


Introduction to Routing

This presentation covers the basic concepts of routing including types of routing, routing protocols, and the default route. Examples are given of routing tables, how routing progresses as a network gradually comes on line, and what happens when there is a network outage.

- What is routing?
- What is a router?




Routing

In relation to a TCP/IP network: The act of getting a packet from the source to the destination over some arbitrary intermediate architecture. This includes choosing the best path based on some pre-defined criteria and forwarding the packet to the next step along the way.

Router

In relation to a TCP/IP network: A device which provides the function of routing; a device that determines the best next step to the destination and forwards a packet to that best next step.

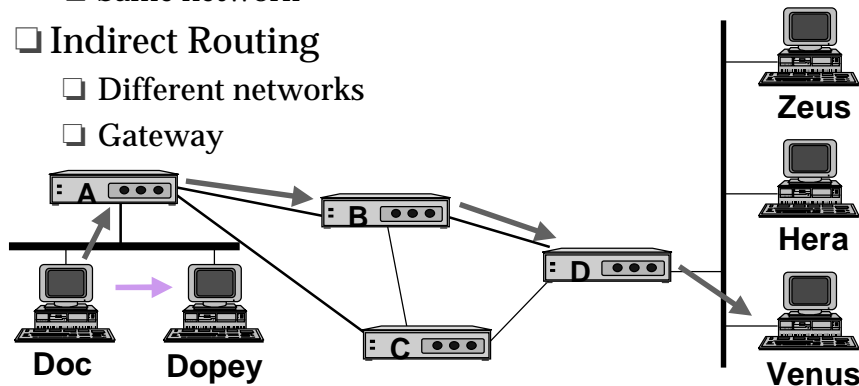
- Direct Routing
 - Same network



In a **direct routing** situation packets can be passed directly from the source to the destination; there are no intermediate steps or routers.

In the example above, packets go from Doc directly to Dopey without traversing any intermediate devices.

- Direct Routing
 - Same network
- Indirect Routing
 - Different networks
 - Gateway



Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T03.4

Indirect routing is more complicated than direct routing. The source and destination are not on the same network so packets must pass through one or more intermediate gateways. The source must identify an initial step. At each step the router must determine the best next-hop and forward the packet to that router.

As an example, Doc cannot deliver packets directly to Venus. The packets must be routed indirectly through other routers. A packet will be routed indirectly until it reaches a gateway that can deliver the packet directly to the destination. In the above example, D will deliver the packet directly to the destination Venus.



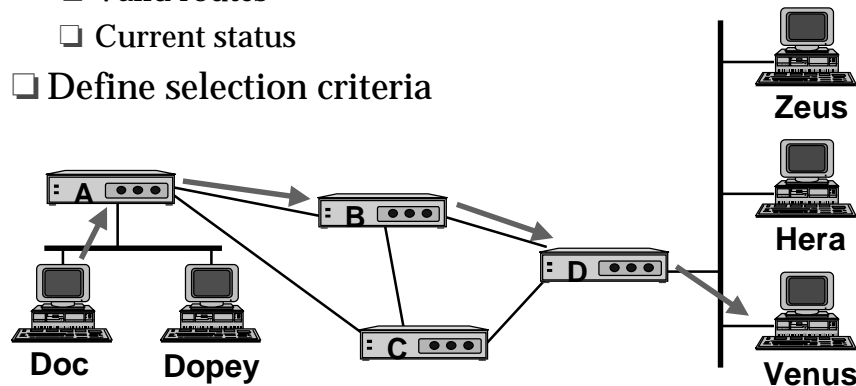
Static vs. Dynamic

- Static Routing
 - Pre-configured
 - Unchanging
- Dynamic Routing
 - Changes with the environment

Static routing is “hardwired.” Each packet always takes the same path. The routing information is predetermined and pre-configured into a device.

Dynamic routing is not pre-configured. The routing information changes as the state of the network changes. Packets that are routed dynamically can take different paths between the same source and destination.

- Informational exchange via routing updates
 - Valid routes
 - Current status
- Define selection criteria



A **routing protocol** is the language that allows routers to communicate and share information. Information is exchanged during **routing updates**. The frequency and exact format of the routing update is specified by the routing protocol.

In addition to communicating valid routes with a measure of their worth, routing updates can also provide information about the state of the network such as **load** and **error rate**.



Selection Criteria

- Possible criteria
 - Network Load
 - Packet Length
 - Shortest Paths
 - Error Rate
- Protocol specific

Each routing protocol has a different way of measuring how good a route is; the best next hop may be different depending upon which routing protocol is used.

Some of the factors that can be used to measure the validity of a route are:

Network load: how busy is the path?

Packet length: will the packet have to be segmented?

Shortest path: how many intermediate gateways must be traversed?

Error rate: how many errors are occurring on the path?

- Destination - Next hop - Metric
- Destination is usually a network, not a host.
 - Feasibility
 - Performance
- Next hop is a neighbor - generally direct.
- Metric measures route.

The basic components of a routing table are:

Destination - the destination IP address of the network.

Next hop - the IP address of the neighbor that the packet should be forwarded to.

Metric - measure of how good the given entry is.

A routing table cannot store an individual entry for every given host; multiple host entries must be condensed into a single network entry. This results in a smaller routing table which improves performance of the router.

In most cases, the next hop must be on a directly connected network, that is, the router must be able to deliver the packet **directly** to the next hop. Some routing protocols, such as **BGP** (Border Gateway Protocol) operate differently and do not enforce this restriction.

The next hop is always a neighbor (another router or device that is exchanging information with the given router) or the destination host.

The destination IP address of the packet is compared to the destination IP address of the network in each routing table entry. The matching entry with the best (lowest) metric is chosen. The packet is then forwarded to the next hop specified in the given routing entry.

PSINet *The Smaller the Better*

Router B

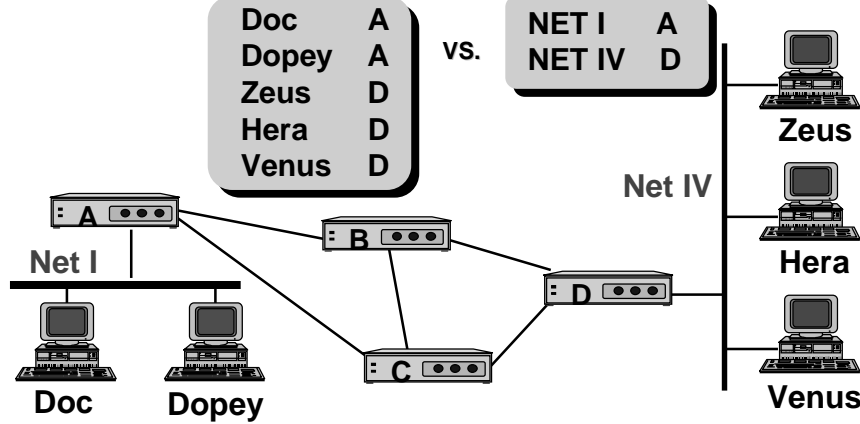
By Host:

Doc	A
Dopey	A
Zeus	D
Hera	D
Venus	D

vs.

By Network:

NET I	A
NET IV	D

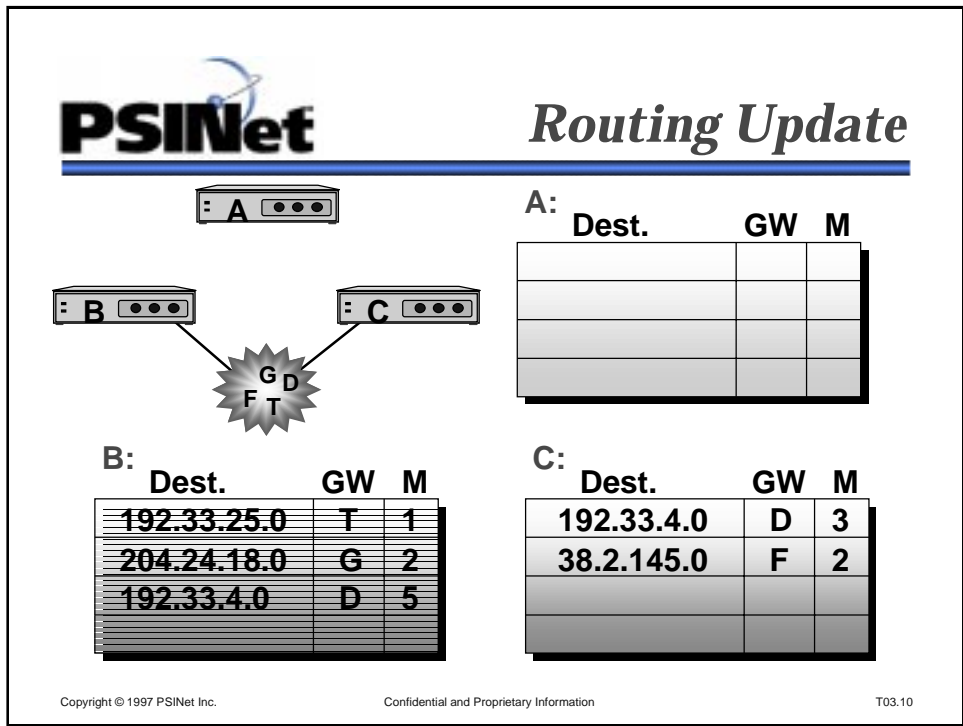


Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T03.9

This example illustrates the benefit of routing tables based on host information versus routing tables based on network information, even for a simple network.




Consider the above situation - A, B and C are routers. C and B route traffic to a collection of networks through one or more connections. The collection of networks is represented by the star shape. D, F, G and T are routers contained within the star.

The current routing tables for B and C are detailed above.

Dest. is the destination network, GW is the gateway or next hop, and M is the metric. Assume the metric is a simple hop count and multiple entries for a single destination network are not allowed.

In this example, GW is the router, for example D. In reality, the GW would be the IP address of an interface on router D. The IP address for D would belong to a network that C and D have in common; packets can be delivered directly between C and D over this common network.



A:

Dest.	GW	M


B:

Dest.	GW	M
192.33.25.0	T	1
204.24.18.0	G	2
192.33.4.0	D	5

C:

Dest.	GW	M
192.33.4.0	D	3
38.2.145.0	F	2

What happens when A connects to B and C?



A:

Dest.	GW	M
192.33.25.0	B	2
204.24.18.0	B	3
38.2.145.0	C	3
192.33.4.0	C	4

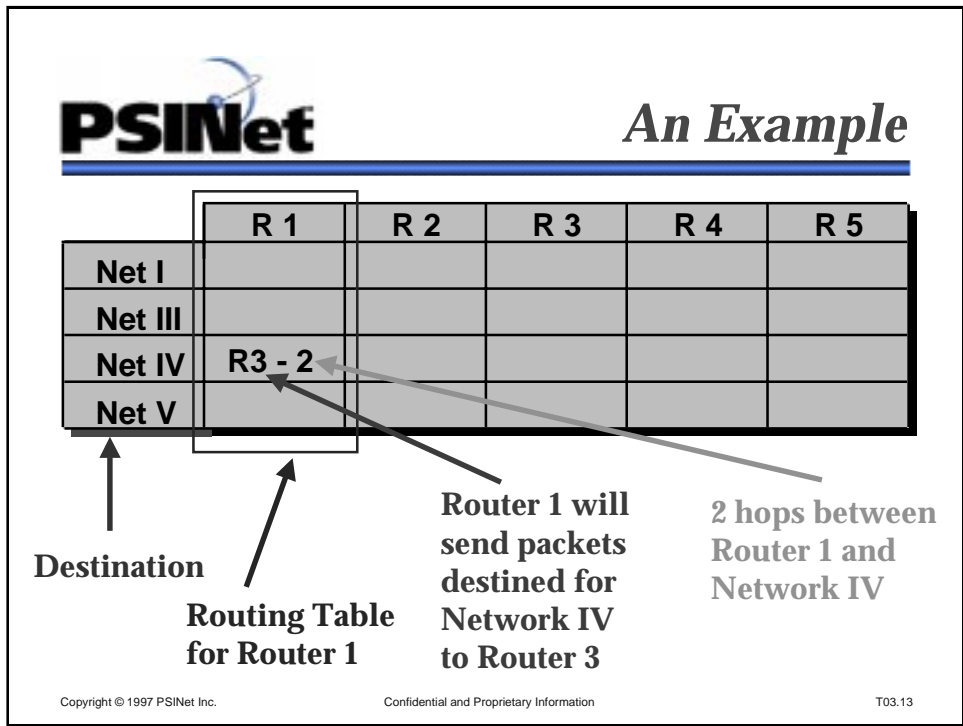
B:

Dest.	GW	M
192.33.25.0	T	1
204.24.18.0	G	2
192.33.4.0	D	5

C:

Dest.	GW	M
192.33.4.0	D	3
38.2.145.0	F	2

Examine the resulting routing table for router A. Notice that both B and C advertise a route to the network 192.33.4.0. Router A has to examine both routing advertisements and choose the one with the best route, in this case the one with the fewest intermediate steps.



First, an explanation of the notation used in the following example.

The columns represent a routing table for the router indicated in the label for the column (R1 -> Router 1, R2 -> Router 2, etc.).

The rows represent the destination portion of the routing table. For example, the row labeled Net I contains all the routing information for packets destined for Network I.


Within a particular block there may be 2 pieces of information:

Next hop or gateway - where a packet destined for the given network will be forwarded.

Metric - how many hops (routers) between the router and the destination network.

A block may also contain a special entry, **direct**. This entry means the destination network is directly connected to the router; the router can route to the destination using direct routing.


	R 1	R 2	R 3	R 4	R 5
Net I					
Net III					
Net IV					
Net V					



The following example will illustrate how routing tables are updated as the network changes. The routers will be turned on one at a time and the resulting routing tables will be detailed.


With all routers off, there is no routing information.

	R 1	R 2	R 3	R 4	R 5
Net I					
Net III			Direct		
Net IV					
Net V					



When R3 is turned on, it begins advertising a route to Net III. Currently, there is nobody listening to R3's advertisements.

	R 1	R 2	R 3	R 4	R 5
Net I	Direct		R1 - 1		
Net III	R3 - 1		Direct		
Net IV					
Net V					



Copyright © 1997 PSINet Inc.


Confidential and Proprietary Information

T03.16

When R1 is turned on, it begins advertising a route to Net I. Since R1 and R3 are neighbors, they will hear each other's routing advertisements and update their routing tables accordingly.


R1 will learn that Net III is 1 hop away via R3. R3 will learn that Net I is one hop away through R1.

	R 1	R 2	R 3	R 4	R 5
Net I	Direct		R1 - 1	R3 - 2	
Net III	R3 - 1		Direct	R3 - 1	
Net IV			R4 - 1	Direct	
Net V					



When R4 is turned on, it will exchange information with its neighbors - in this case, just R3. R3 and R4 will update their routing tables based on information they learn from each other.

	R 1	R 2	R 3	R 4	R 5
Net I	Direct		R1 - 1	R3 - 2	
Net III	R3 - 1		Direct	R3 - 1	
Net IV	R3 - 2		R4 - 1	Direct	
Net V					




Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T03.18


During the next routing update, R3 and R1 will exchange information. Only at this point will R1 learn about the route to Net IV. R3 has to learn about Net IV and pass this information to R1 during a subsequent routing update. This illustrates propagation delay - routing information is passed from router to router, it is not learned simultaneously by all routers in the network. (Some routing protocols function differently. This will be discussed in the Routing Protocols presentation.)

	R 1	R 2	R 3	R 4	R 5
Net I	Direct		R1 - 1	R3 - 2	R4 - 3
Net III	R3 - 1		Direct	R3 - 1	R4 - 2
Net IV	R3 - 2		R4 - 1	Direct	R4 - 1
Net V	R3 - 3		R4 - 2	R5 - 1	Direct



This is what the routing tables will look like with R5 turned on, after information has propagated throughout the network.

	R 1	R 2	R 3	R 4	R 5
Net I	Direct	R1 - 1	R1 - 1	R3 - 2	<i>R2 - 2</i>
Net III	R3 - 1	R1 - 2	Direct	R3 - 1	R4 - 2
Net IV	R3 - 2	R5 - 2	R4 - 1	Direct	R4 - 1
Net V	<i>R2 - 2</i>	R5 - 1	R4 - 2	R5 - 1	Direct



Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T03.20

Finally, this is the state of the routing tables with all routers active. Notice that the entry in R1's routing table for Net V has been updated. Previously, R1 was routing traffic for Net V through R3 with a metric of 3. Once R2 is active, R1 will learn there is a better route to Net V through R2 with a metric of 2. The entry for Net I in R5's routing table will be updated similarly.

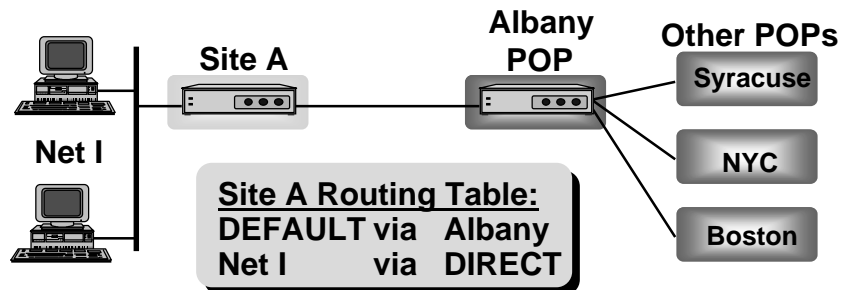


Default Route

- Special destination
- Condenses the routing table
- Packet destination IP address is compared to routing table entries.
 - If no entry matches exactly, the packet is sent via the default route.
- Ideal for many PSINet customers

The **default route** is a special routing table destination. Each packet that enters the router is examined against the routing table entries. If a specific routing table entry matches the destination IP address of the packet, the packet is forwarded using this specific routing entry. If the destination of the packet does not match the destination in any routing table entries, the packet is forwarded using the default route.

- Site A has a single decision to make.
- Outbound packets - *default* route
- What about Albany?



Examine the above example which is similar to the configuration of many PSINet customers - a single network that is to be connected to the Internet. If traffic is not destined for the local network (Net I), it should be sent to the Internet via the Site A router.

The Albany router can also use a default route but not to the same advantage. Albany has a more complicated decision to make when routing a packet - should it be sent to Site A, Syracuse, Boston or NYC? Depending upon the traffic patterns, NYC might be chosen as the default route.

- Recall the previous example

	R 1	R 2	R 3	R 4	R 5
Net I	Direct	R1 - 1	R1 - 1	R3 - 2	R2 - 2
Net III	R3 - 1	R1 - 2	Direct	R3 - 1	R4 - 2
Net IV	R3 - 2	R5 - 2	R4 - 1	Direct	R4 - 1
Net V	R2 - 2	R5 - 1	R4 - 2	R5 - 1	Direct

20 entries

Consider how the default route would effect the routing tables from our previous example.

Using the default route


	R 1	R 2	R 3	R 4	R 5
Net I	Direct	R1 - 1	R1 - 1	R3 - 2	R2 - 2
Net III	R3 - 1	R1 - 2	Direct	R3 - 1	R4 - 2
Net IV	R3 - 2	R5 - 2	R4 - 1	Direct	R4 - 1
Net V	R2 - 2	R5 - 1	R4 - 2	R5 - 1	Direct
Default	R3	R1	R4	R3	R4

15 entries

Adding a default route to the routing table for each router results in a net loss of one entry per router. This is an overall savings of 5 routing table entries - and this is just a simple example.

Notice that R2 could specify R1 or R5 as the default route with the same number of resulting entries. R1 was chosen as the default route at random.

- Dynamic Routing
- Redundancy




What allows traffic to travel on the Internet despite network outages?

Routing must be dynamic - the path to a destination must be flexible in order to get around an outage in the “normal” path.

The network must contain redundancy; if there is an outage, there must be another viable path.

- B has a route to Net1 via A.




Consider the above example.

Router B is connected to routers A and C. Routers A and C are connected to some network structure, the details of which we are not concerned with. Routers A and C may or may not be able to route traffic back and forth through the unknown network structure.


Router B will route packets destined for Network I to Router A.

- B stops hearing updates from A.



If there is an outage between Router A and Router B, they will stop hearing periodic routing updates from each other.


- B stops hearing updates from A.
 - Route to Net1 is marked invalid.
- Router B listens to its other neighbors.



If B stops hearing from A, B will mark all routing table entries with a gateway of A invalid since the route to A is no longer valid.

B will then listen to its other neighbors, in this case C, for routes to these invalid destinations.


- Router B listens to its other neighbors.
 - Routing table updated with new routes.
 - Loss of connectivity if there are no valid routes.



If C is announcing such a route, the gateway will be changed to C and the metric will be recalculated.

If C is not announcing such a route, the routing entry for the given destination will be removed from B's routing table.


- ❑ R5 sends packets to Net IV via R4.



Returning to our example network, R5 is sending packets destined for Net IV via R4.


The arrows indicate how each router is forwarding packets destined for Net IV.

❑ What if the link between R4 and R5 is down?



What happens if the link between R4 and R5 goes down?


- ❑ R5 stops hearing routing updates from R4.
- ❑ R5 marks the next-hop to Net IV (R4) invalid.
- ❑ R2 will also mark its next-hop to Net IV (R5) invalid.
- ❑ R2 will find a new route via R1



R5 will stop hearing updates from R4. The route to Net IV via R4 will be marked invalid. This invalid route status will be passed on to R2. As a result, R2 will mark its route to Net IV via R5 invalid.


R2 and R5 will both listen to other neighbors to determine another route to Net IV. R2 will learn of an alternate route via R1.

- ❑ R5 then listens to its other neighbors (R2) and finds a valid route to Net IV.



R2 will pass this alternate route to R5 in the next routing update. R5 will update its routing table so all traffic destined for Net IV is sent via R2.

- The link between R4 and R5 is back up.
 - R5 hears updates from R4 again.
 - R5 sees that the best route to Net IV is via R4.
 - R5 updates its routing table.
 - R2 updates its routing table.



What happens when the link between R4 and R5 is repaired?

R5 will again receive routing updates from R4. R5 will analyze the metrics on the routes to Net IV (metric of 4 via R2 - metric of 1 via R4). Based on this information, R5 will update its routing table to send traffic destined for Net IV to R4. R5 will then pass this updated information to R2 causing R2 to update its routing table entry for Net IV.

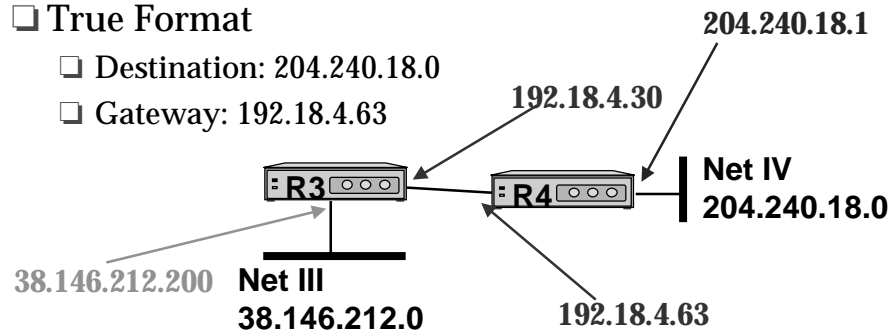
Routing will be restored to the state it was in before the outage.

Simplified Format (Router 3)

- Destination: Net IV
- Gateway: R4

True Format

- Destination: 204.240.18.0
- Gateway: 192.18.4.63



We have been using a simplified notation for destination and gateway in the preceding examples. We will now expand our notation to match real routing tables.

Each interface on a router must have a distinct IP address. In our example, Router 4 is really known by two different IP addresses, 204.240.18.1 and 192.18.4.63.

Destination networks will not be named (Net IV) but rather numbered (204.240.18.0).

Thus, Router 3 will have an entry for packets destined to 204.240.18.0 (Net IV) and the next hop or gateway will be the interface of Router 4 that Router 3 can directly route to, 192.18.4.63.

