Network Scalability and Resilience

Afnog 2001

Many networks look like this...

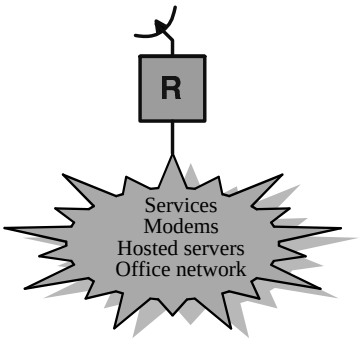

Problem: Layer 2 networks do not scale

- Volume of broadcast traffic
- Vulnerability to broadcast storms
- Machines can steal IP addresses (accidentally or maliciously)
- Huge MAC tables and ARP tables
- Troubleshooting is hard
  - No "layer 2 traceroute"
- Building resilience is hard
  - Spanning tree is no match for OSPF

Need to separate networks at Layer 3

- Multiple IP subnets
- Separate different classes of machines - especially different levels of trust
  - Access networks: for customers to connect to the Internet (leased lines, modems etc)
  - Service networks: machines which we own and manage (mail servers etc)
  - Hosted servers: machines which customers own but locate in our facilities
  - Office network - should be firewalled anyway
- Can also gain some resilience
  - e.g. put DNS caches on separate networks

Solution 1

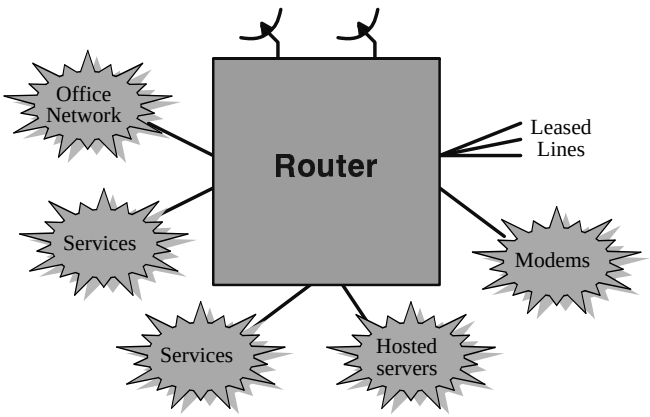

Solution 1

- Pros
  - Simple to build
  - Resilience is the vendor's problem
- Cons
  - Expensive
  - Wasteful - either buy a bigger/faster router than you need now, or throw away and upgrade
  - Selection of line cards may be limited
  - No router is resilient against software bugs and reloads

Solution 2

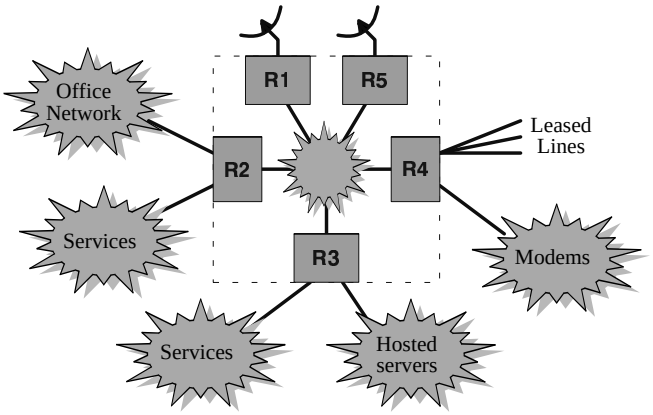

Solution 2

- Maintain investment in existing equipment and add to it as required
- Mix interface types and/or vendors
- Upgrading is less intrusive
- Careful design can give a high level of resilience