

This workshop session will explore the key issues related to effective engineering of DPN-100 token ring ISRB solutions. Particular importance will be paid to the key factors that affect network scalability of ISRB bridged virtual rings. Broadcast traffic sources and their impact on network link traffic, and broadcast server utilization will be explored in detail. Use of a simple Wingz-based engineering tool to simplify calculation of traffic and PE utilization will be demonstrated. Interworking with Passport native IP and IPX routing function as a means of addressing scalability limits will also be addressed.

**About the presenter:**

Gary Palmer has more than 23 years experience in the data communications industry, particularly with IBM SNA protocols. His Nortel career started in 1986 in the IBM services product planning and product management areas. The last three years have been focused on developing network engineering tools and training materials, for use by Nortel systems engineers and customers worldwide.

## Agenda

- **ISRB Overview**
- **Engineering Constraints**
- **Broadcast Traffic Sources**
- **Network Topology Considerations**
- **LanCalc Tool**
- **Interworking with Routers**

2

This presentation will begin with an overview of the ISRB service, particularly from a broadcast traffic perspective as broadcast traffic tends to determine the scalability of any particular ISRB solution.

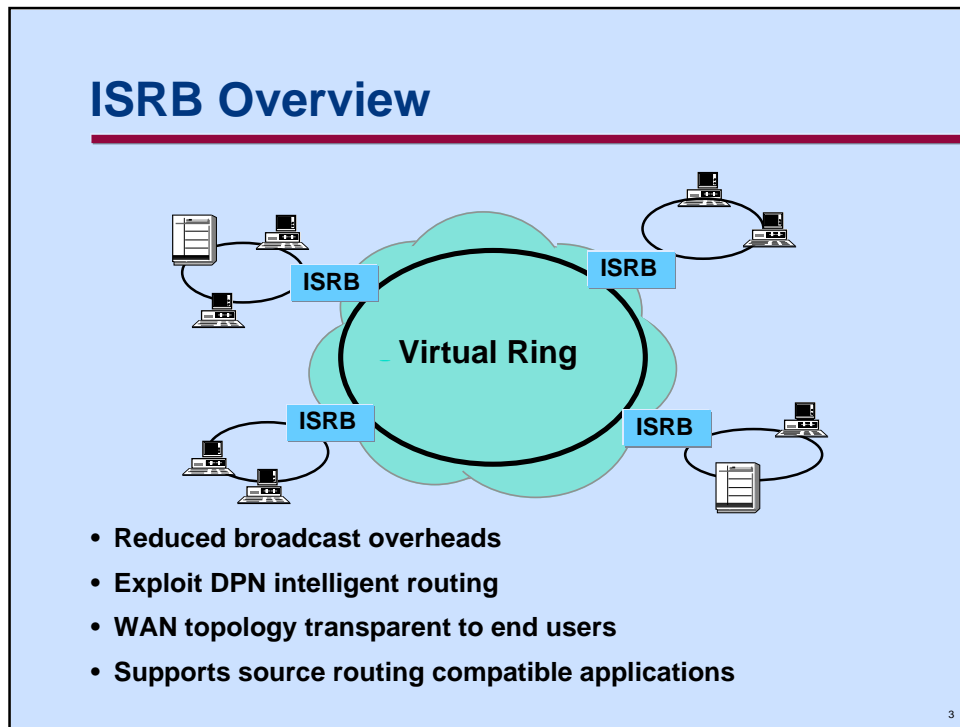
A review of engineering constraints in general, and for the ISRB in particular, will provide a necessary background for the engineering exercises and examples that follow later.

Broadcast traffic sources for the Novell Netware and IP environments will be explored.

Basic network topology considerations as they apply to placement of broadcast servers will also be explored.

The overall structure of the LanCalc tool, which reduces the burden of calculating link and PE utilizations, will be covered in some detail.

Finally an interworking scenario with router functionality provided by a Passport backbone will be introduced as a means to expand the scalability of ISRB-based token ring solutions.



This slide provides a high level overview of the **Intelligent Source Routing Bridge (ISRB)** feature. The ISRB feature can co-exist on the same PE and share a PI with the token ring SNA PAD feature.

The key aspect of this service is that it automatically links each of the remote token rings via an internal **virtual ring** creating a simplified topological view for each of the attached devices.

The actual network topology is completely hidden from the end token ring stations which behave as if they were all linked together by a single backbone ring.

This feature works with all token ring applications which use source routing but will not work with applications which assume transparent bridging capability.

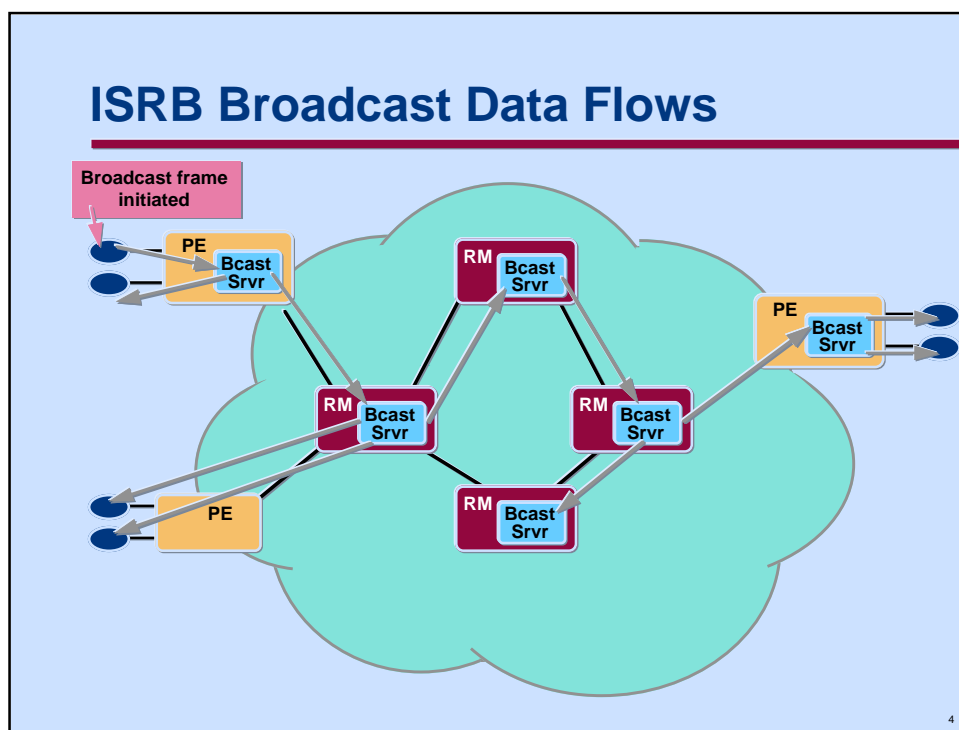
The ISRB feature makes use of an internal broadcast server hierarchy to efficiently distribute broadcast frames to all rings attached to the virtual ring.

Features such as **MAC Header Compression** and **Proxy Server** are available to reduce bandwidth required over the backbone network to a minimum.

This presentation will look at the key engineering issues that help determine the **scalability** of this solution. In other words, it will help you determine how many external rings can be bridged together to form a single virtual ring network.

Not only will the engineering issues and principles be covered, but an easy-to-use engineering tool, **LanCalc**, will be demonstrated.

Note that in subsequent illustrations, the term ISRB usually refers to a single token ring PI attached to a PE. Up to four instances of the ISRB feature can exist on a single PE.



This slide depicts how broadcast frames would flow in a typical ISRB environment overlaid on a typical simple network. This slide captures a broad range of typical configurations, and illustrates multiple configuration options.

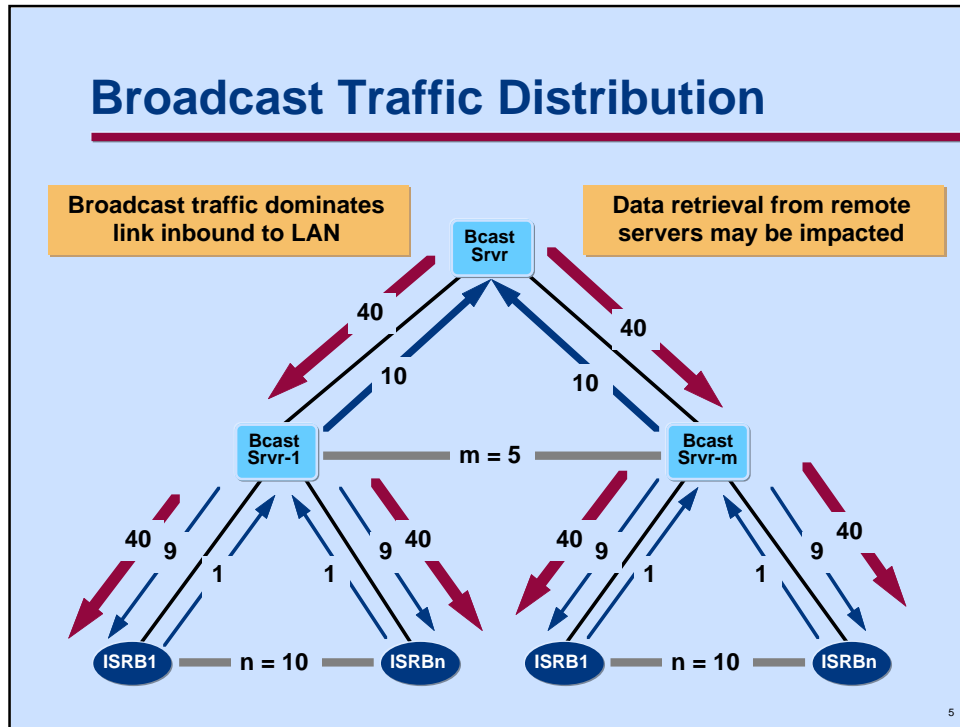
The access node at the upper left corner includes a broadcast server process on the same PE (MAS option only) which is supporting the ISRB service for two PIs. A broadcast frame originates from a station on the uppermost of these two rings.

The broadcast server on this node replicates the frame and sends it to the other ring on this same node while also forwarding it upwards to the server located on the backbone RM to which the node is connected.

The broadcast server on this backbone RM then forwards a copy to the next RM containing a broadcast server, and makes two copies to forward to the PE in the lower left. Since this PE has not been configured with a broadcast server, a copy of the broadcast frame must be generated for each ISRB Token ring PI installed on the node. This illustrates the value of providing broadcast server functionality on any access node supporting multiple ISRB token ring PIs.

Note that the logical arrangement of broadcast servers in the backbone will, in general, mirror the physical design of the backbone, but must not contain any routing loops. The network designer is responsible for defining the broadcast server topology, and will often find a logical ring with a single break is an acceptable compromise.

Since any given broadcast server processes each broadcast frame once for each VC connecting it to either another server or to subservient ISRBs; it is good practice to try to establish a balanced number of VCs for each server. Given that a PE can support up to four PIs, a value of five VCs per broadcast server is a good initial starting design point for determining the scalability of an ISRB solution



This slide depicts in a graphical manner the flows and consolidation of traffic. A simple symmetrical two level broadcast server hierarchy is used to illustrate the basic principles.

At the top of the hierarchy in this example is a single broadcast server located near the geographic center of the network. There are  $m = 5$  lower level servers connected to this top level server. At the lowest level of the hierarchy, each server supports  $n = 10$  ISRB token ring LANs; for the purposes of simplified numerical computation assume each LAN generates  $X = 1$  broadcast frames per second and each frame is **1,000 bits** in length. Thus each link connecting an ISRB LAN to its superior server forwards **1 Kbit/s** of traffic.

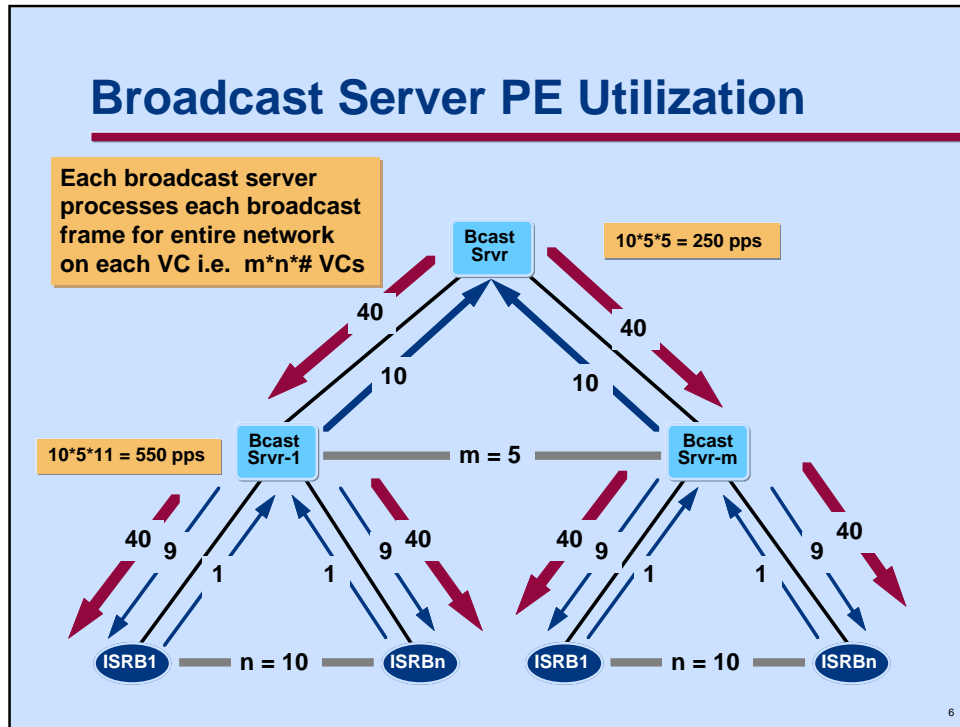
The thin arrowed lines indicate the local broadcast traffic flows within a segment of the hierarchy, while the thicker lines represent the remote broadcast traffic flows.

Each of the five first-level servers receive  $n * x = (1 * 10) = 10$  Kbit/s of traffic from the 10 attached LANs. This 10 Kbit/s of traffic is forwarded on each of the physical links traversed by the **5 Virtual circuits** connecting the lower level servers to the top level server. Nine Kbit/s of local traffic will flow from the first level server back down to each of the source LANs

The top level server will distribute each broadcast frame received from a lower level server down each of its virtual circuits to the lower level servers. Consequently each link will see  $(m-1) * n$  i.e.  $(5-1) * 10 = 40$  Kbit/s of traffic downward.

At the bottom of the hierarchy each link will receive broadcast traffic from every other LAN except its own. This is computed as  $((m * n) - 1) * X$  i.e.  $((5 * 10) - 1) * 1 = 49$  Kbit/s of traffic.

From this simple example it is obvious that potential contention exists for traffic **inbound** to any LAN as would be the case for retrieval of data from a remote server. This characteristic may provide guidance in establishing engineering constraints as discussed later.

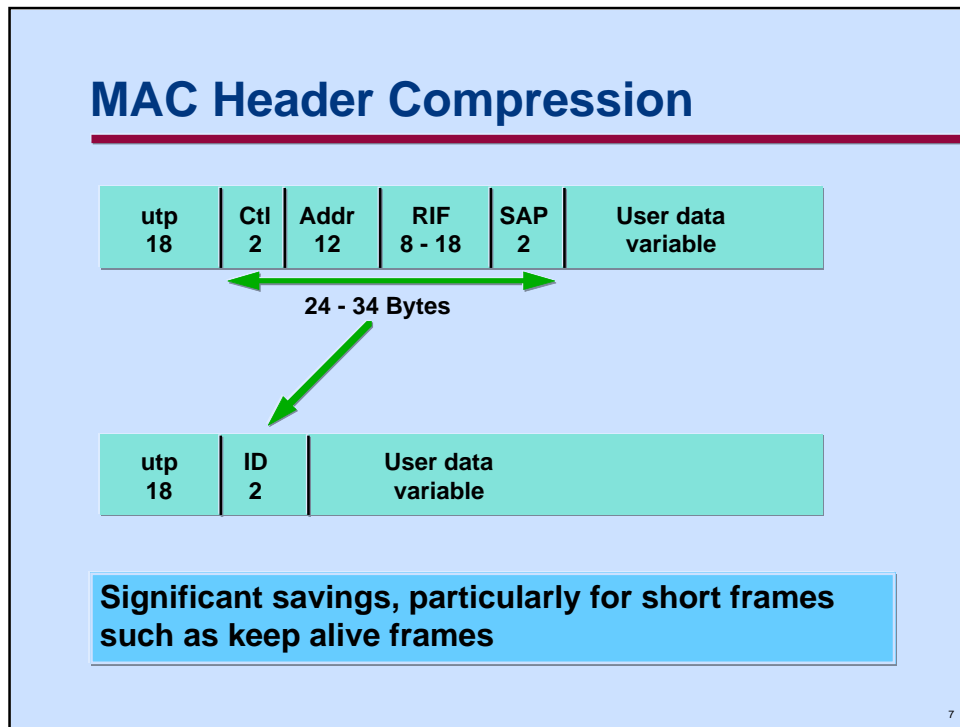


This slide looks at the prior diagram from the perspective of determining the load on the broadcast server PEs. Note that broadcast server PE utilization is a function of the packet/sec rate (pps) and the average packet size. For illustration purposes we are simply using the pps rate which is assumed to be **1 pps** per ISRB LAN.

The thin arrowed lines indicate the local broadcast traffic flows within a segment of the hierarchy, while the thicker lines represent the remote broadcast traffic flows.

For each of the five first-level servers; each broadcast frame received from an ISRB LAN is broadcast over each of the other VCs connecting the server to other ISRBs and the superior server. The 10 attached LANs contributing 1 pps each result in 10 pps of traffic being forwarded upward to the superior server, and 9 pps being sent out on each of the lower links for local broadcast traffic. In addition to locally generated traffic; the first level servers will also receive broadcast frames from the superior server. The number of frames received will be equal to the number of other first level servers ( $m-1$ ) i.e. four in this case times the number of broadcast frames passed upwards to the top level (i.e.  $10 \times 4 = 40$ ).

The text blocks adjacent to the broadcast servers show a simplified calculation based on the observation that each broadcast server will process each broadcast frame from all sources once on each VC connecting it to either ISRB LANs or to other broadcast servers. Thus with 10 LANs in each of five regions we have a total of 50 LANs generating 1 pps each = 50 pps of broadcast traffic. The lowest level broadcast servers have 11 VCs defined (10 down, 1 up) hence process 550 pps, while the top level server process  $5 \times 50 = 250$  pps.

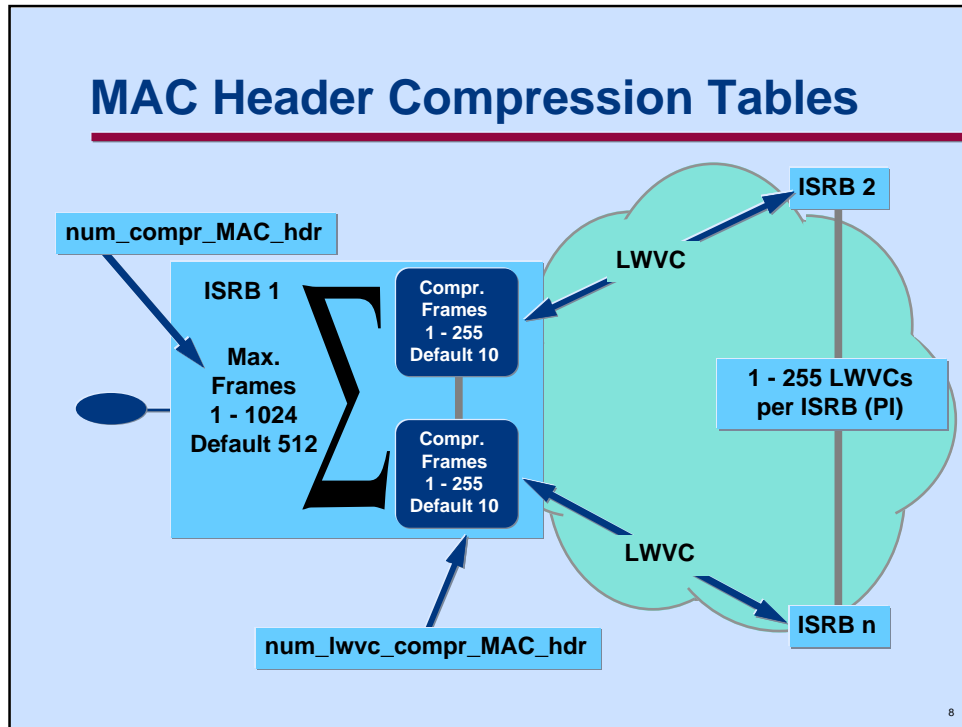


MAC Header Compression reduces the fixed header content of frames which may be between 24 to 34 bytes, to two bytes. This can be a significant savings, particularly where there are a large number of short frames repetitively transmitted such as keep alive frames. The variation in the Routing Information Field (8 to 18) bytes is a function of the number of hop counts. The RIF consists of a two byte header followed by up to eight two-byte route designators (bridge and ring number). The route designators indicate the bridge and ring to use to make the next hop; in the IBM implementation, there can be up to eight route designators. In the normal ISRB scenario with no external bridges defined, the RIF will be 8 bytes in length (i.e. two byte header plus three route designators); thus the MAC header field is normally 24 bytes.

Note that MAC header compression can be used for any specifically routed frame, unlike the Proxy feature and should be used in all cases. Limits do exist on the maximum number of frames that can be compressed, but these are controlled by user specification, and are limited only by the amount of memory available on a PE. There is a very small additional processing overhead of less than 5% PE utilization worst case for this feature that can generally be ignored.

Typical bandwidth savings range from 83% for keep alive frames, to 38% for a typical interactive login sequence with average frame sizes of 60 bytes, to 2.4% for file transfer with an average frame size of 900 bytes.

The utp header of 18 bytes (12bytes VC, 6bytes HDLC) as would exist on a network link or trunk is added to the MAC Header.



Note there are two parameters with associated limits that must be specified for each ISRB service. Remember that there is an ISRB service for each PI that supports the ISRB feature. The values which need to be specified are:

**num\_compr\_MAC\_hdr:** **Default 512;** range 1 - 1024

This field defines the maximum number of unique frames, originating from or destined to a specific local PI that will have their headers compressed.

**num\_lwvc\_compr\_MAC\_hdr:** **Default 10;** range 1 - 255

This field defines the maximum number of unique frames, originating from or destined to a specific remote ISRB(i.e. PI) that will have their headers compressed.

Each compressible unique frame is defined by a unique SMAC/SSAP, DMAC/DSAP address pair. In general this can be defined as the number of sessions established per protocol given that each protocol uses a unique SAP (i.e. X'F0' for Novell, X'E0' for Netbios, etc).

Note that there is only one entry in the LWVC (Light Weight Virtual Circuit) tables for a unique address pair. This entry has the header fields adjusted appropriately for the direction of travel of a frame (i.e. swapping of SMAC/DMAC, and the direction indicator in the Routing Information Field).

The maximum number of LWVCs defined per ISRB is specified via the provisioning parameter **max\_num\_lwvc**. The LanCalc tool keeps track of this value.



## MAC Header Compression Benefit

Transaction	Protocol	Options		Kbps Traffic		PE Utilization %			Link Util %	
		MHC	Proxy	In	Out	ISRB	Beast	Total	In	Out
<b>20/50 bytes</b>										
	TCP Interactive	N	N	2.9	2.5	0.6%	0.0%	0.6%	29.7%	25.6%
	TCP Interactive	Y	N	2.3	1.9	0.5%	0.0%	0.5%	23.6%	19.4%
<b>100/1000 bytes</b>										
	TCP Interactive	N	N	3.5	1.1	0.3%	0.0%	0.3%	36.9%	11.9%
	TCP Interactive	Y	N	3.3	0.9	0.3%	0.0%	0.3%	34.4%	9.4%

Approx 20% Savings

Approx 7% Savings

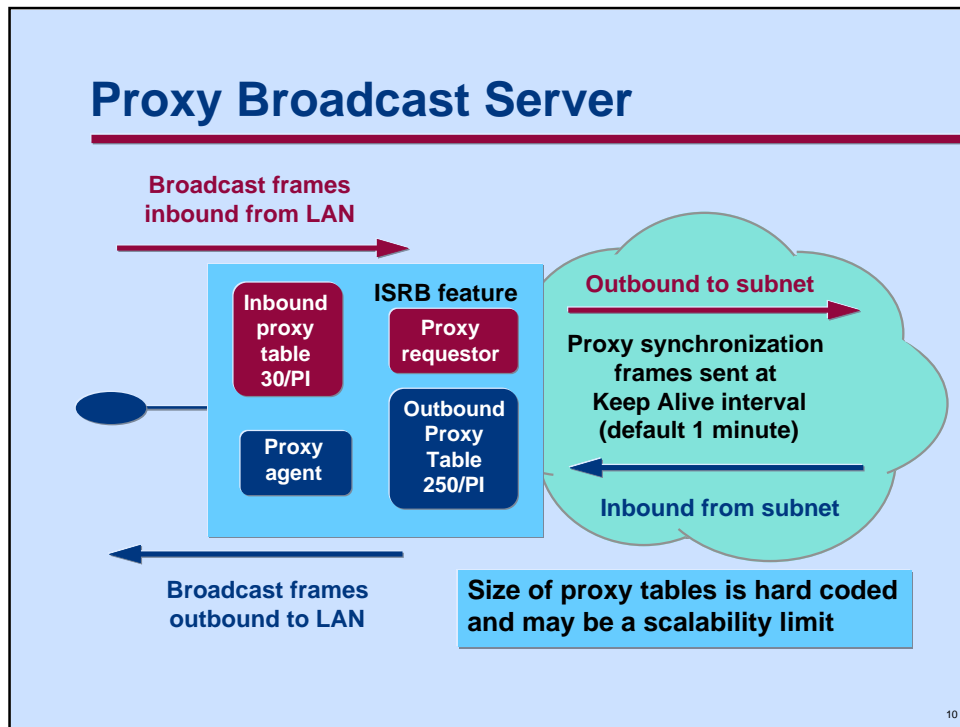
**Compression savings greatest for small frames such as LLC2 Keep Alive frames; approximately 83%**

This extract from the output of the LanCalc tool illustrates relative savings derived from use of the MAC Header Compression feature. This reduces network link utilization primarily, and has a very small beneficial impact on Access PE utilization.

The two scenarios shown are a user data 20 byte outbound, 50 byte inbound TCP interactive scenario, compared to a larger 100 byte outbound, 1000 byte inbound scenario.

From this example it is clear that there is a smaller benefit derived as frame sizes increase.

The greatest savings occur with the small MAC level keep alive frames that are generated by LAN stations using LLC2 protocol. In this extreme example, savings of 83% are experienced.



The Proxy Broadcast function can be used to significantly reduce the amount of periodic broadcast traffic flowing through the network. This diagram provides a high-level overview of the general operation of the Proxy function. Note that Proxy broadcast and MAC Header Compression are mutually exclusive options since compression only applies to specifically routed frames.

All broadcast frames received from the LAN (i.e. Inbound from LAN) are examined to determine if they are proxiabile. If they are, then a copy is stored in the **Inbound Proxy table** which is used and maintained by the Proxy Requestor. Note that there is currently a hard coded limit of **30 frames per PI** that can be stored in this table.

The Proxy requestor will then send a copy of this Proxy frame to all other Proxy agents on all ISRBs associated with the same virtual ring. This Proxy initialization frame is then stored in the **Outbound Proxy table** of all ISRBs. Note that there is currently a hard coded limit of **250 frames per PI** that can be stored in this table. The Proxy agent then broadcasts the frames in this table on to the local LAN at a fixed one minute frequency.

In the event that either of these tables overflows, then additional broadcast frames will not be proxied and the full overhead will be incurred.

After this initialization process Proxy synchronization frames are broadcast at the frequency of the Broadcast Keep Alive interval (default 1 minute). The size and number of these frames is kept to a minimum by using frame correlators and blocking multiple correlators in a single frame as will be described in more detail later.

Since the size of these tables is currently a fixed size, they may potentially become an engineering scalability constraint. The LanCalc tool will keep track of the total number of proxiabile frames to allow evaluation of this potential constraint.

## Proxy Eligible Frames

- Novell Netware RIP (Routing Information Protocol)
- Novell Netware SAP (Service Advertisement Protocol)
- IBM LanServer \MAILSLOT\NET\NETLOGON
- IBM LanServer \MAILSLOT\LANMAN
- Netbios Name Query \*\*\*

**Note: Other RIP frames such as IP RIP are NOT proxy eligible in current ISRB implementation**

11

Note that while Novell RIP frames are supported by the Proxy feature of ISRB, other RIP frames such as IP RIP, and AppleTalk RTMP frames are not. Likewise IP ARP frames are not supported.

\*\*\* **Note** that the Netbios Name Query broadcast command is a special case. In normal Netbios operation, a Name Query frame is transmitted up to six times in one second intervals whenever a Netbios station is attempting to locate another station. With ISRB, only the first of these six frames will be broadcast via the broadcast server, and the remaining 5 duplicates will be suppressed. There is transient use of the Inbound Proxy table for this function, but otherwise, this is a special case.

For the non supported frame types, only the MAC Header Compression option can be used to reduce traffic frequency.

This presentation will focus on the Novell and TCP/IP environments, consequently there will be no additional information on the Netbios environment.

## Proxy Synchronization Frames

utp 18	hdr 22	Frame correlators n x 3
-----------	-----------	----------------------------

- UTP header of 18 bytes
- Proxy broadcast header is 22 bytes
- Each proxiable frame is represented by a 3 byte frame correlator
- Maximum size of frame (excluding utp header) is limited to packet size defined between ISRB and broadcast server
- Proxy Requestor limited to 30 frames in Inbound table
- Maximum frame size =  $18 + 22 + (30 \times 3) = 130$  bytes

12

This is a high-level overview of the Proxy server/requestor frames involved in broadcast traffic as seen on a network link including the utp header of 18 bytes (12bytes VC, 6bytes HDLC).

The Proxy broadcast synchronization frame has a 22 byte header followed by a variable number of 3 byte correlators. Each Proxy requestor is limited to a maximum of 30 frames that it can Proxy on behalf of each attached LAN, consequently the maximum frame size is 130 bytes.

## Proxy Broadcast Savings

Protocol ID	Options		Kbps Traffic Summary				Cumulative Frame Counts				PE Utilization %			Link Util %	
	MHC	Proxy	Specific Route		Broadcast		Proxied		MHC	ISRB	Beast	Total	Acc +	In	Out
			In	Out	In	Out	In	Out	Total	perVC	Access	Server	Server		
Novell SAP	N	N	0.0	0.0	10.4	0.1	0	0	0	0	0.6%	11.8%	12.4%	18.6%	0.1%
Novell SAP	N	Y	0.0	0.0	1.2	0.0	1	200	0	0	0.2%	9.4%	9.6%	2.2%	0.0%
<b>Proxy Server savings</b>											<b>61.7%</b>	<b>20.3%</b>	<b>22.2%</b>	<b>88.2%</b>	<b>90.5%</b>

**Significant savings from use of Proxy Server for example:**

Access PE	60%
Server PE	20%
Access Link	90%

13

This extract from the LanCalc tool shows how significant the savings from the use of the Proxy broadcast server can be.

It also shows the advantage of using a spreadsheet base for the tool development since the relative savings calculations are simple spreadsheet formulae entered in a few seconds work.

## Agenda

---

- ISRB Overview
- **Engineering Constraints**
- Broadcast Traffic Sources
- Network Topology Considerations
- LanCalc Tool
- Interworking with Routers

14

This section looks at the possible engineering constraints that must be evaluated to determine the scalability of the ISRB solution.

## Engineering Constraints

- **Multiple possible constraints**
  - network link bandwidth
  - access PE utilization
  - broadcast server PE utilization
  - Proxy Table sizes
  - MAC Header Compression Table sizes
- **Weakest link determines scalability limit**

1. **Establish criteria for each potential constraint**
2. **Use LanCalc Tool to evaluate each resource**
3. **Determine scalability based on limiting constraint**
4. **Each network and application is unique**

15

As is usually the case, there is no single constraint that can be evaluated to determine scalability. In the ISRB case, the normal constraints of access PE and network link utilization need to be assessed as in all access services.

In addition, the impact of broadcast traffic on both the access link, and broadcast server PE must be considered. Not immediately apparent is the fact that broadcast traffic will tend to affect predominantly the Inbound to LAN side of the network link; whereas non-broadcast traffic will tend to be more balanced, depending on application characteristics.

The fixed limits for Proxy table sizes must also be monitored, and considered in the evaluation. This is not a hard limit since exhaustion of the limit will simply result in a larger amount of broadcast traffic.

The MAC Header Compression table sizes must be specified via service data. The LanCalc tool will keep track of the number of different frame types to provide an estimate of the appropriate values for the two parameters involved.

In common with most engineering situations, the scalability of a solution is based on the most limiting constraint (i.e. **weakest link**).

Note that memory consumption is not normally a constraint, however it may be if a complex multi-protocol MAS is part of the virtual ring. In that case the DPN-100 MET and/or SNA MET tools should be used to provide an evaluation of memory requirements.

Given the complexity of the LAN environment, each network and application scenario is going to be unique. It is very difficult to generalize in this situation.

## Network Link Utilization

- **Function of network size, number of devices, application protocols, frame sizes, and ISRB options**
- **Separate criteria for broadcast traffic and specifically routed traffic**
- **Protocols such as TCP and SPX burst mode use all available bandwidth**
- **Slowest speed network link normally the limiting constraint**

**Need a methodology to establish maximum allowable broadcast traffic**

16

Network link utilization in the ISRB scenario is a complex function of the size of the virtual ring, the number of devices, the different protocols and applications, frame sizes and ISRB options.

Broadcast traffic and specifically routed (i.e. non-broadcast) traffic will have different criteria.

Since the predominant application protocols TCP and Novell Burst mode SPX dynamically adjust their window sizes; application traffic will tend to drive network link utilization to 100%. Consequently, it is important to use mechanisms that place bounds on the broadcast traffic, allowing a certain guaranteed minimum for the productive non-broadcast traffic.

As with a chain, where the weakest link determines the ultimate strength, so in the ISRB, the lowest speed network link normally determines the limiting link speed constraint. Remember however that the overall network constraint may reside elsewhere (i.e. Broadcast Server PE utilization).



## Broadcast Traffic Limitation

- **Constrain broadcast traffic to lowest priority of DPN 4 Q mechanism**
- **Determine percentage share of network link for each priority level**
  - function of average frame size for each level
  - evaluate statistics to determine average frame sizes
  - use LanCalc tool to determine percentage share
- **Determine percentage to allow for broadcast**
- **Consider other features sharing Broadcast Server**
  - T2.1 Router; ITI, or X.25 Broadcast

17

The DPN 4 Q quota mechanism provides a minimum guarantee to each of the four traffic priority levels as shown in the following table:

- |   |           |
|---|-----------|
| 1. High Priority Delay Sensitive        | 3 packets |
| 2. High Priority Throughput Sensitive   | 2 packets |
| 3. Normal Priority Delay Sensitive      | 2 packets |
| 4. Normal Priority Throughput Sensitive | 1 packet  |

Given this packet-based priority scheduling mechanism, it is necessary to know the average frame sizes to determine the **minimum guaranteed** percentage share of network link bandwidth. These average frame size values can be determined by processing DPN-100 statistics.

The LanCalc tool makes it easy to calculate the percentage share if the average frame sizes are known. An excerpt from the dialog box is shown on the next slide to illustrate the calculations.

Typically this lowest priority level may not be used exclusively for broadcast traffic, but may also be used for low priority application traffic. Users will often allocate a portion of this amount to serve as an engineering constraint.

Remember that there may be other features such as the T2.1 router that are also using a broadcast server. The load generated by these applications must also be factored into the equation.

## Broadcast Traffic Limit Example

**SNA traffic**

**ISRB traffic**

**Broadcast traffic**

**Define Average Frame Sizes**

Select Broadcast Prio

1 High, Delay

2 High, Thruput

3 Normal, Delay

4 Normal, Thruput

Specify Util % Value

Frmsz	Util %
128	20.7
256	27.6
256	27.6
450	24.2
	10.0

Enter average frame sizes for each of the four DPN priority levels; These will be used to calculate the minimum guaranteed share of link bandwidth based on the DPN quota priority mechanism.

- Each network will be different
- Exploit DPN prioritization mechanism

18

The LanCalc tool will simplify determination of the limiting broadcast traffic utilization threshold. This becomes one of the more significant potential constraints.

The basic approach is to determine the minimum guaranteed bandwidth share based on the DPN 4 queue quota scheduling algorithm. This algorithm allocates bandwidth on a packet basis, of 3,2,2,1 for each of the four DPN priority levels. Consequently the actual share of bandwidth is a function of the average frame sizes on the trunk or network link in question. Note that the LanCalc tool can help assess the average frame sizes for the broadcast traffic sources. In the case where broadcast traffic is relegated exclusively to the lowest priority queue, the average frame size may be less than the next higher priority queue which may have all file transfer traffic allocated to it.

Relegating broadcast traffic to the lowest priority is an excellent way to prevent broadcast traffic from having a serious impact on your network operation. This is a particularly important consideration in the event of broadcast storms which may ensue if a particular LAN station goes berserk, and generates a large number of broadcast frames.

Depending on the overall network design, you may wish to assign a higher priority to broadcast traffic; particularly if you are using broadcast as a means to disseminate useful information through the network. Applications such as Network News sometimes use broadcast protocols in this manner.

As always, it is difficult to generalize, and each network is unique hence the rationale for using a tool such as LanCalc and exploiting the unique traffic prioritization capabilities of Magellan DPN-100.

In a typical mixed SNA interactive and ISRB network users may want to follow this example for mapping VCs to the DPN Priority mechanism to ensure that high priority interactive traffic is not swamped by LAN traffic

## PE Type and Broadcast Server

- **Use of Broadcast Server on a MAS is a double edged sword**
  - netlink utilization will be reduced
  - broadcast server must process all broadcast frames, and may consume too much of PE utilization, becoming the weakest link
- **Only applies if there are multiple ISRB capable PIs on the MAS, or downstream MAS-based nodes in a cascaded configuration**

**LanCalc tool makes it easy to examine this tradeoff**

19

It is important to note that deployment of the broadcast server function on a local MAS in the case where there are multiple ISRB capable PIs may or may not be a good idea when the network is large. Generally use of the local broadcast server in this situation is recommended since broadcast traffic on the network link can be substantially reduced.

The cost of this savings is offset, however by a PE utilization impact of processing all broadcast frames on a PE that is also supporting ISRB and possibly other access traffic as well as netlink traffic. This multi-protocol access environment is normally the lowest throughput situation and may become the weakest link.

The LanCalc tool can be easily used to look at the tradeoffs in this area.

## Agenda

---

- ISRB Overview
- Engineering Constraints
- **Broadcast Traffic Sources**
- Network Topology Considerations
- LanCalc Tool
- Interworking with Routers

20

This section looks at some of the key broadcast traffic sources.

## Novell Netware Broadcast Sources

- **SAP (Service Advertising Protocol)**
  - broadcast every 60 seconds by every server
  - frame size a function of number of services
- **RIP (Routing Information Protocol)**
  - broadcast every 60 seconds by every IPX router
  - frame size a function of topology of network
- **Service query**
  - broadcast by every Novell user station to locate nearest servers
  - use Sniffer to determine relevant frequency

21

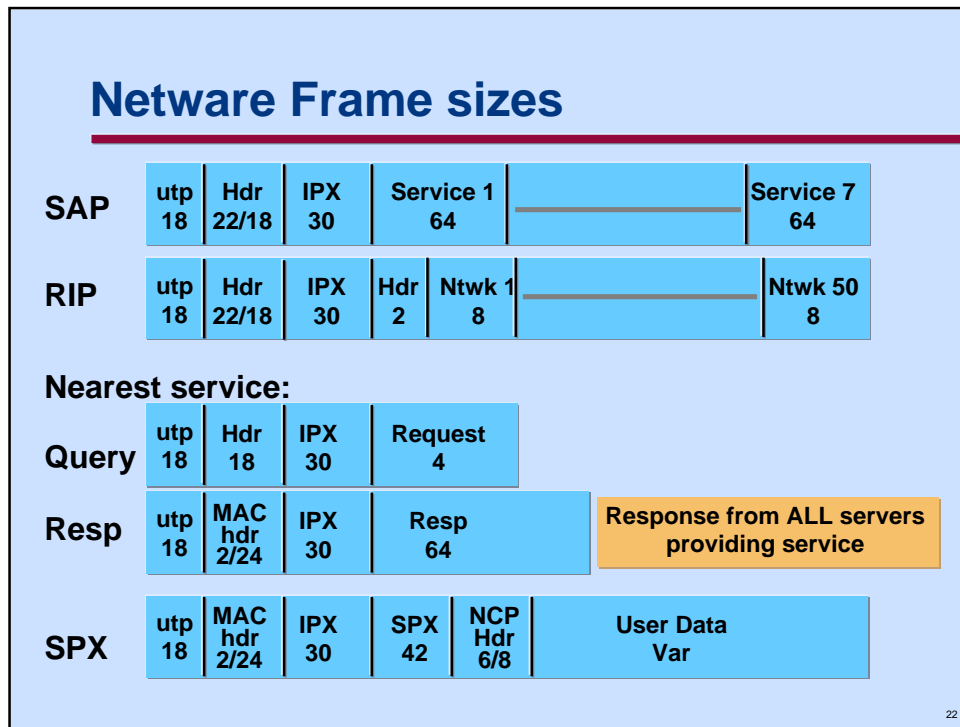
In the Novell Netware environment, there are three key sources of broadcast traffic.

The Service Advertising Protocol, or SAP is used by Novell servers to advertise the services available from the server. These frames are broadcast every 60 seconds, and their frame size is a function of the number of services being advertised.

The RIP frames are broadcast by Novell routers every 60 seconds by each IPX router in the network. The size of these frames is a function of the topology of the network.

The Service query frame is used by individual Novell end-user stations to locate their nearest server. This type of broadcast is not a periodic broadcast like SAP and RIP, consequently it is necessary to do a little analysis on a real network to determine the frequency; a LAN Sniffer can be used to trace this type of activity; alternatively, assume it is related to the frequency with which server requests are issued by end stations, or the frequency with which end-user stations are connected to the network.

- Note: the initial version of the LanCalc tool does not provide specific support for the Name Query protocol. A generic broadcast capability will be added later to address specific cases such as this. A reasonable approximation of impact may be obtained by using the Netware Interactive protocol with appropriate user data sizes.



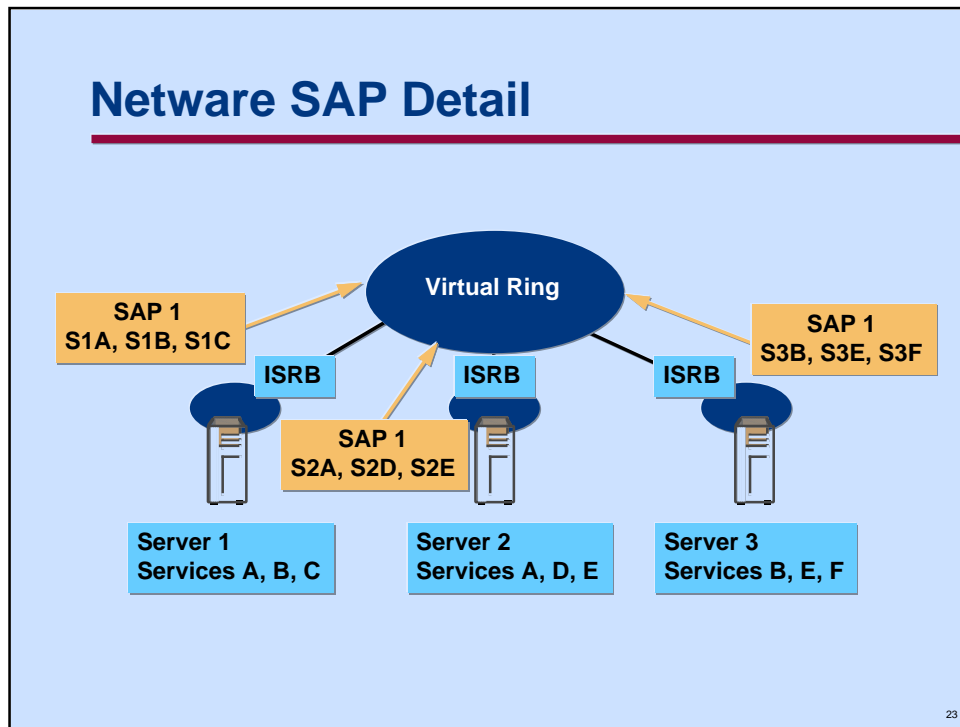
This is a high-level overview of the main Novell frame types involved in broadcast traffic as seen on a network link including the utp header of 18 bytes (12bytes VC, 6bytes HDLC) and the media header or broadcast header as appropriate. For broadcast frames the 22 byte Proxy header applies, unless the frame is not proxied (either not supported as in Nearest server query, or tables overflowed), in which case an 18 byte MAC header applies. The MAC header of 24 bytes uncompressed or 2 bytes compressed applies to specifically routed frames.

The SAP frame is the general service broadcast issued every 60 seconds by Novell servers. Up to 7 services can be advertised by a server in each frame. If there are more services to advertise, then additional broadcast frames will be constructed by the server.

The RIP broadcast frame is used by routers to advertise their knowledge of the topology of the network every 60 seconds. Up to 50 networks (i.e. LANs) can be advertised in a single RIP packet. As with the SAP packet additional RIP frames may be generated if there are more than 50 routers in the network.

The Nearest service query is used by individual LAN stations to locate servers in the network that provide services desired. The **query is broadcast**, and **each server offering the type of service requested will respond**. This query is typically used at station startup time, but may be used at other times depending on application design. There are approximately 20 different categories of servers that could be requested by a station, see Novell documentation for more information. Use of a Sniffer on an operating Novell network is perhaps the most pragmatic way to get information on the number of requests and frequency of requests for this type of broadcast. Note that the **response packet is not broadcast**, only the query packet.

**User data packets** in the Novell environment normally use the **SPX** protocol to provide reliable transport. SPX is similar in many respects to the TCP protocol. In addition to the **42 byte SPX header** is a 6/8 byte request/response header known as the Netware Core Protocol (**NCP**) header.

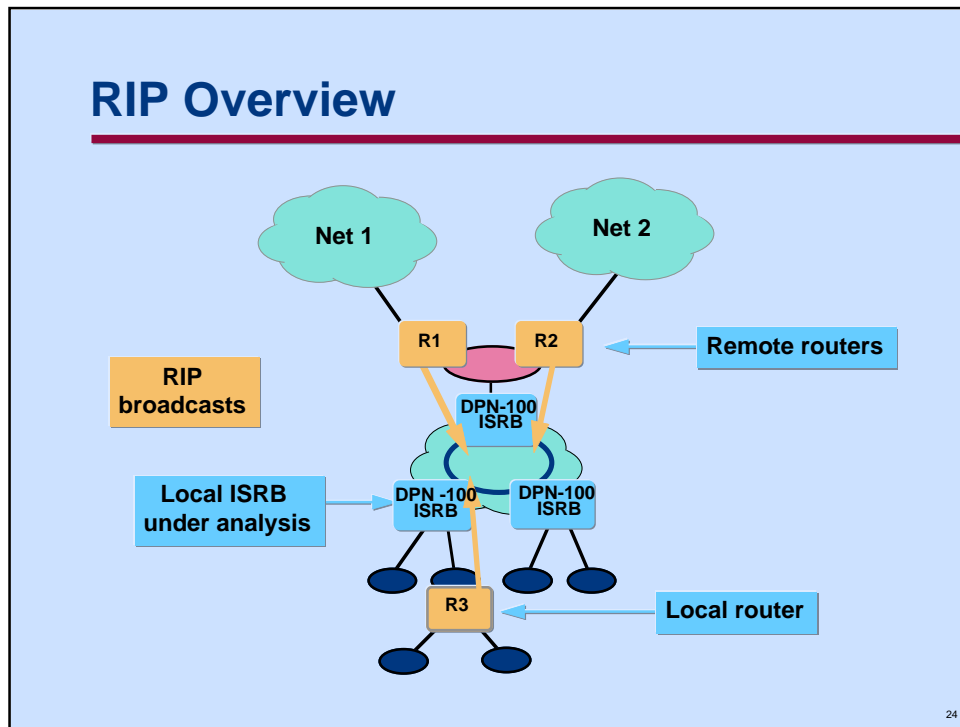


This slide provides an overview of the operation of Novell SAP servers on a bridged LAN segment. Each server broadcasts the services that it supports every 60 seconds.

A detailed look at Novell SAP broadcasts, as shown in Laura Chappell's book, *Novell Netware Lan Analysis*, indicates that after the initial broadcast at server startup where each server broadcasts its name and the services; each of the other servers in the network immediately learns about all the other servers and services. All subsequent broadcasts from each server, in that example will contain entries for each of the services on each of the servers. This is obviously an  $n*(n-1)*s$  problem which if not carefully considered and controlled can easily swamp a bridged Novell token ring environment. This example in the Novell book can be misleading, and **does not apply** to the ISRB situation. The reason for this is that in the Novell example the server is situated as a gateway between independent LAN segments. In that case, it must broadcast not only its own services, but also those of servers on the other LAN segment.

Users are encouraged to aggressively employ ISRB Filtering capabilities, to keep local server SAP traffic from being unnecessarily broadcast through the network.

Given that a SAP frame can only advertise seven services, then multiple frames will be generated by each server at each 60 second interval if the number of services per server exceeds seven.



This diagram provides a high-level overview of the RIP process which is common to most dynamic routing protocols such as Novell Netware, TCP/IP, AppleTalk, etc.

Note that the broadcast traffic is a function of the number of routers and the number of network addresses being summarized by each router. In the case depicted by router R3, this is simply the number of LAN segments directly attached to the router. In the cases depicted by routers R1 and R2, the number of segments summarized depends on the size of the networks represented by the clouds Net1 and Net2. Remember that R1 and R2 may employ some filtering and/or summarization capabilities to reduce the number of net entries that need to be broadcast. The local and remote terminology employed here is the same as used by the LAN traffic protocol selection dialog box of the LanCalc tool.



## IP Broadcast Sources

- **RIP (Routing Information Protocol)**
  - broadcast by every router every 30 seconds
  - frame size a function of network topology
- **IP ARP (Address Resolution Protocol)**
  - broadcast by individual stations to determine MAC address associated with a given IP addr
  - frequency a function of appl'n characteristics
  - use Sniffer to determine frequency
- **IP RARP (Reverse Address Resolution Protocol)**
  - used by diskless workstations to locate their server
  - potential source of broadcast storms

25

The TCP/IP protocol suite which is rapidly becoming the default for client server applications does not typically generate vast quantities of broadcast traffic, unlike Novell Netware or Netbios.

### **RIP:**

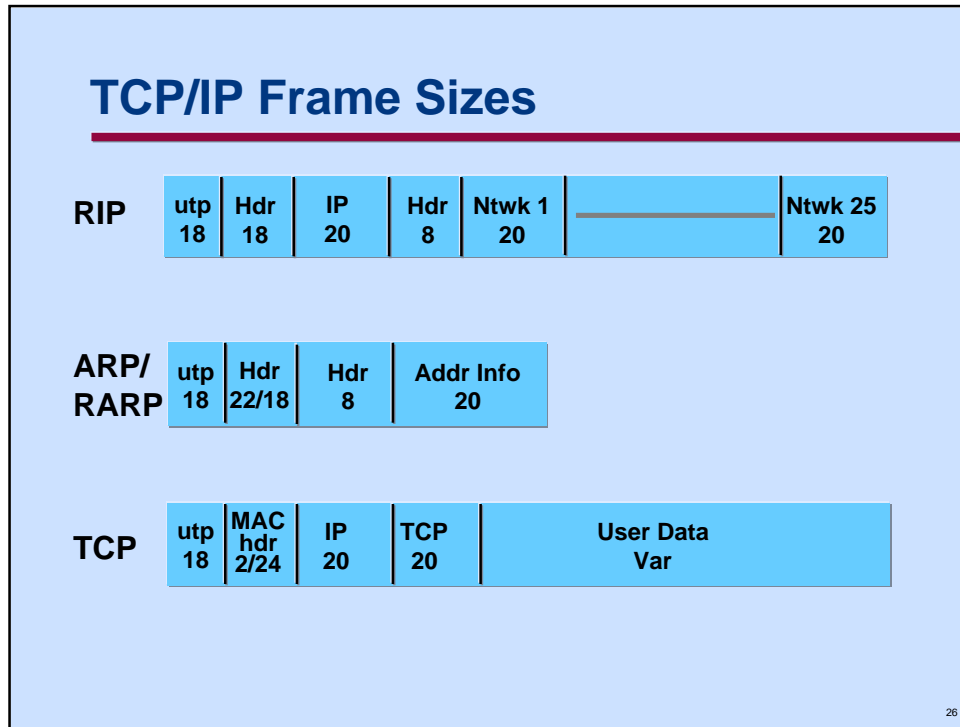
RIP is broadcast every 30 seconds by every router in the network to provide network topology and routing information. The frame size is a function of the topology of the router network. In most cases an ISRB virtual ring will not see broadcasts unless two or more routers are connected to the ring and to a backbone router network.

### **ARP:**

ARP is used by end-user stations and routers to determine the mapping between IP and hardware addresses. The frequency is largely a function of application characteristics such as the rate of session establishments, cache sizes, and cache aging timer values. Consequently, it is normally necessary to use a sniffer on an operating LAN segment to determine this value. The default value for the cache timer is 20 minutes for an entry for an active station, and three minutes for an entry for an inactive station. Assuming that each station will ARP once every 20 minutes for each remote destination is a reasonable engineering assumption if no more detail is available.

### **RARP:**

RARP is an inverse form of ARP used by diskless workstations to locate their server so that they can be downloaded with an operating system and application code. Diskless stations may be a potential source of broadcast storms if no server is available since poor implementations may retry indefinitely. In general it is best to avoid use of RARP in a bridged network if at all possible.

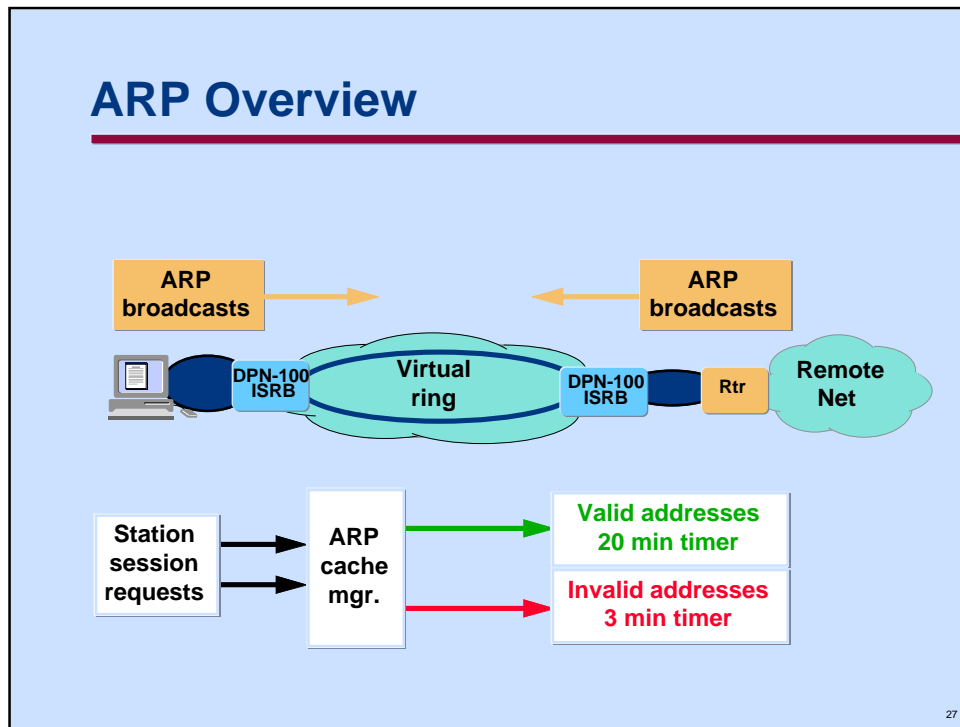


The utp header of 18 bytes (12bytes VC, 6bytes HDLC) as seen on a network link and the media header or broadcast header as appropriate is common for all frames. For broadcast frames the 22 byte Proxy header applies, unless the frame is not proxied (either not supported as in RIP, or tables overflowed), in which case an 18 byte MAC header applies. The MAC header of 24 bytes uncompressed or 2 bytes compressed applies to specifically routed frames.

The RIP frames are generated every 30 seconds by IP routers. While an ISRB bridged environment may not include any routers in the ideal case, in the general case there are likely to be routers that are connected to either external networks or possibly to a high-speed router backbone such as Passport. In that case the ISRB virtual ring will see RIP broadcast packets which are a function of the topology of the IP subnetworks connected to the ring. Note that RIP packets can only hold 25 entries in a maximum size 512 byte packet. Consequently, in a large network, multiple RIP packets may be generated every 30 seconds by each router.

The individual stations on a ring using the TCP/IP protocol suite use the Address Resolution Protocol (ARP) to dynamically locate stations and associate their IP address with a MAC address. The frequency is largely a function of application characteristics such as the rate of session establishments, cache sizes, and cache aging timer values. Consequently, it is normally necessary to use a sniffer on an operating LAN segment to determine this value. The default value for the cache timer is 20 minutes for an entry for an active station, and three minutes for an entry for an inactive station. Assuming that each station will ARP once every 20 minutes for each remote destination is a reasonable engineering assumption if no more detail is available.

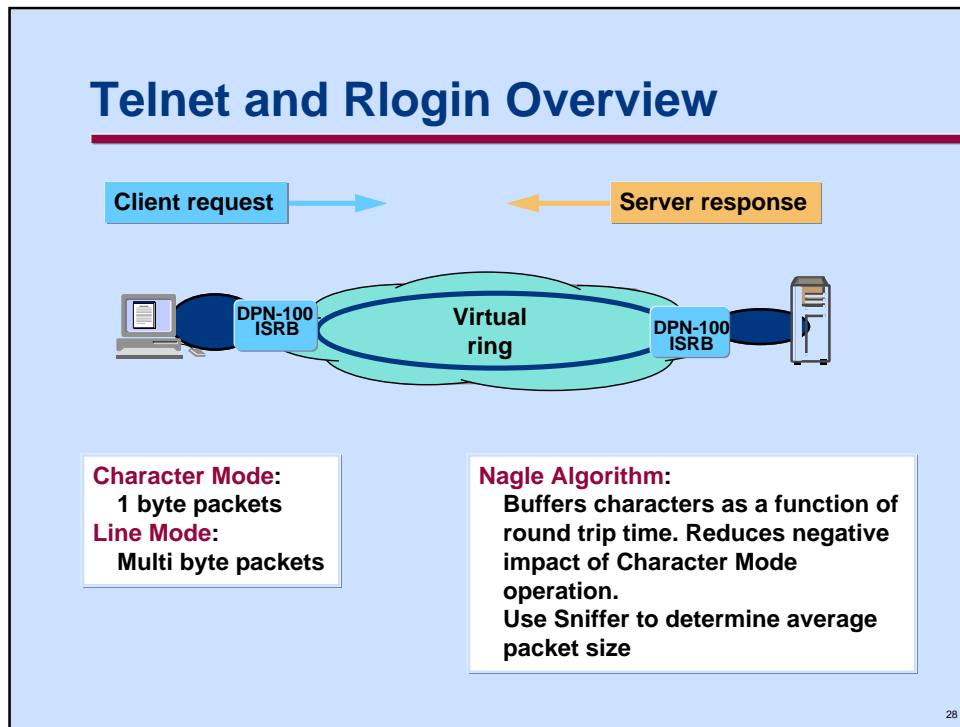
User data packets are normally TCP packets to provide reliable transmission. The TCP protocol will segment user data streams into MTU (Maximum Transmission Unit) sized segments. In addition TCP uses an adaptive windowing algorithm that responds to network congestion and attempts to keep WAN link utilization as high as possible. For more information on this dynamic windowing mechanism refer to the *Predicting End-User Performance* presentation given at Inform 95.



This diagram provides a high-level overview of the ARP process which is common to many dynamic routing protocols such as TCP/IP, AppleTalk, etc. Note that Novell Netware does not use the ARP process since the SAP process provides most of the same functionality.

Note that the broadcast traffic is a function of the number of individual LAN stations, and the average number of sessions established by each station. It is also a function of the number of routers attached to the virtual ring and the number of sessions established from the remote net(s).

Note that each station and router normally employ an ARP Cache. The cache is intended to reduce the frequency of ARP broadcasts and operates with two different aging timers. Session requests are aged in the cache depending on the response to the initial ARP broadcast. If a response was received from an ARP, then the *valid* or *completed* address is aged using a 20 minute timer. If however the ARP times-out (default 75 second TCP timer”, then the address is marked as *invalid* or *incomplete* and these entries are timed out at a three minute frequency. Note that the LanCalc tool does not automatically differentiate between these two timers, allowing the user the flexibility to specify the frequency of the timer for each row entry.



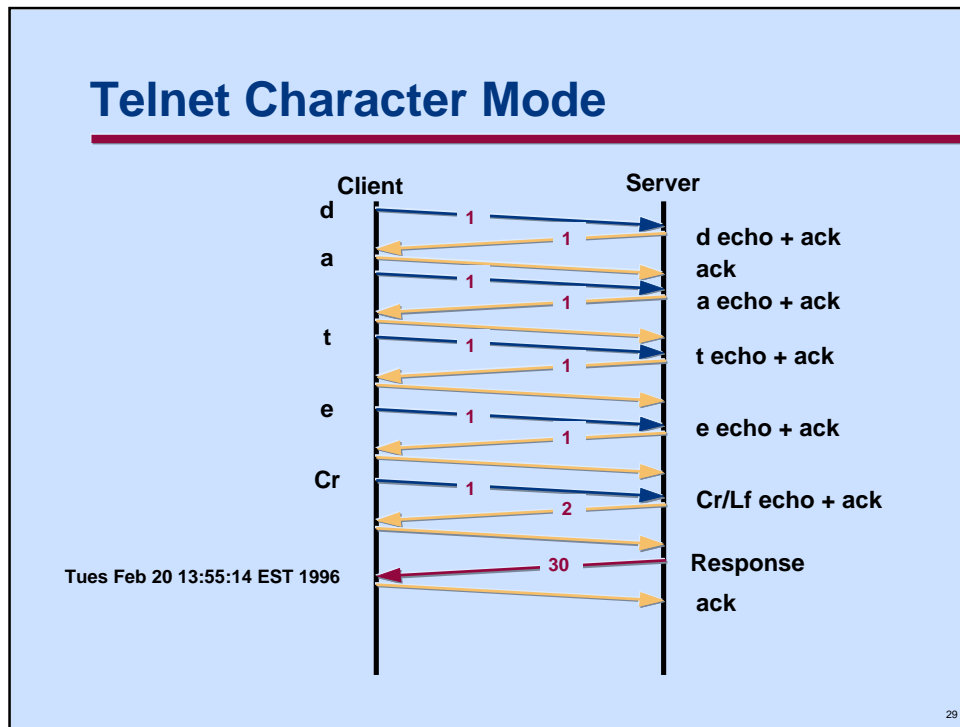
Telnet and Rlogin are two of the commonest LAN applications. They operate in a classical client server relationship. The end-user at a client terminal typically issues short commands and the server responds with a typically larger response. For example the user may request a directory listing requiring 10 keystrokes, and may receive several hundred characters in response. A study by Verne Paxson in 1993 on typical Internet traffic found that the **ratio** of response to request size is **20:1**.

There are many options for Telnet, but for the purposes of a scalability engineering analysis there are two significant modes of operation, character mode, and line mode.

In *Line Mode*, a packet is generated for each line of input or output. The line length is variable, but typically the number of characters which can be displayed in a window (72-80).

In *Character Mode*, a separate packet is generated for each character typed by the client. The server will echo each of these input characters with separate packets. This mode of operation can generate very large amounts of traffic given the 20 byte TCP and 20 byte IP headers plus the media headers. On a high-speed LAN this is not an issue, but on a lower-speed WAN it can become a very significant issue.

The *Nagle* algorithm was developed to deal with the *tinygram* problem introduced by Character Mode operation. This algorithm says that a TCP connection can have only one outstanding small segment that has not yet been acknowledged. TCP will collect these small segments of data and send them in a single larger segment when the acknowledgment for the previous segment arrives. The beauty of this arrangement is that it is self-clocking, and adjusts to the delay in the network. In effect, this means that both client and server echo data packets will contain the number of characters that can be typed in in one round trip time. This is the engineering approximation that is used in the LanCalc tool. Use of a Sniffer to determine the average packet size for client requests and server responses is recommended.



This diagram shows a typical simple Telnet or Rlogin client server interaction showing the number of packets flowing and the amount of user data included in each packet.

The user in this example types in **d a t e** followed by carriage return. Each of these characters is sent as a TCP data packet containing just one byte of data.

The server echoes each of these packets and piggy-backs a TCP acknowledgment in the one byte data packet.


The client then sends a TCP acknowledgment packet with zero data bytes.

This pattern repeats for each character until the terminating carriage return, which the server echoes with a two byte carriage return/line feed response.

The server then responds with the 30 character date which is acknowledged by the client to complete the interaction. Note that while we are operating in character mode, the Nagle algorithm results in the server response being sent as if it was operating in line mode.

This general sequence is the one modelled by the Character Mode option of the LanCalc tool.

Note also that the Nagle algorithm could result in 2 or 3 characters from the client being grouped together in a single packet based on the round trip transit delay and the speed of typing.



This chart, borrowed from the Inform '95 presentation *Predicting End-User Performance* shows the cumulative data rate as a function of the cumulative number of bytes transferred. This is perhaps one of the best charts to appreciate the effective throughput using the TCP and SPX protocols. Note how much slower the Novell algorithm reaches a steady state throughput compared to the TCP algorithm. In particular, the rate for small files of less than 30K bytes drops off precipitously. The exact shape of this curve depends greatly on the overall round trip delay, and the maximum advertised window size.

The Y axis in this graph represents the average data rate in Kbits/sec. The X axis is the cumulative number of Kbytes transferred.

For the purposes of the initial version of the LanCalc tool, a simplifying assumption is made that either protocol will drive the access netlink to saturation at 95% utilization. This 95% is approximately the utilization level that congestion management and packet discard starts to occur, which triggers the dynamic backoff algorithms of TCP and SPX Burst Mode.

Consequently, use of either TCP File transfer or Novell Burst Mode traffic types should be the last entry in the LanCalc input area. The tool will assume full mtu sized frames, and calculate a frame/sec rate which saturates the balance of the link. This is done to assess PE utilization in this worst case scenario to determine if it is a limiting factor.

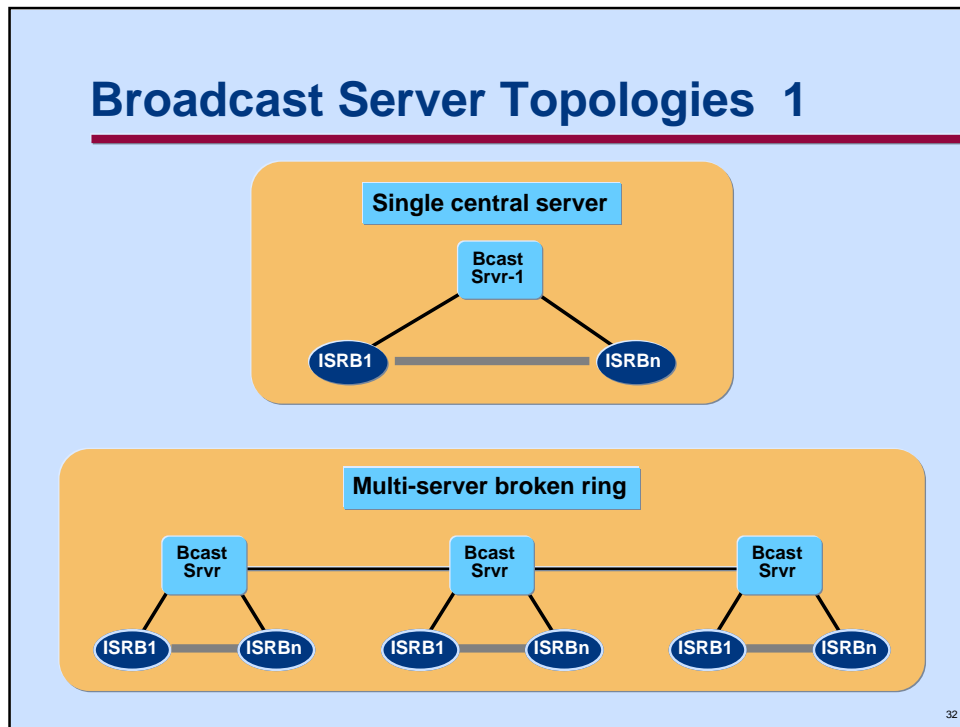
## Agenda

---

- ISRB Overview
- Engineering Constraints
- Broadcast Traffic Sources
- **Network Topology Considerations**
- LanCalc Tool
- Interworking with Routers

31

This section looks at some topology design issues associated with placement of broadcast servers.



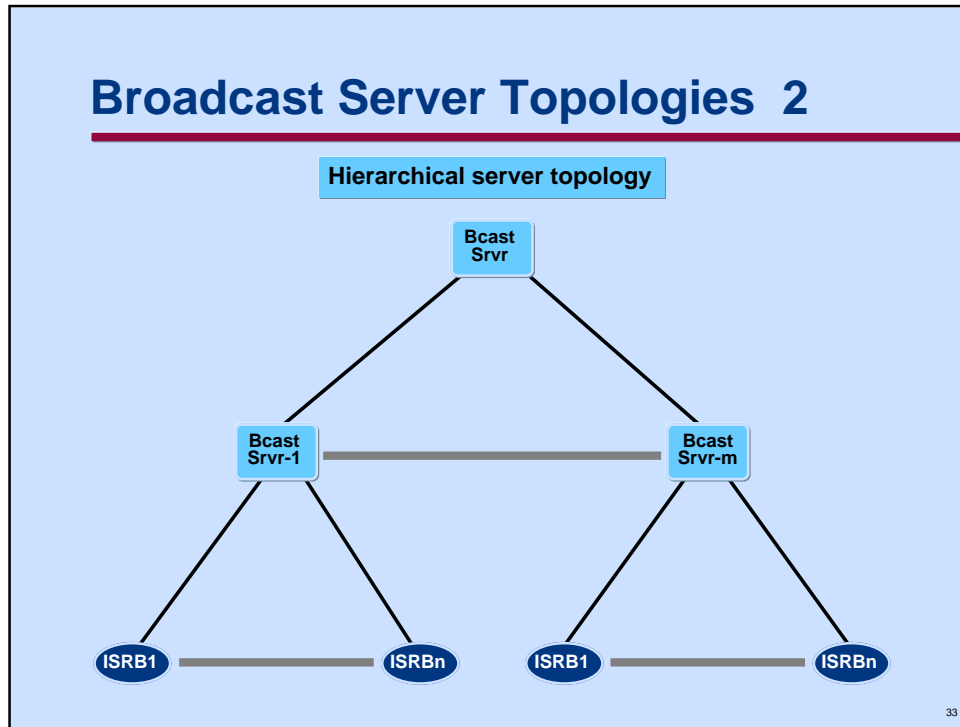
Two typical starting logical topologies are shown in this diagram.

The simplest configuration is the single central server which is suitable for simple topologies. The LanCalc tool can be used to assess the maximum number of VCs that can be supported.

The multi-server broken ring topology can be used to extend beyond the single central server concept without going to a hierarchical structure. This approach keeps the number of broadcast servers to a minimum, but needs to be analyzed closely to ensure that broadcast traffic mapped onto the physical topology does not introduce trunk loading problems.

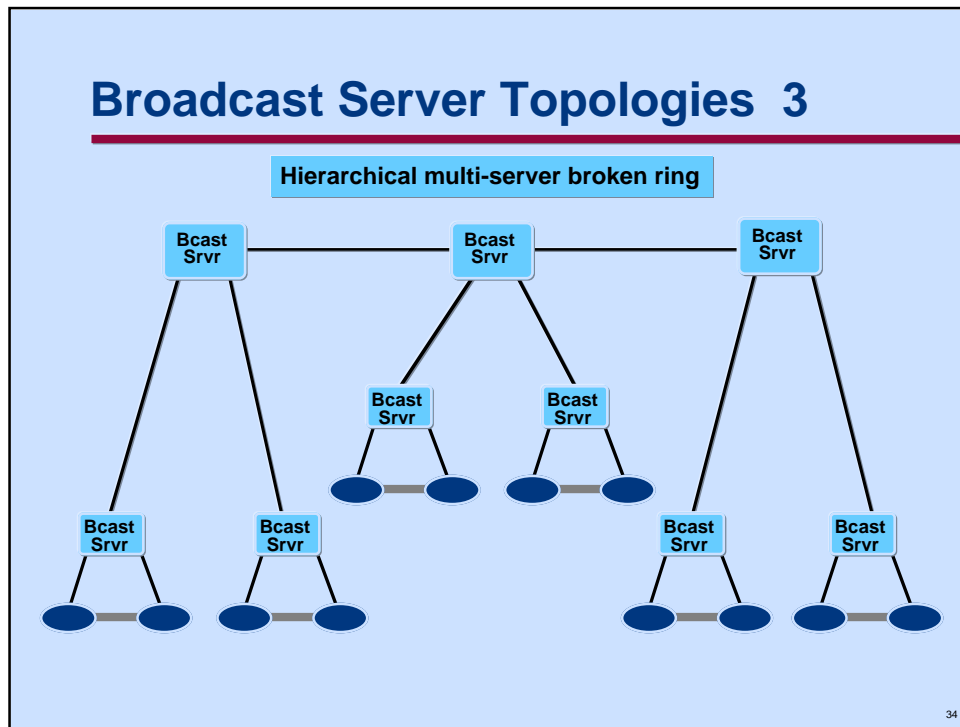
The outputs from the LanCalc program can be used as input to topologically-aware network planning tools such as NetCalc to determine how any given ISRB traffic maps to a specific physical network topology. Details on how to do this are outside the scope of this presentation.





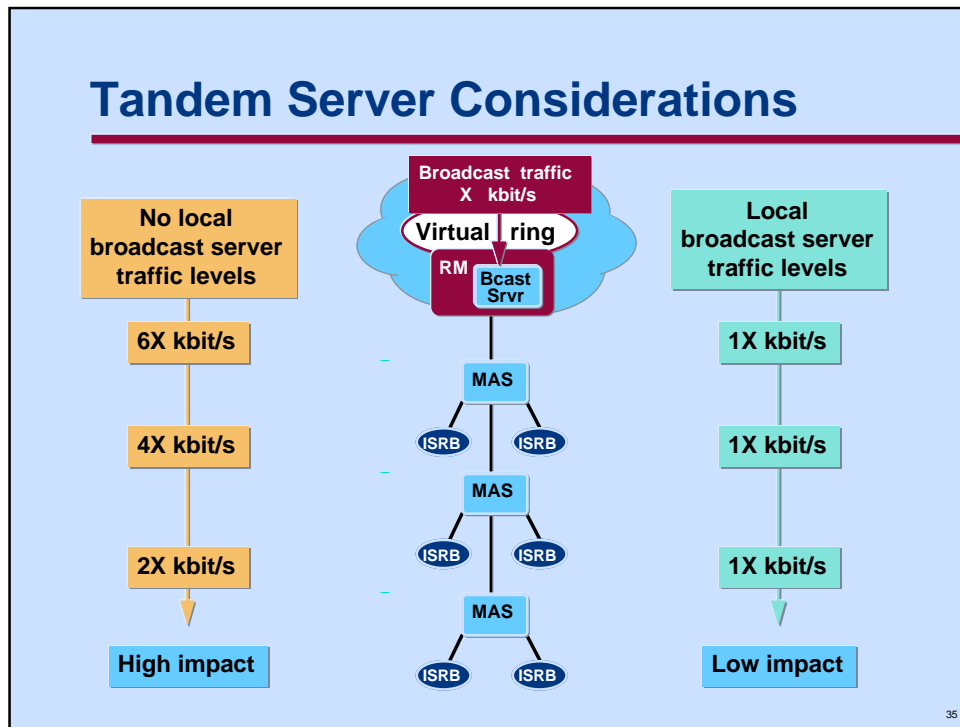
This slide depicts a simple two level hierarchy as was used for the earlier examples of broadcast traffic flow.

This hierarchical approach may be continued to several levels as necessary to handle broadcast traffic, but it is seldom necessary to go beyond this two-level example.



The two earlier topologies can be combined to create a hybrid broken ring at the top level of the hierarchy, and a hierarchy below it. This provides a savings of a single broadcast server over simply continuing the hierarchy to a higher level.

This approach may be a better solution than a three level hierarchy when the traffic is mapped to a specific physical topology. As before, a combination of LanCalc info and NetCalc topology analysis will provide a quantitative assessment of the benefits of one design alternative over another.



This slide depicts how broadcast frames would flow in a tandem ISRB environment.

Examining only the broadcast traffic from the balance of the network, the most significant component, for the sake of simplicity; we can see that broadcast traffic is progressively additive on each subsequent network link depending on the number of ISRB LAN segments attached to each MAS.

For the purposes of illustration we have assumed that there are **two ISRBs** per MAS. Given that the total broadcast traffic is **X Kbit/s** being broadcast downwards from the RM located broadcast server; we can see **6X**, **4X**, and **2X** broadcast traffic on each of the network links proceeding downwards from the top of the chain in the case where local broadcast servers are not employed.

Where local broadcast servers are employed, only **1X** Kbit/s of broadcast traffic is seen on each of the tandem network links. This is a significant savings.

Analyzing the two scenarios for PE utilization indicates that the topmost MAS is the most heavily loaded in both cases, although there will be minor differences in the local server case as tandem trunk traffic has relatively lesser PE utilization.

The bottom line is that local broadcast servers are considered mandatory in most tandem ISRB scenarios.

## Agenda

---

- ISRB Overview
- Engineering Constraints
- Broadcast Traffic Sources
- Network Topology Considerations
- **LanCalc Tool**
- Interworking with Routers

36

This section provides an overview of the *LanCalc* Tool. Users are encouraged to attend one of the *Whiteboard Engineering Clinics* for a hands-on test drive of the tool. A limited supply of free copies of the tool for either MAC or PC environments will also be available at the clinic.

## LanCalc Overview

- **Wingz spreadsheet-based tool available for PC, MAC, and Unix platforms**
- **Dialog boxes for all data entry**
- **Context sensitive help**
- **Table driven to make extensions easy and permit key parameters to be changed**
- **Detailed error log for parameters and inputs out of range**
- **All logic coded in Hyperscript to prevent inadvertent change and speed execution**

**Wingz available from IISC worldwide**

37

The LanCalc tool is a free Nortel-developed Wingz spreadsheet based tool that incorporates data entry dialog boxes and context sensitive help facilities to simplify the task of quantitatively assessing the impact of various ISRB configuration options, traffic types, PE types, and link speeds.

Data entry and primary results are contained in a single page of the sheet. Additional pages are used to provide a more detailed breakdown of input parameters for each traffic type, and the intermediate calculation of PE and link loading for both broadcast, and non-broadcast traffic both inbound to the LAN and outbound from the LAN.

An error log provides detailed information on input errors.

The internal design of the tool is table driven in many areas to simplify extension of the tool to cover additional LAN traffic types, and different PE types, etc.

All the logic associated with the various calculations is written in the Hyperscript language. This provides protection from inadvertent change, improves quality and maintainability, and greatly speeds execution.

Wingz is the standard engineering tools platform used for internally developed tools from the Network engineering group. It is available internationally for PC, MAC, or Unix platforms from IISC at the contacts listed below. Note the current Wingz is a standardized product, and users of the older international PC version from Informix should upgrade to the new version to fix bugs.

North and South America:

- Doug Fielder: Tel: 1-800-494-9464 fielder@wingz.com

Europe and Asia Pacific:

- Keith Andrews: Tel: 44-171-628-6960 keitha@wingz.com

## LanCalc Main Page

Global parameter dialog boxes (3)

Global Access Variables				Global Broadcast Variables				Link Limits		Frame	Util	Kbps
12-Feb-96	PE Type	MAS	HPPS: SEL	UTP	PE Type	MAS	20s	High, Delay 1	100	17.6%	11.3	
	Link Speed Kbps	64.0			Max # Serv VCs	10		High, Thrupt 2	200	23.5%	15.0	
	Subnet Pktsz	512			Proxy Keep Alive	1min		Normal, Delay 3	300	35.3%	22.6	
	Avg # LWVCs	4			MTU/Frame Size	512		Normal, Thrupt 4	400	23.5%	15.0	
								Limiting Broadcast LMTS	5.0%	3.2		

Transaction Description	Protocol	Options		Kbps Traffic Summary				Cumulative Frame Counts		PE Utilization %		Link Util %		Errs			
		B	MHC	Proxy	In	Out	In	Out	In	Out	Total	per VC	Access Server		Ass # Server		
Case1	NovellSAP	N	N		0.0	0.0	10.4	0.1	0	0	0	0	0.0%	16.0%	16.3%	0.1%	1 err
	NovellSAP	N	Y		0.0	0.0	1.2	0.0	1	200	0	0	0.2%	12.2%	12.4%	1.0%	0.0%
	NovellRIP	N	Y		0.0	0.0	0.6	0.0	1	100	0	0	0.1%	6.1%	6.2%	0.0%	0.0%

Input parameters via dialog box

Summary output results area

This slide provides an overview of the layout of the main data page format for the LanCalc tool.

The upper portion of the page contains three separate global parameter areas. From left to right these are:

- The access PE-related global variables such as access PE type, access link speed, subnet packet size, and average number of Light Weight VCs (LWVCs) connecting an ISRB LAN to other ISRB LANs.
- The broadcast server PE-related global variables such as PE type, Maximum number of VCs to other servers and/or ISRBs, the Proxy keep alive interval, and the Maximum Transfer Unit (MTU) used for segmentation by attached LAN stations and routers.
- The global link limits area which is used to determine the minimum guaranteed bandwidth for each of the four DPN priority levels, and the user selected limit for broadcast traffic.

The body of the main sheet contains the row organized key input variables and the key summary results. The first four columns are used to define the LAN traffic protocols selected. These selections are made by a dialog box invoked by hidden buttons located in columns 2 - 4. The values entered for the various parameters are shown on the next page.


The summary results area contains cumulative values for both specifically routed and broadcast traffic, counts of Proxy and MAC Header compressed frames, and broadcast server and access PE utilizations as well as total link utilization.

Context sensitive help is provided for each input and output field on the sheet. Any values that exceed limits are highlighted in red and a more detailed description logged in the error log.

Note a blank row resets the accumulators allowing easy comparison of alternatives as shown here.

## LanCalc Input Parameters Page

Input Parameter Names & Values											
Trans'n ID	Protocol ID	Parm 1		Parm 2		Parm 3		Parm 4		Parm 5	
		Name	Val	Name	Val	Name	Val	Name	Val	Name	Val
<b>Case 1</b>											
	Novell SAP	# Local Servers	1	# Local Services	6	# Rmt Servers	100	# Rmt Services	10		
	Novell SAP	# Local Servers	1	# Local Services	6	# Rmt Servers	100	# Rmt Services	10		
	Novell RIP	# Remote Rtrs	100	RmtNetsPerRtr	10	# Local Rtrs	1	LocNetsPerRtr	10		
	Novell RIP	# Remote Rtrs	100	RmtNetsPerRtr	10	# Local Rtrs	1	LocNetsPerRtr	10		
	IP RIP	# Remote Rtrs	10	RmtNetsPerRtr	100	# Local Rtrs	1	LocNetsPerRtr	25		
	IP ARP	Local Stations	10	Remote Stations	100	Sessions/Station	10	Frequency in min	1		
	TCP Interactive	Local Stations	10	Sessions/Station	1	Msg Size Outbound	100	Msg Size Inbound	100	Interval in min	1



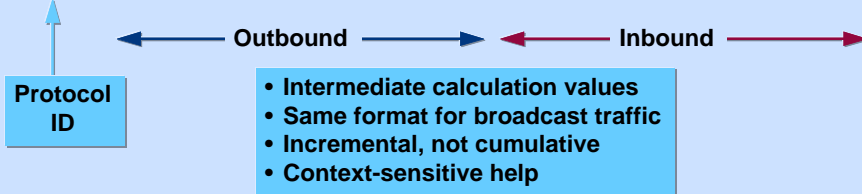
39

This slide shows an extract of the Input parameters page. This section is the second page on the sheet, and stores the values that have been entered via dialog box for each protocol selected.

While the primary use of this area is to save and document these parameters and values; users may change parameter values directly before recalculating results. Note however that any change in the parameter name fields will result in errors at recalculation time.

## LanCalc Intermediate Results

Protocol ID	Specific (Non-Broadcast) Traffic Intermediate Calculations															
	Outbound from LAN							Inbound to LAN								
	frames/sec	pkts/sec	Kbytes/sec	frames/Sec	pkts/Sec	Link Util	PE Util	frames/sec	pkts/sec	Kbytes/sec	frames/Sec	pkts/Sec	Link Util	PE Util		
Microsoft	0.02	0.02	0.007	456	494	0.00	0.0%	0.0%	3.55	3.55	1.240	572	590	0.00	0.0%	0.6%
Microsoft	0.02	0.02	0.000	25	45	0.00	0.0%	0.0%	3.55	3.55	0.055	25	46	0.00	0.0%	0.2%
Microsoft	0.02	0.02	0.000	25	45	0.00	0.0%	0.0%	1.67	1.67	0.042	25	45	0.00	0.0%	0.1%
Microsoft	0.02	0.02	0.002	152	190	0.00	0.0%	0.0%	1.67	1.67	0.220	152	190	0.00	0.0%	0.2%
IP ARP	0.07	0.07	0.019	278	296	0.00	0.0%	0.0%	1.67	1.67	0.715	428	446	0.00	0.0%	0.3%
IP ARP	1.67	1.67	0.047	25	46	0.00	0.0%	0.1%	16.55	16.55	0.915	52	70	0.25	1.0%	1.2%



40

This slide shows an extract of the Intermediate Calculations page. There are two identically formatted pages on the sheet:

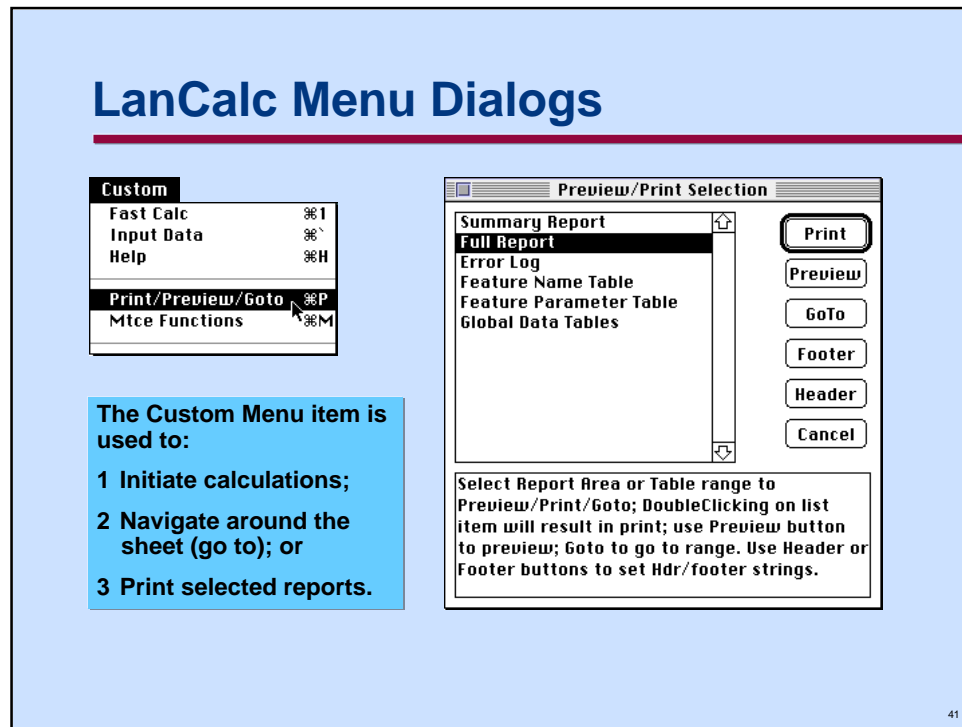
- The first page as shown above, (page 3) provides the intermediate calculations for the specifically routed traffic (i.e. non broadcast) for the protocol selected on any row.
- The second page (page 4) provides the intermediate calculations for the broadcast traffic for the protocol selected on any row.

The information is grouped into two large, identically formatted sections. The first section is for traffic outbound from the local LAN, and the following is for traffic inbound to the local LAN. The values on each row are incremental, not cumulative from prior rows. As with all input or output cells on the sheet, context sensitive help is available.

For each of these segments the following intermediate calculations are provided:

- The first three columns contain values used to calculate PE utilization. Note that the Kbyte/s field is based on the user data; whereas the later Kbit/s field is based on the total data.
- The two frame size fields provide the user data and total frame sizes. User data includes all LAN headers, as applicable, as transmitted on the network link. The Total frame size includes the utp link headers as well.
- Network link traffic in Kbit/s, and link utilization are also provided.
- The PE utilization; either access PE for specifically routed traffic, or broadcast server PE utilization as appropriate to the page is also shown. Note that for broadcast protocols, there are values on both pages; whereas for specifically routed traffic, there will be values only on the first page.





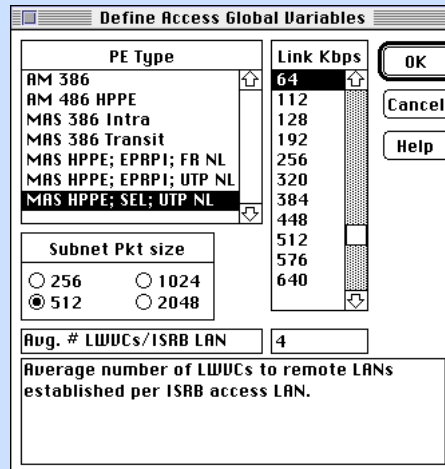
The Custom Menu item in the menu bar is used to select custom functions specific to the LanCalc tool. The options are:

- Perform a fast calculation, **command 1**. This results in all non-blank rows of data being evaluated and all calculations performed by the internal tool logic. No cell-based formulae are used by the tool, ensuring reliable consistent results and fast calculation.
- The input data selection, **command `**, is used to reposition to the first page of the tool.
- The Help selection, **command H**, can be used to obtain context-sensitive help for any input or output cell on the sheet.
- The Print/Preview/Goto selection, **command P**, is used to pop up the dialog box shown at the right.
- The Mtce Functions selection, **command M**, is normally greyed, and is used by the tool developer to ensure internal parameter table integrity.

The Print Preview dialog box allows a user to print, preview, or go to selected ranges on the sheet.

- The Summary Report provides a two page summary, Results, and Inputs.
- The Full Report provides all four pages including the intermediate calculations.
- The Error Log contains detailed error text for any errors detected in any cell during Fast Calculation.
- The last three entries are used to peruse the tables that are used by the tool for its internal operation. Most users can ignore these areas.

## LanCalc Global Access Dialog Box

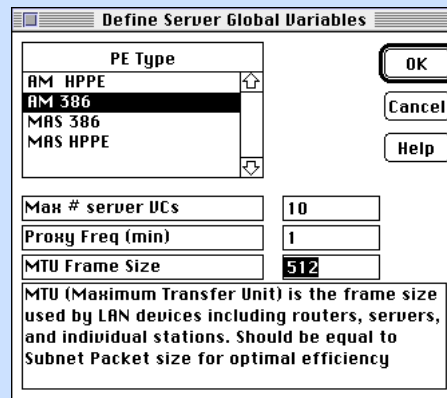


42

The access global dialog box is used to define access-related global parameters. The key parameters defined are:

- The access PE type selected from a list box. This is used for access PE utilization calculations. Note that the tool is hard coded in its initial release to flag PE utilizations exceeding 80% as errors.
- The access link speed selected from a list box. This is used for access link utilization calculations. Note that the tool is hard coded in its initial release to flag link utilizations exceeding 80% as errors.
- The subnet packet size is selected from a radio button box. This is intended for calculating PE segmentation; however this capability is not included in the first release of the tool.
- The average number of LWVCs (Light Weight Virtual Circuits) established to other ISRBs is used to determine the number of entries in MAC Header Compression tables.

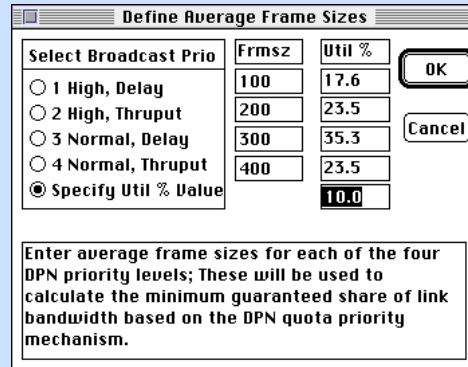
## LanCalc Global Server Dialog Box



The server global dialog box is used to define broadcast server-related global parameters. The key parameters defined are:

- The broadcast server PE type selected from a list box. This is used for PE utilization calculations. Note that the tool is hard coded in its initial release to flag PE utilizations exceeding 80% as errors.
- The maximum number of server VCs. This is used to calculate the total broadcast traffic load on the server PE.
- The Proxy server frequency is used to calculate the traffic for Proxy synchronization frames. The default value is one minute.
- The MTU (Maximum Transfer Unit) value is used for segmentation calculations for LAN originated traffic. The initial version of the tool performs segmentation analysis where appropriate, but does not perform packetization. As noted in the help field, the MTU value should be the same as the subnet packet size for optimal performance, or alternatively a simple multiple.

## LanCalc Global Link Dialog Box



Select Broadcast Prio	Frmsz	Util %
<input type="radio"/> 1 High, Delay	100	17.6
<input type="radio"/> 2 High, Thruput	200	23.5
<input type="radio"/> 3 Normal, Delay	300	35.3
<input type="radio"/> 4 Normal, Thruput	400	23.5
<input checked="" type="radio"/> Specify Util % Value		10.0

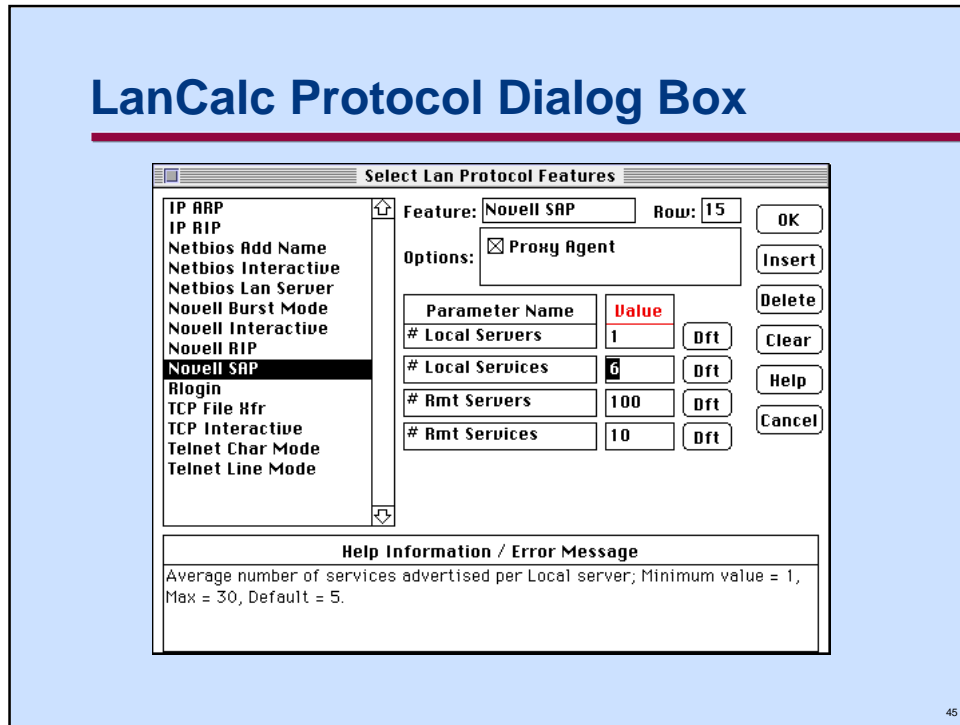
Enter average frame sizes for each of the four DPN priority levels; These will be used to calculate the minimum guaranteed share of link bandwidth based on the DPN quota priority mechanism.

44

The Link global dialog box is used to determine the minimum guaranteed bandwidth share of the access link for each of the four DPN priority levels. The key parameters defined are:

- The average frame size is specified for each of the four priority levels in the fields underneath the *Frmsz* heading. The percentage share of the access link is dynamically calculated as values are entered. DPN statistics may be used to obtain good averages for an existing network or weighted averages may be developed using intermediate calculations from *LanCalc* and other tools such as *SNA Delay*.
- Selecting a radio button from the left hand box will place the associated utilization value in the fifth utilization box as an input field value. Alternatively selecting the last choice *Specify Util % Value* allows the user to specify some other value directly in the box. Note that on exiting the dialog box, this value will only be transferred to the global portion of the sheet if the *Specify* button was selected.

On exiting the box, not only are the utilization values placed in the global area, but also equivalent Kbit/s limits based on the access link speed entered in the access global area.



This dialog box is the major workhorse for the LanCalc tool. It is invoked via hidden buttons located in columns 2 - 4 in each row of the main sheet. The selection options are as follows:

- The list box on the left is used to select one of the LAN protocols supported by the tool. Selection of one of these entries results in context-sensitive help being displayed in the help box at the bottom of the dialog box. Simultaneously, the input parameter fields applicable to the protocol chosen are dynamically drawn in the area to the right of the list box. At the same time the Options box is drawn to allow check box selection of either the Proxy Agent or MAC Header Compression options.
- For each of the displayed parameters there is an input field and an associated Default button (**Dft**). Selecting the input field will result in context sensitive help being displayed including acceptable range and default values. Selecting the **Dft** button will place the default value in the input field. Values entered in any of the input fields are range checked each time a new field is selected. If the value is out of range it is highlighted in red and an explanatory error message shown.

Exiting the dialog box by selecting the OK button results in range checks for each input field. If any errors exist they are highlighted and the box remains open until all errors are corrected. This problem is most likely to occur if one of the limits defined in the parameter tables is inappropriate. Changes to these tables are easy to make if necessary.

The **Insert** button will Insert a row with the same information as the existing row.

The **Delete** button will delete a row and shuffle subsequent rows up.

The **Clear** button will clear all information in the selected row.

The **Help** button displays general dialog box navigation help.

## Agenda

---

- ISRB Overview
- Engineering Constraints
- Broadcast Traffic Sources
- Network Topology Considerations
- LanCalc Tool
- **Interworking with Routers**

46

This section looks at the interconnection to router based backbone functionality as a means of expanding the scalability of ISRB solutions.

## Interworking with Routers

- **Familiar maxim: “*Bridge when you must, route when you can*”, still applies**
- **ISRB extends the application of bridging**
- **Magellan native LAN Routing can be exploited to extend overall solution**
  - Passport interLAN switching
    - IP
    - Novell IPX
  - DPN-100 embedded router
    - IP, IPX, XNS, Decnet, AppleTalk, etc.
- **External token ring gateway**
  - can be stub ring
  - consider hop count constraints

47

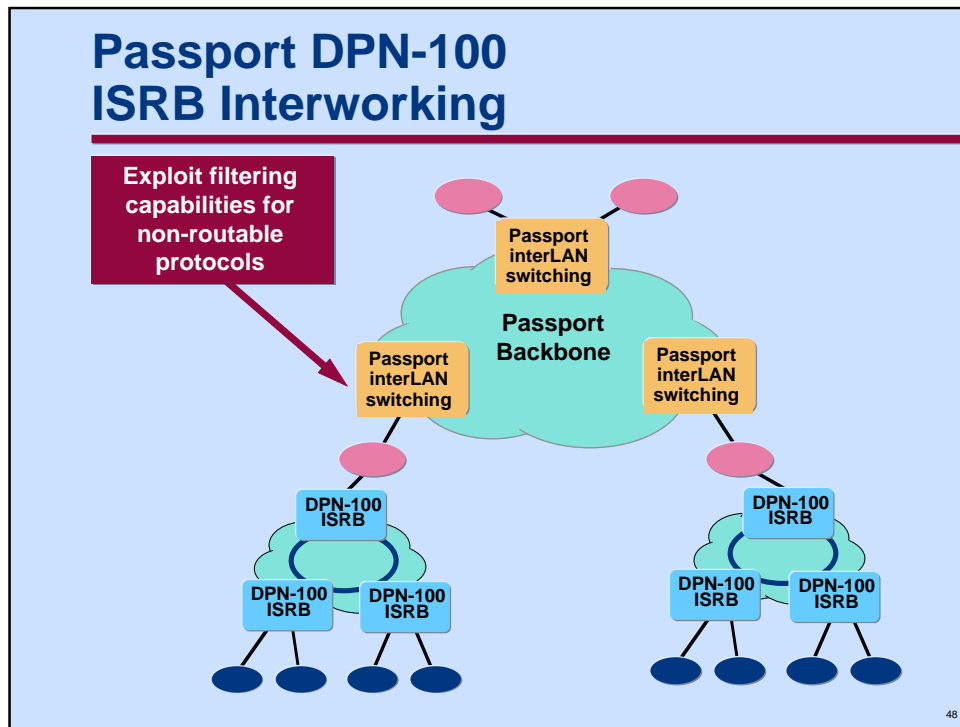
While the DPN-100 ISRB solution extends the applicability of simple LAN bridging techniques over conventional bridges, it is still preferable to use routing wherever possible.

The use of native Magellan-based LAN routing capabilities provides a natural evolutionary path for ISRB-based LAN solutions.

Native Passport routing function currently exists for the two leading LAN protocol stacks; TCP/IP, and Novell Netware SPX/IPX. Additional protocol support will be added later as market demand dictates.

The Magellan DPN-100 Embedded Router provides LAN routing support for all popular routable LAN protocols.

An external token ring gateway; typically a stub ring can be used to connect virtual rings to a backbone router ring as is explored further in the next slide. When using this technique, it is important to consider possible hop count constraints; particularly if there are downstream bridges connected to the access level ISRB LANs.



This diagram provides a high-level overview of a collection of DPN-100 logical access virtual rings or subnets connected via Passport interLAN switching.

Multiple ISRB based access virtual ring subnets can be constructed from a DPN-100 based access layer surrounding a Passport backbone. This is a logical migration step for Magellan DPN-100 networks requiring the higher performance and extended LAN functionality offered by Passport.

The access layer virtual rings can be engineered to most cost effectively exploit existing investments in Magellan DPN-100 technology.

The additional hop count introduced by this solution is not normally a problem, but may be relevant where there are multi-hop, externally bridged token rings connected to the ISRB access nodes.

Planned upcoming enhancements to the DPN-100 ISRB service will eliminate the need for external token ring LAN attachments to connect regional virtual rings to a Passport based LAN routing backbone. This enhancement will provide RFC1490 based frame relay encapsulation as the link between the ISRB and Bifrost router, allowing existing trunk ports and links to be shared.



## Summary

- **Knowledge is the key**
  - LAN protocol suites
  - application characteristics
- **Use LanCalc to quantitatively assess alternative network designs**
- **Exploit native Magellan Passport routing**
  - upward migration path
  - reduces impact of broadcast storms
- **References:**
  - TCP/IP Illustrated                      Stevens
  - Netware LAN Analysis                Chappell
  - LAN Protocol Handbook            Miller
  - ISRB Specification and Guide      Magellan

49

As noted earlier, there is no substitute for knowledge of the LAN protocol stacks, and the specific key applications.

The LanCalc tool provides a user with the ability to quantitatively assess the impact of alternative proposed network designs quickly, easily, and effectively.

Magellan native Passport routing capability can be exploited to provide an upwards growth path for ISRB-based virtual rings. This also provides a means of fire-walling broadcast storms, even though the effective use of DPN-100's unique priority system limits the worst-case impact of these storms.

Finally, there are some references that I have found to be particularly valuable to the Network Engineer.

- *TCP/IP Illustrated*, Volume I, by W. Richard Stevens, ISBN 0-201-63346-9 is an absolutely excellent reference text for this key protocol suite; highly recommended.
- *Netware LAN Analysis* by Laura Chappell, published by Novell Press, ISBN 0-7821-1362-1 provides an excellent overview of the Novell Netware protocol suite; also highly recommended.
- *LAN Protocol Handbook* by Mark A. Miller, ISBN 1-55851-099-0 is a useful reference for a number of protocols, including Netbios. This book tends to focus on the MAC layer.
- Finally, the ISRB Specification, and User Guides are required reading for ISRB specific information.