



Computer Networks Design & Planning

Dr.-Ing. Ali Diab

Outline

- Goals & Objectives
- Design Approaches
 - Traditional analytic design
 - Building block approach
- Requirements: Boring, However Essential
- Campus Networks Design

Goals & Objectives

Key Questions

- Our network design must answer some pretty basic questions
 - What stuff do we get for the network?
 - How do we connect it all?
 - How do we have to configure it to work right?
- This meant mostly capacity planning
 - Having enough bandwidth to keep data moving
 - May be effective, but result in over engineering
- While some uses of the network will need a lot of bandwidth (multimedia), we may also need to address
 - Security
 - Considering both internal and external threats
 - Possible wireless connectivity
 - Reliability and/or availability
 - Like speed for a car, how much are you willing to afford?

Goals & Objectives

- **Good designs should**
 - Deliver services requested by users
 - Deliver acceptable throughput and response times
 - Be within budget and maximize cost efficiencies
 - Be reliable
 - Be expandable without major re-design
 - Be manageable by maintenance and support staff
 - Be well documented


Design Approaches

Basic Approaches


- Two typical methods
 - Traditional analytic design
 - Building block approach
- Both use a similar iterative approach



Traditional Design Process




Building Block Approaches



Design Approaches

Traditional Design Process

Traditional Design Process



Agree Requirements

- Engage end users
- Translate requirements
 - Business objectives → technical specification
- Phasing the requirements
 - Right level of detail at each design stage
- Designing the requirements
 - Aim for completeness
 - Prioritise with a hierarchical system such as
 - [M] - Mandatory
 - [H] – Highly desirable
 - [D] - Desirable
 - [N] - Note

Information Gathering

- Need to find details of user behaviour, application use and location information, e.g.
 - User: location, numbers, services used, typical access
 - Sites: number, location, constraints on traffic (security, political or cost)
 - Servers and services: location, level of distribution
 - WAN/backbone predicted link traffic
 - Protocol support: bridged, routed or switched – Gateways needed?
 - Legacy support: equipment, protocols or services
 - Specific availability needs? 24-hour/backup links, etc.
 - Five-year plan – changes to population or business requirements
 - Budgetary constraints
 - Greenfield or existing site
- Information is refined and leads to a requirements database and capacity plan

Site Constraints

- Greenfield site
 - Greenfield sites have no legacy constraints
 - It is difficult to determine the real network loads and stresses
 - Needs more detail of application use and underlying protocols
 - Could use simulation to predict performance
- Existing site
 - Limited access
 - Access to live network could be restricted
 - Bottlenecks more obvious
 - Can use traffic/network analysis tools

Planning

- Uses information on
 - Hosts, users, services, and their internetworking needs
- Iterative process of
 - Conceptual design
 - Analysis
 - Refinement
- Involving
 - Brainstorming, design reviews, modelling tools
- Leading to final draft design

Design Specification & Implementation

- Detailed document of the design
 - Acts as a benchmark for design changes
 - Final design choices and changes need justification and documenting
 - Should include change history to aid maintenance
 - Used for the implementation
- Needs a project plan to include
 - Phased introduction of new technology
 - Educating the users (what to expect)
 - Pilot installation (test for possible problems)
 - Acceptance testing (to prove performance meets requirements)
 - Deployment (provide support on going live and provide fall-back position)


Connectivity Options

- Technology choices
 - LANs (Ethernet, Token ring, ATM)
 - MANs (FDDI, SMDS, ATM, SONET/SDH)
 - WANS (Frame relay, ATM, ISDN, X.25, PDCs, Satellite)
 - Wireless (802.11, Bluetooth, GPRS, GSM)
 - Dial-up lines
 - Serial links
- Determinants
 - Packet, cell or circuit switching
 - Wired or wireless
 - Distance
 - Performance
 - Bandwidth
 - Quality of Service
 - Availability

Design Approaches

Building Block Approaches

Building Block Approaches



Requirements: Boring, However
Essential

Requirements?


- Determining the requirements for a network probably isn't as much fun as shopping for really expensive hardware
 - that may be why many networks are poorly designed – no one bothered to think through their requirements!
 - Many people will jump to a specific technology or hardware solution, without fully considering other options – the obvious solution may not be the best one
- **Low level** design and **higher level** architecture must be developed
 - The environment in which the network operates is important
- The design chosen must be proved that it is '**just right**'
 - Is that \$2 million network backbone really enough to meet our needs?
 - How do we know \$500,000 wouldn't have been good enough?

Requirements?

- Part of this process is managing the customer's expectations
- A system approach must be used for understanding the network
 - The system goes far beyond the network hardware, software, etc.
 - Also includes understanding the users, applications or services, and external environment
- How do these need to interact?
- What does the rest of the organization expect from the network?


Requirements?

- Consider how devices communicate



Requirements?

- Consider how devices communicate




Requirements?

- What services are expected from the network?
 - Typical performance levels might include capacity, delay time, reliability, etc.
 - Providing 1.5 Mb/s peak capacity to a remote user
 - Guaranteeing a maximum round-trip delay of 100 ms
 - ...
- Functions include
 - Security, accounting, scheduling, management, etc.
 - Defining a security or privacy level for a group of users or an organization

Requirements?

- Service requirements could include the QoS guarantees
 - ATM
 - Intserv
 - Diffserv
 - ...
- This connects to network management monitoring of network performance




Requirements?

- Service requirements could include the QoS guarantees
 - ATM
 - Intserv
 - Diffserv
 - ...

- This

- **Bandwidth** is the theoretical capacity of some part of the network
- **Throughput** is the actual capacity, which is less than the bandwidth
 - Due to protocol overhead, network delays, etc.
- Actual capacity is always less than advertised, due to formatting




Requirements Analysis

- Requirements can come from many aspects of the network system
 - User Requirements
 - Application Requirements
 - Device Requirements
 - Network Requirements
 - Other Requirements

User Requirements

- User requirements are often qualitative and very high level
 - What is 'fast enough' for download? System response (RTT)?
 - How good does video need to be?
 - What's my budget?
 -




Application Requirements

- What types of applications are we using?
 - Mission-critical
 - Rate-critical
 - Real-time and/or interactive
- How sensitive are applications to remote access (reliability, maintainability, availability)?
- What capacity is needed?
- What delay time is acceptable?

Application Requirements

- What groups of applications are being used?
 - Telemetry/command and control - remote devices
 - Visualization and simulation
 - Distributed computing, web development,
 - ...
- Where are the applications located?
 - Are some only used in certain locations?




Device Requirements

- What kinds of devices are on your network?
 - Generic computing devices include
 - Normal PCs, Macs, laptops, handheld computers, workstations, etc.
 - Servers include all flavors of server
 - File, print, app/computation, and backup
 - Specialized devices include extreme servers (supercomputers, massively parallel servers), data collection systems, industry-specific devices, networked devices (cameras, tools), stoplights, ATMs, etc.


Device Requirements

- Specialized devices are often location-specific



Device Requirements

- We want an understanding of the device's performance – its ability to process data from the network
 - Device I/O rates
 - Delay time for performing a given application function




Device Requirements

- Performance results from many factors
 - Storage performance, that is, flash, disk drive, or tape performance
 - Processor (CPU) performance
 - Memory performance (access times)
 - Bus performance (bus capacity and arbitration efficiency)
 - OS performance (effectiveness of the protocol stack and APIs)
 - Device driver performance

Device Requirements

- The device locations are also critical
 - Often generic devices can be grouped by their quantity
 - Servers and specialized stuff are shown individually



Network Requirements

- Network requirements are the requirements for interacting with existing network(s) and network management concerns
- Most networks have to integrate into an existing network, and plan for the future evolution of the network
- Issues with network integration include
 - Scaling dependencies
 - How will the size of the existing network affect the new one?
 - Will the existing network change structure, or just add on a new wing?
 - Location dependencies
 - Interaction between old and new networks could change the location of key components

Network Requirements

- Performance constraints
 - Existing network could limit performance of the new one
- Network, system, and support service dependencies
 - Addressing, security, routing protocols and network management can all be affected by the existing network
- Interoperability dependencies
 - Changes in technology or media at the interfaces between networks need to be accounted for, as well as QoS guarantees, if any
- Network obsolescence
 - Do protocols or technologies become obsolete during transition?
- Network management and security issues need to be addressed throughout development

Network Requirements

- Network management and security issues need to be addressed throughout development
 - How will the network be monitored for events?
 - Monitoring for network performance?
 - What is the hierarchy for management data flow?
 - Network configuration?
 - Troubleshoot support?
- Security analysis can include the severity (effect) of an attack, and its probability of occurrence

Other Requirements

- Requirements can come from other outside sources
 - Your customer, legal requirements, larger scale organization (enterprise) requirements, etc.
- Additional requirements can include
 - Operational suitability
 - How well can the customer configure and monitor the system?
 - Supportability
 - How well can the customer maintain the system?
 - Confidence
 - What is the data loss rate when the system is running at its required throughput?

Other Requirements

- Financial requirements can include not only the initial system cost, but also ongoing maintenance costs
 - System architecture may be altered to remain within cost constraints
 - This is a good reason to present the customer with design choices, so they see the impact of cost versus performance
- Enterprise requirements
 - Typically include integration of your network with existing standards for voice, data, or other protocols

Requirements Specification & Map

- A requirements specification is a document which summarizes the requirements for (here) a network
 - Often it becomes a contractual obligation, so assumptions, estimates, etc. should be carefully spelled out
- Requirements are classified by Status, as noted earlier (core/current, future, rejected, or informational requirement)
- Priority can provide additional numeric distinction within a given Status (typically on a 1-3 or 1-5 scale)
- Sources for Gathering requirements can be identified, or give basis for Deriving it
- Type is user, app, device, network or other

Requirements Specification & Map

- A requirements specification is a document which summarizes the requirements for (here) a network
 - Often it becomes a contractual obligation, so assumptions, estimates, etc. should be carefully spelled out
- Requirements are classified by Status, as noted earlier (core/current, future, rejected,


Requirements Specification							
ID/Name	Date	Type	Description	Gathered/Derived	Locations	Status	Priority

on a 1-3 or 1-5 scale)

- Sources for Gathering requirements can be identified, or give basis for Deriving it
- Type is user, app, device, network or other

Requirements Specification & Map

- Requirements Mapping can show graphically where stuff is, what kind of apps are used, and existing connectivity



Campus Networks Design

Terminology

- How does switch work?
- How does Hub work?
- How does bridge work?
- How does router work?

Terminology

- **Workgroup switch**

- In a local area network (LAN), a **workgroup switch** is a relatively low capacity switch that serves the needs of a **workgroup**, or **small group of workers** who generally are geographically clustered
- A workgroup switch is the LAN equivalent of an edge switch in a public wide area network (WAN)

Used symbol




Samples



Terminology

- **High-End switch / Network switch**
 - A network switch (also called **switching hub**, **bridging hub**, officially **MAC bridge**) is a computer networking device that connects devices together on a computer network by using **packet switching** to receive, process, and forward data to the destination device
 - Ethernet switches are the most common form of network switch

Used symbol



Samples




Terminology

- **Multilayer switch**

- A multilayer switch is a network device that has the ability to operate at higher layers of the OSI reference model
- Unlike the Data Link Layer (DLL) traditionally used by switches. A multilayer switch can perform the functions of a **switch** as well as that of a **router** at incredibly fast speeds

Used symbol



Samples




Route Architecture

- Where do we route?
 - At the point where we want to limit our layer-2 broadcast domain
 - At IP subnet boundary
 - We can create more complex topologies using routers and at the same time keep things simple
 - If we start with the right topology it will make our network more stable
 - Use a hierarchical approach that makes good use of your traffic patterns and IP address allocations
 - Be aware that topology and logical design are not the same

Route Architecture

- What is the right topology?
- Continue to think of three layers
 - Access
 - Distribution
 - Core
- Thinking of layers helps reduce convergence time because of the scope of information to process These layers should not be confused with your L2 architecture

Route Architecture



Route Architecture

- **Access Layer**
 - Minimum routing information
 - Feeds traffic into the network
 - Link sizing
 - Provides network access control
 - No spoofing
 - No broadcast sources
 - No directed broadcasts
 - Provides other edge services
 - Tagging for QoS
 - Tunnel termination
 - Traffic metering and accounting
 - Policy-based routing

Route Architecture

- **Distribution Layer**
 - Goals
 - Isolates topology changes
 - Controls the routing table size
 - Aggregates traffic
 - Strategies
 - Route summarization
 - Minimize the number of connections to the core


Route Architecture

- **Core Layer**
 - Goal
 - Forwarding packets fast
 - Strategies
 - Clear of network policies
 - Every device has full reachability to every destination
 - Facilitates core redundancy
 - Reduces suboptimal routing
 - Prevents routing loops


Route Architecture

- Depending in how large your campus is you could use the typical hierarchical model or a subset
 - Two collapse core models
 - Single router acts as the network core
 - All other routers in the distribution layer
 - Single router acts as the network core
 - No distribution layer
 - All access layer routers connected to the core

Route Architecture




Traditional Campus Networks



Traditional Campus Networks

- Campus Network
 - A building or group of buildings connected into one enterprise network that consists of or more LANs
 - Generally deploy a campus design that is optimized for the fastest functional architecture over existing wire
- Network Administrator Challenges
 - LAN run effectively and efficiently
 - Availability and performance impacted by the amount of bandwidth in the network
 - Understand, implement and manage traffic flow
- Useful services
 - Broadcasts, multicast traffic (traffic propagated to a specific group of users on a subnet), video conferencing, multimedia traffic
 - Security and traffic flow


80/20 or 20/80 Rule



- Traditional 80/20 rule
 - 80% traffic local to subnet, 20% remote
- Remote traffic
 - Traffic across the backbone or core to enterprise servers, Internet, remote sites, other subnets (more coming)


80/20 or 20/80 Rule

The New 20/80 Rule



- New 20/80 rule
 - 20% traffic local to subnet, 80% remote
- Traffic moving towards new 20/80 rule due to
 - Web based computing, servers consolidation of enterprise and workgroup servers into centralized server farms due to reduced TCO, security and ease of management


Evolving Campus Structure




Address Space

- What to do about your address space
 - Assign it as you need it WRONG!
 - Poor summarization has an impact on your network's stability
 - Very difficult to correct poor allocations
 - Spend some time thinking about how you will assign address space
 - Routing stability is affected by the number of routes propagated through your network

Address Space



Address Space



Address Space

- Where should you summarize?
 - Only provide full topology where it is needed
 - Core routers don't need to know about every single network
 - Access routers don't need to know how to get to every other network
 - They should only carry enough information to reach one (or a couple of) distribution router(s)
 - Summarize at the hierarchy edges
 - Distribution layer to core
 - Distribution layer to access

Address Space

- Strategies for Successful Addressing
 - First come, first serve
 - Start with a large pool and hand them out as needed
 - Politically
 - Divide the space so each group within the organization has a pool of addresses available
 - Geographically
 - Divide the space so that every location has a pool of addresses available
 - Topologically
 - Assign addresses based on the point of attachment to the network (maybe same as geographically)

High Availability

- How can we achieve high availability?
 - Introduce hardware resiliency and backup paths into your network
 - Depending on the layer, you will use techniques differently
 - The idea is to protect your network against a single device failure affecting all of your network
 - Direct relationship between reliability, complexity and costs
 - The trick is to balance all variables and come up ahead

High Availability

- Evaluate your needs
 - Minimal need
 - Network just needs to be up for a portion of the day
 - Downtime is easily schedule after working hours
 - Business is not impacted if the network is down
 - Users' productivity is not impacted by a network failure
 - Medium need
 - Network needs to be available for most of the day
 - Only centralized servers need to be up 24 hours/day
 - Downtime needs to be scheduled on weekends
 - If critical parts of the network fail, the business operation is impacted
 - A network failure affects user productivity

High Availability

- High need
 - Network needs to be up 24x7
 - Downtime are scheduled well in advance and completed within schedule
 - A network failure causes major loss of business
 - User productivity drastically impacted by a network failure

High Availability

- Methods
 - Component Redundancy
 - Duplicate or backup parts
 - Power supplies, fans, processors, etc.
 - Have spares handy
 - Server Redundancy
 - Protect your data with backups
 - Use of hot standby servers
 - Or better yet use load balancers to distribute access
 - Network Link & Data Path Redundancy
 - Provide physical redundant connections between devices
 - Allow for hot backup paths (STP) and parallelism (routing)

High Availability

- Core layer
 - Build a dual router core and provide dual paths to it from your distribution layer
 - These could be either L2 or L3 paths
 - Make sure that you have redundant power supplies in your devices
 - This also assumes two different sources of power
 - Think of UPS protected circuits
 - Maybe even a power inverter solution for emergencies
 - Think about the possibility of dual routing/forwarding engines
 - Weight this against the use of two devices
 - Or just throw that in there as yet another layer of reliability

High Availability

- You want to also balance
 - Reduction of the hop count
 - Reduction of the available paths
 - Increase of the number of failures to withstand
- Easy to do in a single location but complexity and costs directly proportional to the number and distance between the locations

High Availability

- Distribution Layer
 - Provide dual connections to the core
 - Or provide a redundant link to other distribution layer devices
 - Doubles the core's routing table size
 - Possible use of the redundant path for traffic transiting the core
 - Preferring the redundant link to the core path
 - Routing information leaks
 - Allow for dual-homing of Access layer devices


High Availability

- Make sure that you have redundant power supplies in your devices
 - This also assumes two different sources of power
 - Think of UPS protected circuits
 - Maybe even a power inverter solution for emergencies
- Think about the possibility of dual routing/forwarding engines
 - Weigh this against the use of two devices
 - Or just throw that in there as yet another layer of reliability
 - Increases the cost of the distribution layer


High Availability

- Access Layer
 - Same challenges and solutions as the distribution layer
 - Dual home to the same distribution layer branch
 - Make sure to restrict destinations advertised to prevent transit traffic through the access layer router
 - Alternate path to another access layer device
 - Don't use the redundant link for normal traffic
 - Make sure to restrict destinations advertised to prevent transit traffic through the access layer router
 - Dual home to different distribution layer branches
 - Don't use the redundant link for normal traffic
 - Make sure to restrict destinations advertised to prevent transit traffic through the access layer router


High Availability



High Availability



High Availability



High Availability

- So I built all this redundancy and high availability in my network, how can my end users take advantage of it?
- You are already providing more than one router for a segment
- You want to provide your users with a way to move their traffic from one default gateway to another
- If one of the routers fails the other one will continue to provide services to the segment
- Be aware that redundancy is not the same as load balancing

High Availability

- How can we accomplish that?
 - Have the routers do proxy-ARP
 - Run a routing protocol between your workstations and the routers
 - Split your workstations into two groups
 - One uses one router as its default gateway
 - The other group uses the other router
 - Use ICMP Router Discovery Protocol (IRDP)
 - There is got to be a better and simpler way to do this

High Availability


- Current solutions:
 - Hot Standby Redundancy Protocol – HSRP (Cisco Proprietary, RFC2281)
 - Virtual Router Redundancy Protocol – VRRP (RFC3768)
 - Gateway Load Balancing Protocol – GLBP (Cisco Proprietary)

High Availability

- The concept is very similar
 - Workstations get configured with a single default gateway
 - The routers in the segment will negotiate who will provide services to the workstations and keep track of the state of the other routers
 - In the event of a primary/active router failure, one of the standby routers will take over the task of forwarding traffic for the workstations and become the primary/active
 - Traffic to the workstations will go to the primary/active router
 - Incoming traffic into the segment will follow the routing decisions made by routers in the network


High Availability

- HSRP




High Availability

- VRRP



High Availability

- GLPR




High Availability

- Which one should I use?
 - They all allow for a common default gateway and MAC address
 - VRRP is standardized
 - HSRP/GLBP are Cisco proprietary
 - GLBP provides load balancing
 - HSRP/VRRP do not (without introducing complexity)
 - GLBP/HSRP can track an uplink interface
 - VRRP does not


High Availability

- VRRP can reuse the default gateway IP
 - HSRP does not
- HSRP/GLBP support IPv6
 - VRRP does not yet
- VRRP uses protocol 112 & 224.0.0.18
 - HSRP uses UDP/1985 & 224.0.0.2
 - GLBP uses UDP/3222 & 224.0.0.102


IP Route Lookup




IP Route Lookup




IP Route Lookup



IP Route Lookup



IP Route Lookup



IP Route Lookup

