi Router Traffic Grapher (MRTG) is a tool to monitor the traffic load on network-links. MRTG generates HTML pages containing GIF images which provide a LIVE visual representation of this traffic.

Check <http://www.ee.ethz.ch/stats/mrtg/> for an example. MRTG is based on Perl and C and works under UNIX and Windows NT. MRTG is being successfully used on many sites arround the net. Check the MRTG-Site-Map.

MRTG is available under the GNU PUBLIC LICENSE.

Highlights of MRTG

1. *Works on most UNIX platforms and Windows NT*
2. *Uses Perl for easy customization*
3. *Has a highly portable SNMP implementation written entirely in Perl thanks to Simon Leinen. There is no need to install any external SNMP package.*
4. *MRTG's logfiles do NOT grow. Thanks to the use of a unique data consolidation algorithm.*
5. *MRTG comes with a semi-automatic configuration tool.*
6. *MRTG's query engine checks for port reconfigurations on the router and warns the user when they occur.*
7. *Time critical routines are written in C thanks to the initiative of Dave Rand my Co-Author*
8. *Graphics are generated directly in GIF format, using the GD library by Thomas Boutell.*
9. *The look of the web pages produced by MRTG is highly configurable.*

MRTG Mailing List

There are two mailing lists for MRTG available. One is called 'mrtg' and is a discussion list for users and developers. The other is called 'mrtg-announce' and is a low volume list for MRTG related announcements.

To subscribe to these mailing lists, send a message with the subject line subscribe to either mrtg-request@list.ee.ethz.ch or mrtg-announce-request@list.ee.ethz.ch. For posting to the mrtg list use the address mrtg@list.ee.ethz.ch.

Further information about the usage of the mailing lists is available by sending a message with the subject line 'help' to either one of the request addresses.

For past activity there is also a mailing list archive available: <http://www.ee.ethz.ch/~slist/mrtg>

4.1.1.3 NetSCARF

URL: <http://nic.merit.edu/~netscarf/>

NetSCARF - Network Statistics Collection and Reporting Facility - runs on Unix and NT.

DESCRIPTION

The NetSCARF project Scion package is a set of standalone programs for TCP/IP network statistics collection and reporting. The basic idea is that you can enter some configuration information and have your network statistics appear on the World Wide Web with relatively little effort.

AVAILABILITY

Refer to NetSCARF Project Home Page for how to download software at <http://nic.merit.edu/~netscarf>

The package components include:

- *collect* *Queries network nodes for SNMP data.*
- *scook* *Pre-processes SNMP data.*
- *scache* *Pre-creates standard graphs during off-hours to speed graph display.*
- *sctest* *Serves the data to OpStats (rfc1856) compliant clients.*

The initial work was developed during the time when Merit was operating the U.S. Internet backbone called the NSFNET. The goal of the NetSCARF project is to help instrument the Internet by providing turnkey measurement and reporting software during the post-NSFNET era. Scion means descendent, which is reflective of the fact that the code descends from NSFNET origins.

4.1.2 Commercial Tools

4.1.2.1 CACI

URL: <http://www.caciasl.com>

- Comnet III, Network and application simulation
- Comnet Predictor, Network capacity planning
- Comnet Profiler, IT infrastructure analysis and planning

4.1.2.2 Cisco

URL: <http://www.cisco.com>

- SwitchProbe, RMON probe
- TrafficDirector, Traffic analysis and reporting using data from SwitchProbes or Catalyst switches
- Netsys Enterprise Solver, Analysis, capacity planning and simulation of Cisco¹¹ router networks
- Netsys Service-Level Management Suite, Network and service level management of Cisco router networks

4.1.2.3 Optimal Networks

URL: <http://www.optimal.com>

- Optimal Application Expert, Analysis of application behavior in the network
- Optimal Application Insight, Application monitoring and service level management

¹¹ Network General offers a module that allows Netsys Enterprise Solver to interpret configuration profiles of Bay Network routers.

- Optimal Internet Monitor, Traffic reporting
- Optimal Performance, Network capacity planning
- Optimal Surveyor, Network topology discovery

4.1.2.4 Network General

URL: <http://ngc.com>

- Sniffer, The protocol analyzer
- Distributed Sniffer System, Network fault and performance management
- NetXRay, Software based protocol analyzer
- Distributed NetXRay, Traffic analysis and reporting using NetXRay probes (software on Win95/WinNT systems)
- NetScout Probe, RMON probe
- NetScout Manager, Traffic analysis and reporting using data from NetScout Probes
- Service Level Manager, Service level management of network components
- Netsys Enterprise Solver, Analysis, capacity planning and simulation of Cisco and Bay router networks
- Chariot Console & Endpoint, Network performance measurement
- RouterPM, Health monitoring of Cisco routers
- SwitchPM, Health monitoring of switches

4.2 Standards Bodies And Organizations On The Internet

4.2.1 ACM Special Interest Group On Communications

URL: <http://www.acm.org/sigcom>

SIGCOMM is a professional forum for the discussion of issues in the field of data communications and computer networks. Current emphases are on the architectures, protocols, design, analysis, measurement, maintenance, regulatory policy, standards, applications, and social impact of computer networks and computer communications systems.

4.2.2 ANSI

URL: <http://www.ansi.org/>

ANSI Catalog: http://www.ansi.org/docs/cat_top.html

The ANSI Catalog is an electronic version of our hard-copy Catalog. It includes a searchable database of the over 11,000 approved American National Standards, all of which are available for purchase from ANSI. It also contains information about other ANSI publications and services.

MAB:

This site provides just a catalog. The documents are not available online.

4.2.3 IEEE

URL: <http://www.ieee.org/>

Institute of Electrical and Electronics Engineers standards documents are available online only a subscription basis.

MAB:
Pretty useless site.

IEEE Computer Society: <http://www.computer.org/>

Everyone has free access to all issues of 17 of the society's magazines and transactions from 1995 to the present. Some time in the fourth quarter the library will be open to Computer Society members only. Beginning in 1998, access to each periodical's digital library archive will be available to society members on a subscription basis.

IEEE/ACM Transactions on Networking: <http://dworkin.wustl.edu/~ton/> or <http://www.csrc.wustl.edu/~ton>

The IEEE/ACM Transactions on Networking is an archival, bimonthly journal committed to the timely publication of high quality papers that advance the state-of-the-art and practical applications of communication networks. It is co-sponsored by the IEEE Communications Society, the IEEE Computer Society, and the ACM with its Special Interest Group on Data Communications (SIGCOMM). This page provides information about the journal and its contents, as well as other information of interest to the networking community.

4.2.4 Internet

4.2.4.1 Internet Ad Hoc Committee

URL: <http://www.iahc.org/>

The Internet Ad Hoc Committee (IAHC) is a coalition of participants from the broad Internet community, working to satisfy the requirement for enhancements to the Internet's global Domain Name System (DNS).

Public Information Resources on the Internet Domain Name System: <http://www.iahc.org/dns-refs>

4.2.4.2 Internet Architecture Board

URL: <http://www.iab.org/iab/>

The Internet Architecture Board (IAB) is a technical advisory group of the Internet Society. Its responsibilities include:

- 1. IESG Selection: The IAB appoints a new IETF chair and all other IESG candidates, from a list provided by the IETF nominating committee.*
- 2. Architectural Oversight: The IAB provides oversight of the architecture for the protocols and procedures used by the Internet.*
- 3. Standards Process Oversight and Appeal: The IAB provides oversight of the process used to create Internet Standards. The IAB serves as an appeal board for complaints of improper execution of the standards process.*
- 4. RFC Series and IANA: The IAB is responsible for editorial management and publication of the Request for Comments (RFC) document series, and for administration of the various Internet assigned numbers.*
- 5. External Liaison: The IAB acts as representative of the interests of the Internet Society in liaison relationships with other organizations concerned with standards and other technical and organizational issues relevant to the world-wide Internet.*
- 6. Advice to ISOC: The IAB acts as a source of advice and guidance to the Board of Trustees and Officers of the Internet Society concerning technical, architectural, procedural, and (where appropriate) policy matters pertaining to the Internet and its enabling technologies.*

4.2.4.3 Internet Assigned Numbers Authority

URL: <http://www.iana.org/iana/>

The Internet Assigned Numbers Authority (IANA) is the central coordinator for the assignment of unique

parameter values for Internet protocols.

The IANA is chartered by the Internet Society (ISOC) and the Federal Network Council (FNC) to act as the clearinghouse to assign and coordinate the use of numerous Internet protocol parameters.

The Internet protocol suite, as defined by the Internet Engineering Task Force (IETF) and its steering group (the IESG), contains numerous parameters, such as internet addresses, domain names, autonomous system numbers (used in some routing protocols), protocol numbers, port numbers, management information base object identifiers, including private enterprise numbers, and many others.

The common use of the Internet protocols by the Internet community requires that the particular values used in these parameter fields be assigned uniquely. It is the task of the IANA to make those unique assignments as requested and to maintain a registry of the currently assigned values.

The most recent summary of these assigned parameter values is "Assigned Numbers" which is STD-2 and RFC-1700 published in October 1994. The assignments may also be found online (<ftp://ftp.isi.edu/in-notes/iana/assignments>).

4.2.4.4 Internet Engineering Task Force

URL: <http://www.ietf.org>

The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual.

The actual technical work of the IETF is done in its working groups, which are organized by topic into several areas (e.g., routing, transport, security, etc.). Much of the work is handled via mailing lists. The IETF holds meetings three times per year.

The IETF working groups are grouped into areas, and managed by Area Directors, or ADs. The ADs are members of the Internet Engineering Steering Group, or IESG . Providing architectural oversight is the Internet Architecture Board, or IAB ; the IAB also adjudicates appeals when someone complains that the IESG has failed. The IAB and IESG are chartered by the Internet Society (ISOC) for these purposes. The General Area Director also serves as the chair of the IESG and of the IETF, and is an ex-officio member of the IAB.

The Internet Assigned Numbers Authority (IANA) is the central coordinator for the assignment of unique parameter values for Internet protocols. The IANA is chartered by the Internet Society (ISOC) and the Federal Network Council (FNC) to act as the clearinghouse to assign and coordinate the use of numerous Internet protocol parameters.

First-time attendees might find it helpful to read RFC 1718 "The Tao of the IETF" (<http://www.ietf.org/tao.html>).

4.2.4.5 Internet Research Task Force

URL: <http://www.irtf.org/irtf/>

The Internet Research Task Force (IRTF) is a composed of a number of focused, long-term and small Research Groups (<http://www.irtf.org/irtf/groups.htm>). These groups work on topics related to Internet protocols, applications, architecture and technology. Research Groups are expected to have the stable long-term membership needed to promote the development of research collaboration and teamwork in exploring research issues. Participation is by individual contributors, rather than by representatives of organizations.

The IRTF is managed by the IRTF Chair in consultation with the Internet Research Steering Group (IRSG). The IRSG membership includes the IRTF Chair, the chairs of the various Research Group and possibly other individuals ("members at large") from the research community.

The IRTF Chair is appointed by the Internet Architecture Board (IAB), the Research Group chairs are appointed as part of the formation of Research Groups and the IRSG members at large are chosen by the IRTF Chair in consultation with the rest of the IRSG and on approval of the IAB. In addition to managing the Research Groups, the IRSG may from time to time hold topical workshops focusing on research areas of

importance to the evolution of the Internet, or more general workshops to, for example, discuss research priorities from an Internet perspective.

The IRTF Research Groups guidelines and procedures are described more fully in RFC 2014 (<ftp://ftp.isi.edu/in-notes/rfc2014.txt>).

4.2.4.6 Internet Society

URL: <http://info.isoc.org/index.html>

WHAT IS THE INTERNET SOCIETY?

The Internet Society (ISOC) is the international organization for global cooperation and coordination for the Internet and its internetworking technologies and applications.

WHO ARE ITS MEMBERS?

Its members reflect the breadth of the entire Internet community and consist of individuals, corporations, non-profit organizations, and government agencies.

WHAT IS ITS PURPOSE?

Its principal purpose is to maintain and extend the development and availability of the Internet and its associated technologies and applications - both as an end in itself, and as a means of enabling organizations, professions, and individuals worldwide to more effectively collaborate, cooperate, and innovate in their respective fields and interests.

Its specific goals and purposes include:

- *development, maintenance, evolution, and dissemination of standards for the Internet and its internetworking technologies and applications;*
- *growth and evolution of the Internet architecture;*
- *maintenance and evolution of effective administrative processes necessary for operation of the global Internet and internets;*
- *education and research related to the Internet and internetworking;*
- *harmonization of actions and activities at international levels to facilitate the development and availability of the Internet;*
- *collection and dissemination of information related to the Internet and internetworking, including histories and archives;*
- *assisting technologically developing countries, areas, and peoples in implementing and evolving their Internet infrastructure and use;*
- *liaison with other organizations, governments, and the general public for coordination, collaboration, and education in effecting the above purposes.*

HOW DOES IT OPERATE?

The Internet Society operates through its international Board of Trustees, its International Networking Conferences and developing country workshops, its regional and local chapters, its various standards and administrative bodies, its committees, and its secretariat. The Board of Trustees is headed by a President with the assistance of several officers. The Board consists of 18 eminent individuals drawn from every region of the world - most of whom were instrumental in creating and evolving different components of the Internet and the technology.

WHERE IS IT LOCATED?

The permanent international headquarters and secretariat of the Society is located at Reston VA USA and is headed by the Executive Director.

WHAT IS ITS LEGAL STATUS?

The Internet Society is incorporated as a not-for-profit corporation, with tax-deductible status in Washington DC USA, near its headquarters location.

WHY WAS IT CREATED?

The Internet Society was announced in June 1991 at an international networking conference in Copenhagen and brought into existence in January 1992 by a worldwide cross-section of individuals and organizations who recognized that the Society was a critical component necessary to evolve and globalize the Internet and internet technologies and applications, and to enhance their availability and use on the widest possible scale.

4.2.4.7 RFC Archive Indexed By Topic.

URL: <http://www.it.kth.se/docs/rfc/>

RFC index in reverse numeric order: Fehler! Textmarke nicht definiert.

4.2.4.8 RFC Editor

URL: <http://www.isi.edu/rfc-editor/>

The Requests for Comments (RFCs) are a series of notes, started in 1969, about the Internet (originally the ARPANET). The notes discuss many aspects of computing and computer communication focusing in networking protocols, procedures, programs, and concepts, but also including meeting notes, opinion, and sometimes humor.

The specification documents of the Internet protocol suite, as defined by the Internet Engineering Task Force (IETF) and its steering group (the IESG), are published as RFCs.

The RFC Editor is the publisher of the RFCs and is responsible for the final editorial review of the documents. The RFC Editor is chartered by the Internet Society (ISOC) and the Federal Network Council (FNC).

4.2.4.9 Réseaux IP Européens

URL: <http://www.ripe.net/>

RIPE (Réseaux IP Européens) is a collaborative organisation open to all parties operating wide area IP networks in Europe. The objective of RIPE is to ensure the administrative and technical coordination necessary to enable operation of a pan-European IP network. RIPE does not operate a network of its own.

4.2.5 International Telecommunication Union

URL: <http://www.itu.ch>

ITU-T (formerly known as CCITT) Recommendations are available online and on CD-ROM but only on a subscription basis. An annual subscription is about 2800 Swiss Franks for CD-ROM and 3200 Swiss Franks for online access.

MAB:

Without a subscription is this site useless!

In summer 1996 InfoMagic (info@infomagic.com) published a "Standards" CD-ROM that had a good compilation of ITU-T recommendations. The price was 10\$ or 20\$.

4.2.6 Network Device Test Lab

URL: <http://ndtl.harvard.edu/>

Scott Bradner is a Senior Technical Consultant at Harvard University. He is the director of the Harvard Network Device Test Lab and publishes an annual review of the performance of network interconnect devices for Network Computing (<http://techweb.cmp.com/nc/docs/>).

All tests used in the Harvard Network Device Test Lab are based on work by The Internet Engineering Task

Force (IETF) Benchmarking Methodology Working Group (BMWG), which has published two RFCs (see <http://www.ietf.org>). RFC 1242 defines testing terminology; RFC 1944, testing methodology. All testing is done using special-purpose standalone testers to ensure the most reliable results. The scripts that drive the testers are available on the lab's Web site at ndtl.harvard.edu/ndtl.

Results of the tests are at <http://ndtl.harvard.edu/ndtl/results/data/>.

4.3 Other Useful Internet Sites

4.3.1 Searchable Archives Of Mailing Lists

URL: Fehler! Textmarke nicht definiert.

This site has a searchable archive of some good mailing lists. The lists include BSDi users, Firewalls, Cisco, Ascend, Sun managers, and HP-UX sysadmin.

4.3.2 Cisco Mailing List

To subscribe to the Cisco mailing list, send your request to cisco-request@spot.colorado.edu.

4.3.3 Local Area Networks

4.3.3.1 The BIG-LAN Mailing List

To subscribe to the **BIG-LAN mailing list** send an electronic mail a one-line message of this form to listserv@listserv.syr.edu: `subscribe big-lan <Firstname> <Lastname>`

4.3.3.2 Charles Spurgeon's Ethernet Site

URL: <http://wwwhost.ots.utexas.edu/ethernet/>

This site provides extensive information about Ethernet (IEEE 802.3) local area network (LAN) technology. Including the original 10-Megabit per second (Mbps) system, the 100 Mbps Fast Ethernet system (802.3u), and the Gigabit Ethernet system (802.3z).

MAB:

Probably the most comprehensive site on Ethernet! Very good!

4.3.3.3 Ethernet Type Codes

<ftp://ftp.isi.edu/in-notes/iana/assignments/ethernet-numbers>

<http://www.cavebear.com/CaveBear/Ethernet/>

These same codes are also used on some other IEEE 802 networks (e.g. 802.5 is Token Ring), but may be in slightly different form. Specifically 802.5 sends the bits of a byte in the opposite order to Ethernet or 802.3, so the codes listed below may have the bits of each byte reversed (e.g. exchange 01/80, 10/08, 0C/30, etc.), although some monitors may undo the reversal back to the Ethernet order. Be careful using this reference for other IEEE 802 media.

4.3.3.4 Gigabit Ethernet Alliance

URL: <http://www.gigabit-ethernet.org/>

The Gigabit Ethernet Alliance is an open forum whose purpose is to promote industry cooperation in the development of Gigabit Ethernet.

A number of high-speed LAN technologies have been proposed to provide greater bandwidth and improved client/server response times. Foremost among them is Fast Ethernet, or 100BaseT, a technology designed to provide a non-disruptive, smooth evolution from 10BaseT Ethernet to high-speed 100Mbps performance.

Given the trend toward 100BaseT connections to the desktop, there is a clear need for even higher-speed connections at the server and backbone level.

Gigabit Ethernet will be ideal for deployment as a backbone interconnect between 10/100BaseT switches, and as a connection to high-performance servers. Gigabit Ethernet is a natural upgrade path for future high-end desktop computers that will require more bandwidth than can be provided by 100BaseT.

The Alliance was founded by the following companies: 3Com Corp, Bay Networks, Cisco Systems, Compaq Computer, Granite Systems Inc., Intel Corporation, LSI Logic, Packet Engines, Sun Microsystems, UB Networks, VLSI Technology.

4.3.3.5 Measured Capacity Of An Ethernet

<http://www.research.digital.com/wrl/publications/abstracts/88.4.html>

DEC Research Report 88.4 -- Measured Capacity of an Ethernet: Myths and Reality by David R. Boggs, Jeffrey C. Mogul, and Christopher A. Kent

Ethernet, a 10 Mbit/sec CSMA/CD network, is one of the most successful LAN technologies. Considerable confusion exists as to the actual capacity of an Ethernet, especially since some theoretical studies have examined operating regimes that are not characteristic of actual networks. Based on measurements of an actual implementation, we show that for a wide class of applications, Ethernet is capable of carrying its nominal bandwidth of useful traffic, and allocates the bandwidth fairly. We discuss how implementations can achieve this performance, describe some problems that have arisen in existing implementations, and suggest ways to avoid future problems.

MAB:

See also Mart Molles paper on BLAM, available at Charles Spurgeons site.

4.3.3.6 ASTRAL

URL: <http://www.astral.org>

ASTRAL is an alliance of leading Token Ring technology providers dedicated to supporting Token Ring users as they prepare for the future of high-demand networking. ASTRAL provides vendor-independent education and information about Token Ring and new technology developments that will be important to the future of Token Ring customers.

4.3.3.7 Token Ring Consortium

URL: <http://www.iol.unh.edu/consortiums/tokenring/index.html>

The Token Ring Consortium tests Token Ring (IEEE 802.5) and Dedicated Token Ring (IEEE 802.5r) products and software from both an interoperability and conformance perspective.

They publish a Token Ring newsletter:

"MACs 'n' PHYs" is a tri-yearly technical report published by the Token Ring Consortium of the University of New Hampshire's InterOperability Lab. "MACs 'n' PHYs" discusses testing issues, displays schematics helpful for in-house testing, and reports testing statistics.

4.3.4 Wide Area Networks

4.3.4.1 Dan Kegel's ISDN Page:

URL: <http://www.alumni.caltech.edu/~dank/isdn/>

Note: This information provided for entertainment purposes only :-) This page is mostly a collection of pointers to WWW and FTP documents on other servers. If your favorite online source of ISDN info is not listed here, write me, Dan Kegel <dank@alumni.caltech.edu>. If you need help putting your company's product info on WWW, please read my guide to publishing product info on the Web.

MAB:
Not too bad.

4.3.4.2 SMDS

SMDS SIG: <http://www.cerf.net/smds/>

Switched Multimegabit Data Service (SMDS) is a connectionless, cell-switched data transport service.

4.3.4.3 SONET & SDH

Nortel home page: <http://www.nortel.com/cool>

MAB:
Somewhere on the Nortel site is a document called „SONET101: Intro to Sync Opt Nets; SONET & SDH“. It is a quite good introduction to SONET/SDH. (At least for those who do not have an EE background.)

SONET Interoperability Forum: <http://www.atis.org/atis/sif/sifhom.htm>

I would recommend the books „A Sourcebook of Synchronous Networking“ by Curtis A. Siller and Mansoor Shafi (IEEE Press, ISBN: 0-7803-1168-X) and „Broadband Communications, A professional's guide to ATM, Frame Relay, SMDS, SONET and B-ISDN“ by Balaji Kumar (McGraw-Hill, ISBN: 0-07-035968-7).

4.3.5 Security

4.3.5.1 Computer Emergency Response Team

URL: <http://www.cert.org/>

The CERT Coordination Center is the organization that grew from the computer emergency response team formed by the Defense Advanced Research Projects Agency (DARPA) in November 1988 in response to the needs identified during the Internet worm incident. The CERT charter is to work with the Internet community to facilitate its response to computer security events involving Internet hosts, to take proactive steps to raise the community's awareness of computer security issues, and to conduct research targeted at improving the security of existing systems.

CERT products and services include 24-hour technical assistance for responding to computer security incidents, product vulnerability assistance, technical documents, and seminars. In addition, the team maintains a mailing list for CERT advisories, and provides a web site, www.cert.org, and an anonymous FTP server, info.cert.org, where security-related documents, CERT advisories, and tools are available.

If you would like to be added to the CERT mailing list, please send email to cert-advisory-request@cert.org. On the subject line, type SUBSCRIBE <your-email-address>

4.3.5.2 Electronic Privacy Information Center

URL: <http://www.epic.org/>

EPIC is a public interest research center in Washington, D.C. It was established in 1994 to focus public attention on emerging civil liberties issues and to protect privacy, the First Amendment, and constitutional values. EPIC is a project of the Fund for Constitutional Government. EPIC works in association with Privacy International, an international human rights group based in London, UK and is also a member of the Global Internet Liberty Campaign and the Internet Privacy Coalition.

Online guide to privacy tools: <http://www.epic.org/privacy/tools.html>

Online guide to privacy resources:

4.3.5.3 Firewalls

4.3.5.3.1 Internet Firewalls Frequently Asked Questions By Marcus J. Ranum

URL: <http://www.clark.net/pub/mjr/pubs/fwfaq/>

4.3.5.3.2 Introduction To Firewalls By NIST

URL: <http://www.telstra.com.au/pub/docs/security/800-10/main.html>

This document provides an overview of the Internet and security-related problems. It then provides an overview of firewall components and the general reasoning behind firewall usage. Several types of network access policies are described, as well as technical implementations of those policies. Lastly, the document contains pointers and references for more detailed information.

This work is a contribution of the National Institute of Standards and Technology, and is not subject to copyright.

4.3.5.3.3 Firewall Mailing List

The Firewalls mailing list is a discussion forum for firewall administrators and implementors. To subscribe to Firewalls, send mail to Majordomo@GreatCircle.COM. In the body of the message put "subscribe firewalls".

The Firewalls digest is a compilation of messages from the Firewalls mailing list. To subscribe to the Firewalls digest, send mail to Majordomo@GreatCircle.COM. In the body of the message put "subscribe firewalls-digest".

Compressed back issues are available from <ftp://FTP.GreatCircle.COM/pub/firewalls/digest/>.

A searchable mailing list archive can be found at <http://www.nexial.com>.

4.3.5.3.4 Firewall Products

A copy of Cathy Fulmer's list of products can be found at <http://www.waterw.com/~manowar/vendor.html> or <http://www.access.digex.net/~bdboyle/firewall.vendor.html>.

Another list can be found at [Fehler! Textmarke nicht definiert.](#)

4.3.5.3.4.1 Cisco's Centri Firewall

Source: Cisco Headlines Vol. 2, No. 18

CISCO CENTRI FIREWALL

Download a free evaluation copy of the Cisco Centri Firewall, ideal for small to medium-sized businesses. Using Windows NT, the Centri Firewall combines its unique Kernel Proxy[™] architecture with the Centri Security Policy Builder[™] to provide strong security, high performance, and ease of use.

<http://www.cisco.com/warp/customer/751/centri/index.html>

4.3.5.3.4.2 Cisco's 1605-R Router

Source: Cisco Headlines Vol. 2, No. 22

Cisco announces the dual Ethernet Cisco 1605-R router with the Cisco IOS[™] Firewall feature set, designed for small to medium-sized businesses and remote offices that require secure intranet connections and secure communications over the Internet.

<http://www.cisco.com/warp/customer/146/1977.html>

4.3.5.3.5 Nomad Mobile Research Center

URL: [//www.nmrc.org](http://www.nmrc.org)

This site has information about security/insecurity of systems. It provides security-related tools and FAQs on hacking NetWare, NT and Unix systems.

4.3.5.3.6 Phrack Magazine

URL: <http://www.fc.net/phrack/>

Phrack Magazine is one of the longest running electronic magazines in existence. Since 1985, Phrack has been providing the hacker community with information on operating systems, networking technologies, and telephony, as well as relaying other topics of interest to the international computer underground.

4.4 Template Of A Network Design Document

<td>

4.5 Recommended Reading

Title	Understanding IP Addressing: Everything You Ever Wanted To Know
Author	Chuck Semeria
Publisher	
Year published	
ISBN/URL	www.3com.com/nsc/501302.html

Title	Interconnections
Author	Radia Perlman
Publisher	Addison-Wesley Publishing Company
Year published	1992
ISBN/URL	0-201-56332-0

Title	Routing in the Internet
Author	Christian Huitema
Publisher	Prentice-Hall PTR
Year published	1995
ISBN/URL	0-13-132192-7

Title	Internet Routing Architectures
Author	Bassam Halabi
Publisher	New Riders Publishing
Year published	1997
ISBN/URL	1-56205-652-2

Title	The Art of Computer Systems Performance Analysis, Techniques for Experimental Design, Measurement, Simulation and Modeling
--------------	--

Author	Raj Jain
Publisher	John Wiley & Sons
Year published	1991
ISBN/URL	0-471-50336-3

Title	Internetworking Technologies Handbook
Author	M. Ford, H. K. Lew, S. Spanier, T. Stevenson
Publisher	New Riders Publishing
Year published	1997
ISBN/URL	1-56205-603-4

Title	Internetworking Design Basics
Author	Part of the IOS documentation set
Publisher	Cisco Systems
Year published	
ISBN/URL	CCO or UniverCD

Title	Designing Large-Scale IP Internetworks
Author	Part of the IOS documentation set
Publisher	Cisco Systems
Year published	
ISBN/URL	CCO or UniverCD