SWITCHING TO ATM

TR Switching - System Solutions

A Guide to Supporting IP and IPX Protocols in Switched Networks

Introduction

The emergence of Token Ring switching as the LAN (Local Area Network) interconnect collapsed backbone in a building backbone delivers dramatic performance improvements in client server applications. It also simplifies LAN administration for TCP/IP (Transmission Control Program/Internet Protocol) protocols, by avoiding the moves and changes problems associated with routing in the building backbone.

Madge's Switching White Paper outlines the needs, the functionality, benefits and futures of Token Ring switching. This guide is aimed at technical personnel who require more detail about the system implications of Token Ring switching. It advises on client/server network configurations, the issues associated with supporting TCP/IP addressing and subnetting in switched environment and methods of avoiding local IPX (Internet Packet eXchange) routing in Novell client/server environments. It also clearly marks the dividing line between Switching & Routing and how the two technologies complement each other in the enterprise network.

Token Ring Switching, a new LAN architecture

Traditionally, Token Ring building/campus corporate networks have been based on a *distributed backbone* architecture, in which floor or departmental workgroup rings are source route bridged off of a common backbone ring. The familiar concept is illustrated below:

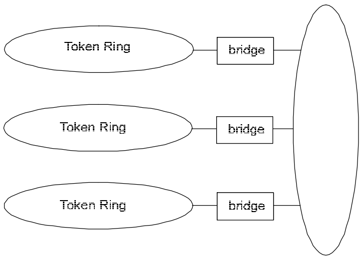

Figure 1: Traditional bridged Token Ring system architecture

This architecture is manageable and has served well while client/server applications have largely remained local to the department or physical area served by the workgroup ring. When fileservers are placed on the same ring as the workstations they serve, the performance of this model is excellent; workstation users see good response from local servers and access to central resources is not slowed by any overload on the 16 Mbits/sec backbone ring.

However, in many organizations this model of working is no longer practical. Increasingly, organizations have moved towards a centralization and consolidation of information access and storage. A significant factor in the move to a centralization of server resources is the cost and difficulty of supporting distributed file servers, but in many cases the centralization trend has been driven by a need to hold reliable, fully backed up, single copy, shared sources of data and application software. Centralized file servers are easier to manage and maintain.

When servers are centralized, the *distributed backbone* is a poor architecture; the hop across a bridge introduces significant latency and increased traffic on the 16M backbone ring results in an overload during high activity periods. Both of these factors contribute to unacceptably slow workstation response times.

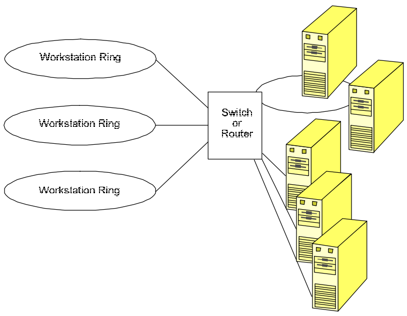

Figure 2: A collapsed backbone System Architecture

For a centralized server strategy and to prepare for other demanding applications, the collapsed backbone network architecture provides the bandwidth and flexible growth for the future.

Traditionally, the collapsed backbone has been based on a router, but Token Ring switching challenges its role in this application and is set to change the system architecture in the building and campus network.

The problems with a routed collapsed backbone

Many network administrators have collapsed their backbone with a dedicated router or have created a collapsed backbone by adding Token Ring ports to an existing router previously used solely for private or public wide area access. This architecture provides the bandwidth for centralized information servers and a future growth path, but suffers from significant drawbacks:

■ The latency associated with the store and forward operation of a router means that the point to point access from a workstation is slowed when communicating with a server located on another router port. Users notice the delay in access compared with the access times they were used to with servers placed locally on departmental or floor based rings.

■ It is too expensive to dedicate a 16Mbits/sec Token Ring router port to key server resources, it is also unlikely that a router port will have the right hardware and software upgrade to support full duplex operation with this mode of connection. The Token Ring connection into a router port is normally presented as a token ring “node”. This is incompatible with a direct connection into a server adapter (as would be required to support a full duplex operation). Building a FDDI ring to accommodate a small group of higher bandwidth server connections is an equally expensive option.

■ There is an unacceptable administration problem associated with local routing and the regular moves and changes common in most organizations. TCP/IP protocols require that either workstation IP addresses have to be changed as they move from one ring to another, or the router needs to be re configured to accept the new positioning of a TCP/IP station on a different router port. These are serious problems, since the adjustment of router configurations is complex. Furthermore, in organizations with large geographically dispersed networks, the control of router configurations is normally delegated to those individuals concerned with global network administration, whereas workstation moves and changes are a local issue.

It is primarily these three reasons that have driven a move away from layer 3 routing to layer 2 bridging or *switching* in the local collapsed backbone, but the cost of routing is also a significant factor: the price of a switched port is significantly lower than that of a router.

The new solution with local switching

Token Ring switching is the new technology for the collapsed backbone. It is targeted at providing an equally high bandwidth collapsed backbone but without the latency and administration issues associated with a router solution. At the same time, switching is also a lower cost solution on a per port basis, which offers the opportunity to run smaller ring sizes, dedicated server connections and dual redundancy as affordable options. Compared with the local routed backbone the switch is:

- Higher performance
  Cut-through switching between rings avoids the point-to-point performance loss associated with a router. Workstation access times to a server sited across a switch are virtually unchanged from those obtained from a locally sited server.

- More flexible
  Switch ports are low enough cost to assign dedicated ports for key server resources. Switch ports can be configured to accept a direct connection to a server adapter and can be software upgraded for 32 Mbits/sec full duplex operation (simultaneous 16 Mbits/sec transmission and reception).

  Dual redundant switches are affordable, easy to configure and a compelling resilience option.

- Easy on moves and changes
  A switch is self learning during moves and changes. It is fundamentally a “layer 2” device, building a flat network architecture. Access to network resources is unchanged wherever a workstation is sited within the local network.

- The migration path for multimedia and ATM
  Forthcoming multimedia applications require a reserved bandwidth and guaranteed delivery timing between the source and destinations of a multimedia session. The non blocking switch fabric in Ringswitch clears ring-to-ring paths and the cut-though operation of Ringswitch reduces network jitter. Ringswitch is also upgradeable for connection into an ATM backbone.

The benefits of taking a LAN switching approach are therefore clear. At the same time, LAN switching forces a new consideration of system issues. These are centred around routeable protocols and how they are handled in a locally switched network. The way that these protocols are supported in a switched LAN potentially differ from those taken with a router approach. LAN switching encourages simplification of LAN administration, which is an ideal approach. These approaches are discussed fully in this document.

Furthermore, in order to ensure that the switching solution scales effectively to serve larger building and campus LANs, the backbone switch needs to offer a means of avoiding a rising level broadcast traffic. The methods used to control broadcast are also discussed in this paper.

The System Issues

Token Ring switches are *complementary* to a routing strategy, not a replacement. The role of routing is in the LAN/WAN boundary and within the geographically dispersed network. Broadcast control issues, limited wide area link bandwidth, dealing with multiple paths through a geographically dispersed network and the need for a hierarchical approach to wide area networking currently favours routing and protocol conversion in the wide area connection. The positioning of the two technologies is summarized in the diagram below:

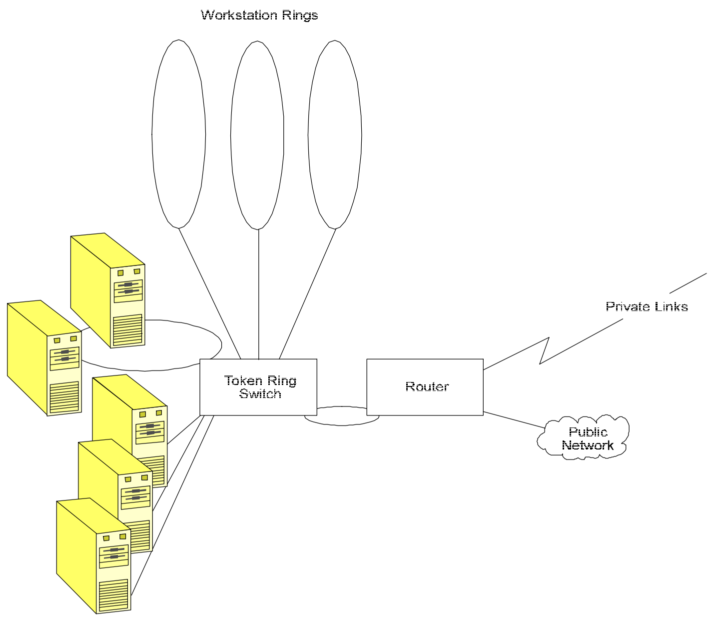

Figure 3: The router's role is at the LAN/WAN boundary in a building/campus network

The above system architecture has a Token Ring switch as a collapsed backbone, with key servers running dedicated connections into the switch. Wide area connectivity is via a single or, in some cases, dual/ multiple Token Ring connection into the router.

Some LAN administrators are concerned about the widespread use of LAN switching in the collapsed backbone. Concerns are centred around routeable protocols (TCP/IP and IPX) in the building/campus LAN and on the control of broadcast traffic. The role of the router in the local/campus LAN has been one of providing broadcast control, a means of segmenting the LAN for firewall or security purposes and a means of communication between routed domains (IP subnets, IPX network numbers). This section discusses the broadcast control, firewalls and security and the support of routeable protocols from a switching perspective. Madge Network's Ringswitch address all of these issues and can be considered as the preferred performance solution for the collapsed backbone.

Broadcast control

Broadcast control is an issue in the LAN because of the disturbance broadcast can cause to attached PCs and workstations on the LAN. Every broadcast frame reaches host software, stealing a small number of CPU cycles each time. High levels of broadcast can cause stations to noticeably slow.

Switching on its own will not check broadcast levels. To overcome this potential problem, Madge's Ringswitch receives and processes broadcast frames from all Token Ring switched ports and forwards them according to pre-set configuration rules. This process is called *active broadcast control* and is outlined in the section of broadcast control strategies detailed later in this guide.

Firewalls and security

In some cases, it is desirable to treat parts of the local network as isolated networks, using gateways or routers to control communications. Again, a pure switched backbone does not necessarily support such configuration needs, but Ringswitch provides for this by supporting the concept of virtual LANs, which can be set up and controlled through network management. Virtual LANs are described in the Madge's switching white paper.

Supporting routeable protocols (IP and IPX)

SNMP management, Internet access and UNIX system access are all based on TCP/IP. Microsoft is promising that Windows NT LAN server and Windows 95 workstations will also use TCP/ IP as the underlying network protocol. These trends indicate a sharp increase in the use of TCP/IP in Token Ring LANs, yet TCP/IP brings with it the worst of all the administration problems associated with a locally routed environment.

The IP moves and changes problem

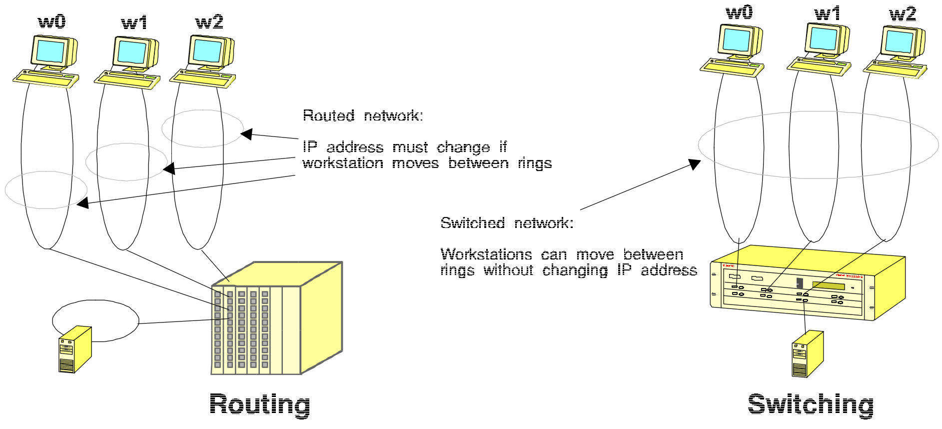

Figure 4: How IP is supported differently in routed and switched LANs

The above example highlights the problem. Routing TCP/IP means that each ring is a different subnet or IP domain. w0,w1 and w2 must have IP addresses to match the subnet or domain to which they belong. If user w0 happens to move to the same ring as w1, he can't talk IP unless his IP address is changed. Such an occurrence is a high probability in routine office moves and changes and with portable PCs gathering momentum, ring to ring moves will be very likely. Changing user IP addresses is a nightmare:

- very often IP addresses are distributed with a hierarchical meaning to them e.g. the a.b.c in IP address a.b.c.x might be "finance", a.b.d might be "sales". This would be lost if the user IP address changes to match the new ring.
- new IP address allocation and documentation needs to be very disciplined to cope with the regular moves and changes. This is an administration overhead.
- Lastly, DNS (domain name server) tables need to be reconfigured to reflect the change to a new IP address.

One strategy to cure this problem is to run several networks into each wiring closet either by running several hubs on each floor or by running a port switching hub. LAN administrators have implemented this strategy in the absence of Token Ring switch technology. It is very expensive solution, but it does enable a workstation to be patched back into the "right" ring after a move. Even so, it still involves significant patching activity or network management administration during the cycle of regular moves and changes.

An alternative and most satisfactory solution is to turn off local routing, retaining routing out to wide area connection only. This avoids the problem of moves and changes for all protocols running in the local network. Over the last two to three years, market pressure has pushed the mainstream router vendors to support just such an operation and many LAN administrators have taken up the option. It works well in mid sized and large networks. A readily apparent observation is that the router is then working in the same way as a switched collapsed backbone, but without the low cost and broadcast control scalability that a Token Ring switch would offer.

In summary, switching rather routing is the optimum solution to support routeable protocols in the building and campus LAN. At the same time switching is naturally the right way to support Non-routeable protocols such as NetBIOS and LLC2.

The System Solutions

The recommended strategy of switching Token Ring networks is a new concept. It does require planning. For many Token Ring users, large scale implementation of TCP/IP will be a new challenge and they will hear a lot about routing. Even for those that are "IP aware", a mountain of outdated books and papers on the subject will recommend routing to be implemented everywhere, including in the LAN. This section sets out the details of a switching strategy for IP and IPX and addresses uncertainties or doubts that a switching rather than routing strategy is right for now and into the future.

The concept of Token Ring switching is new, but the concept of switching is not. The degraded performance of Ethernet networks under load has already forced the deployment of switching. Token Ring bandwidth has lasted out longer. Only now has it become necessary to run switching into Token Ring too. System issues such as supporting multiple overlapping IP domains in the switched LAN or removing IPX domains altogether have been tried and tested in Ethernet switched LANs. Much of the discussion that follows on system architectures has come from experience already gained in the Ethernet environment.

Supporting TCP/IP in a switched LAN

We have already established that TCP/IP will become more widely implemented in Token Ring networks, just as it has become very widely implemented in Ethernet and public wide area networks. Such a wide adoption of TCP/IP was not envisaged when IP was standardised in the early 80's. The current IP address is simply too small and the timescale for standardising and implementing a new bigger address format world-wide is at least 2 or more years out. The squeezed address space is an issue in switched LANs. In this section, the details of IP addresses are examined and the switching system issues are noted together with solution options in support of a switching approach to the collapsed backbone.

The public IP address space and implications

Over the last two years, all of the flexible and easy to manage IP address space has been exhausted. The address space now offered for public inernet access is "class c" format, supporting up to a maximum of 254 nodes.

In order to understand the implications and solutions for a switched Token Ring system architecture, we need to look at the layout and usage of IP addresses:

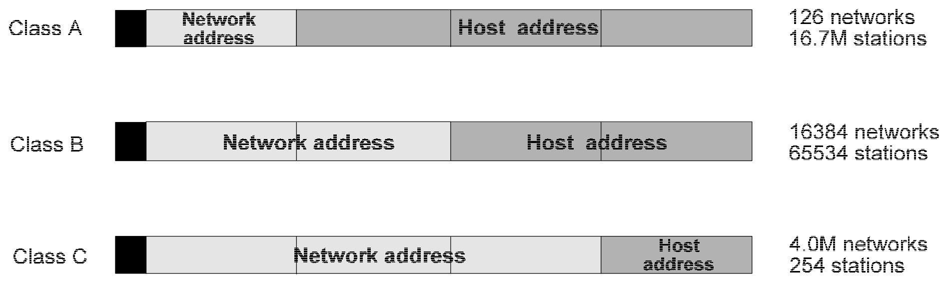

Figure 5: IP address formats

The above diagram shows the traditional IP address formats. Classic IP addresses are 4 bytes in length and split the space into 3 classes. All 3 classes have a dual content; the network address and the host address space. The top 3 bits in the IP address define which class of address the following bits belong. The network part of the address is the IP domain, the host part is the address of a specific host within that domain. In order for a host in one IP domain to reach with a host in another IP domain, the communication must go through a router. Communication within a domain is direct station to station. In the global public IP network (internet) there are many, many networks, each with a unique network address.

The IAB addressing management authority controls the world wide allocation of public IP network addresses. Given the tiny network address space of a class A format address, all public addresses were pre-allocated by the authority years ago. The situation now is that virtually all the public class B address space is used up. It is almost impossible for an organisation to obtain a public class B address from the IAB now.

This leaves class C addresses; fortunately there are still enough of these left. Class C addresses have been left to last because they are difficult to manage. With only 254 hosts per network, a large LAN needs to be served with a number of class C addresses if all workstation are to be served with their own unique address. This leaves corporate network planners concerned about switching in the building LAN. The specific concern is that multiple IP domains in the building LAN might force more routing not less. The reality is that routing and the performance/cost/administration problems that come with it, can be avoided every time. The issue is therefore not about whether to route or switch, but how to build an IP architecture that fits a switched building LAN.

There are several switching solutions for IP based networks. This guide outlines 5 solutions.

IP Solution 1: Multiple overlapping IP domains

More than one IP domain can coexist on a flat bridged or switched network. This is easy method of supporting large numbers of class C based IP stations on a single site without using routing. Wide area access and local IP server solution examples based on overlapping domains are illustrated below.

- A router attached to a switched LAN can support multiple domains off a single port.

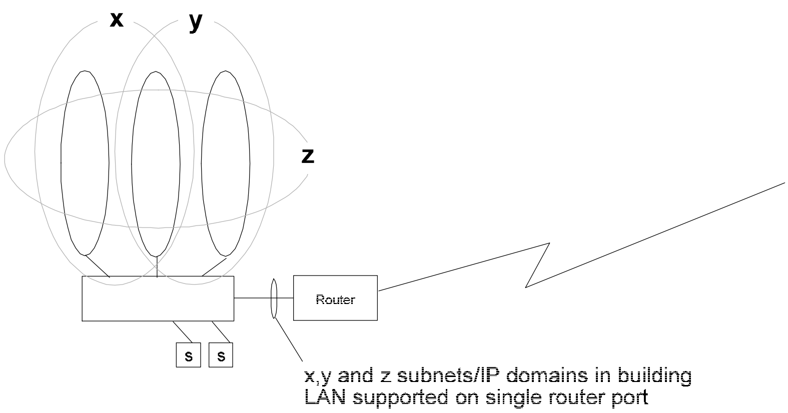

Figure 6: Supporting multiple IP domains in a switched network

Three class C addresses would support a large building network with up to 762 users. Routing is avoided within the LAN by running all three class C networks as overlapping domains in a switched network. This is a perfectly adequate solution. The LAN/WAN divide is supported by a single Token Ring router port into the switch, serving all three IP domains. In this case the Token Ring router port is running a configuration known as *secondary addressing.* Secondary addressing is fully endorsed by all router vendors. The switched LAN solution, with secondary addressing running into a router port, works well when IP access is mainly to the wide area (inter-site email, Internet access). With this architecture, access to local servers is switched with cut through low latency and the IP moves and changes problem is removed.

The switch per port cost is lower compared with a router solution. An added cost benefit is that in most situations, the WAN connection can be a small low cost access router.

- LAN based server resources based on IP can have multiple adapters

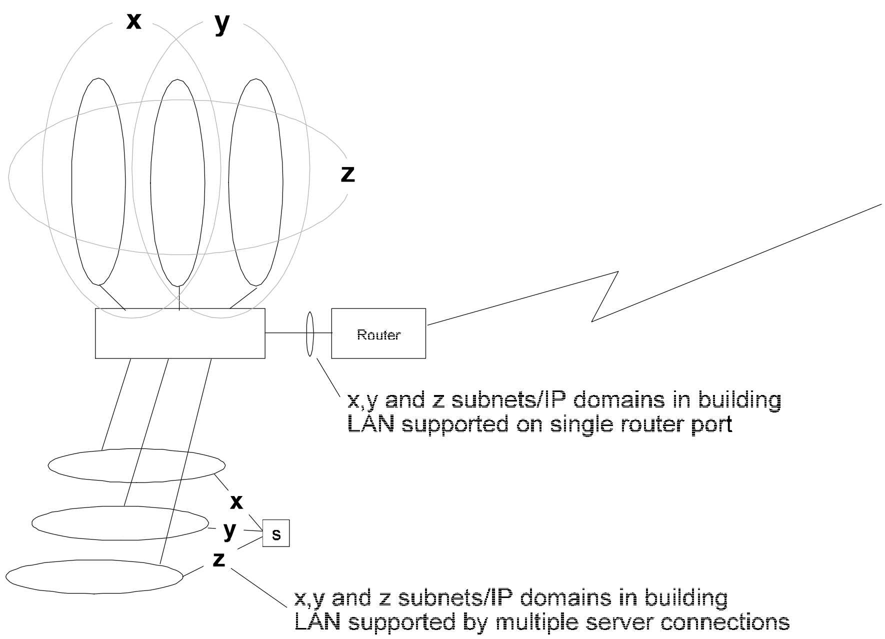

Figure 7: Supporting multiple IP subnets in a server

The overlapping domains solution can be extended to support local IP server resources (e.g. UNIX server or database, Windows NT server). Routing is avoided by supporting multiple adapters into the server resource, with each adapter serving one IP domain. When this architecture in place, all IP workstations (wherever they reside on the switched LAN) can access the server, hence the moves and changes problem is avoided.

IP server access performance is also increased in two ways; the store and forward delays which significantly slow access with a routed IP connection are avoided and the multiple paths into the server improve both bandwidth and resilience.

IP Solution 2: CIDR and "supernetting" Class C addresses

An alternative approach to supporting large number of hosts using public class C addressing has emerged over the last two years. This is a standards activity response. New IP addressing standards have been developed over the past 2 years in order to overcome the running out of address space problem. These are CIDR (Classless Inter-Domain Routing) and the concept of supernetting in PC's, servers and workstations as a software upgrade to IP drivers. Whilst at the time of writing, CIDR is not yet available on all Routers and supernet addressing upgrades are not yet present on workstations and adapters, the solution is documented here because it is a forthcoming standard well suited to a switched Token Ring strategy.

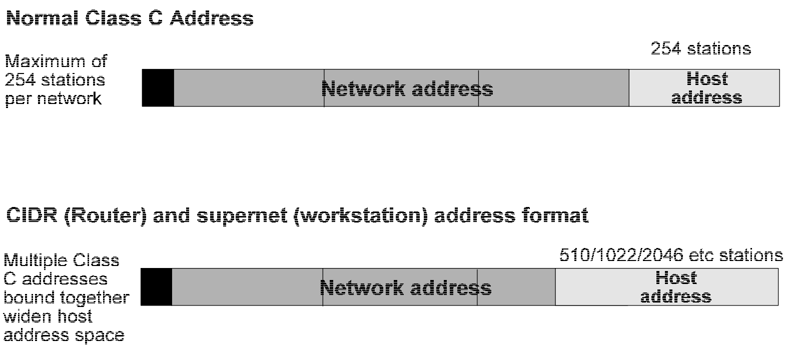

Figure 8: Supernet address format Vs standard Class C addressing

CIDR solves the class C host address limitation problem. CIDR collapses a block of contiguous class C addresses as a single new address running in the LAN and a single address carried over the routed wide area network. Class C IP addresses are now being allocated in contiguous blocks in support this standard. With this standard in place, the wide host address space to support a switched LAN can be created using contiguous class C address blocks. The address space can be wide enough to support all workstations in a large LAN as a single IP network. When the Building LAN is one single IP address, there are no routing requirements within the LAN itself and switching is an optimum solution.

The concept is illustrated below:

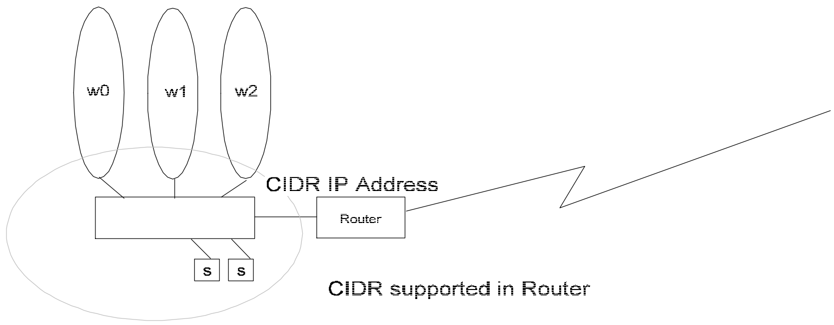

Figure 9: A single CIDR network address can support a large number of users

IP workstation software supporting supernetting references the *IP mask* configuration rather than the class bits for information about how bits in the 4 byte IP address assigned to network and host addressing. This means that even though the workstation has a class C address, it will talk beyond the class C host address boundaries (according to IP mask information) and directly access resources anywhere in the CIDR address space, obviating the need for routing within the building/campus LAN.

IP Solution 3: Using variable subnet masks

Many large organisations will have registered with the IAB for a Public IP address before class B addresses were limited. A raw class B address supports over 65,000 stations, which seems like a lot of space for a single organization. In most established IP implementations however, the class B address would be split into a number of IP networks or domains using “subnetting”. Subnetting allows a single class B address to be used by one organization to support a number of internally routed networks. Subnets are created by moving the network address bits down into the last two bytes of the IP address:

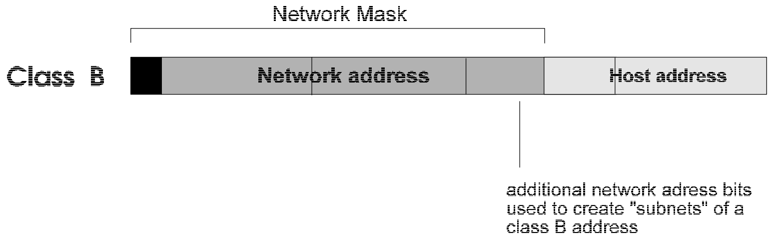

Figure 10: Class B address usage

IP consultants and router vendors have recommended that the network administrator should choose a split of network address space and host address space (a network mask) that is common to all sites supported within the routed enterprise network. This recommendation leads to inefficiency in the assignment of IP addresses. HQ and large regional offices may have large numbers of hosts needing to be supported within a single subnet, small branch offices may need routed IP connections into the main enterprise network, but inevitably have a small number of hosts at the site. Choosing a common network mask to suit both is a compromise. If a wide host address space is chosen on the grounds of simplicity in the HQ and large regional offices, large numbers of IP addresses are wasted on every small branch office introduced into the network. Therefore, given the recommended limitation of a common network mask, most administrators choose a small host address space to avoid wastage. This leads to the situation that even with class B addresses, HQ and large regional networks often need to support multiple domains. Whilst IP Solution 1 (a switched network supporting multiple, overlapping IP domains) adequately deals with this situation, router vendors are now reporting reliable operation of *variable subnet masks* in private enterprise networks.

A variable subnet scheme operates IP addresses with a large host space for HQ/ large regional LANs and small host space for remote offices in a single enterprise network. In this way an enterprise wide class B address is split out more efficiently using variable subnet sizes to suit the premises being served. It also preserves a single IP address in the building LAN, making a switched Token Ring a straightforward choice for the collapsed backbone.

IP Solution 4: Private addressing and address translation

Solutions 1 through 3 have dealt with switching system issues and PUBLIC IP addresses. Running a private IP addressing scheme releases the constraints on IP addressing; all IP address space is now at the disposal of the organization. Private addressing is yet another way to avoid multiple IP domains in the building/ campus LAN, thereby easing the introduction of switching.

It is sensible to run a private scheme if the network is not destined to be directly connected to the public Internet. Many organizations will be in a position to run private addressing, since email is often the only application that touches the public IP internet and this can be adequately handled with a gateway approach.

The emergence of address translation gateways has made it easier to make a decision to run private IP addressing, since these products provide an escape route into public address space should the organization wish to support a direct connection to Internet for terminal, FTP and other services at a later date.

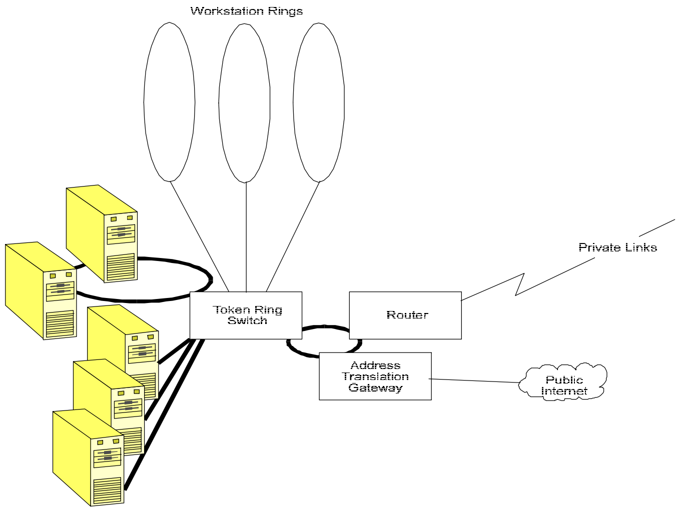

Figure 11: Address Translation Gateways can link private networks to the Internet and provide a security firewall

The IP address translation gateway provides a connection between a private IP network and the public Internet. It holds a limited set of public IP addresses which are statically or dynamically mapped onto IP addresses in the private network; the IP addresses of frames to and from internet are changed as they pass through the gateway. Whilst some network layer services are not always available through the gateway, key internet applications such as Telnet, FTP & email are well supported. Given the constraints of public IP address space, private addressing is likely to gather momentum in the marketplace, especially when vendors start to combine address translation and firewall security for internet access. Naturally, the wide scope of private IP addressing encourages the use of a single IP host space for switching in HQ LAN.

IP Solution 5: IPng or IPv6

The solutions and system rules for switching IP fix a short term problem with a shortage of IP address space. The IAB standards authority for TCP/IP are working on the next generation of IP known as IPng (*IP next generation)* or IP version 6 which is the long term fix to the problem. IPv6 will introduce a number of changes to the protocol, including a shift away from the current 32 bits of addressing to a new 128 bit address. Although the timescale for the global rollout of IPv6 in still in flux, there is a strong commitment behind implementation of the new standard. The wider address space of IPv6 allows even greater flexibility in the use of switching in the IP network.

Supporting IPX in a switched LAN

Out of all protocols running on Token Ring LANs, Novell's IPX suffers the worst degradation in performance when running through a store and forward hop. Cut though switching as a technology is a clear performance leader for Novell environments, but there is a belief that routing is still required to build highly resilient networks. The concerns are twofold:

- multiple connections into a file server have required multiple IPX domains, which in the past has forced a level of routing in the building LAN
- the level of service advertisement (SAP) broadcasts can become excessive in large networks.

This section works through these issues and recommends solutions.

Multiple server adapters and implications

By default, each adapter running into a Novell server requires to be on a different network number. Novell servers will not tolerate overlapping IPX domains. Until quite recently, these two factors presented a significant problem when building a highly resilient switched backbone in NetWare environments. Providing duplicate or multiple connections into a single server offers a compelling fault tolerant and high bandwidth solution for key central file server resources, yet traditional NetWare models suggest that routing is the only way to build such a network configuration.

The problem is outlined below:

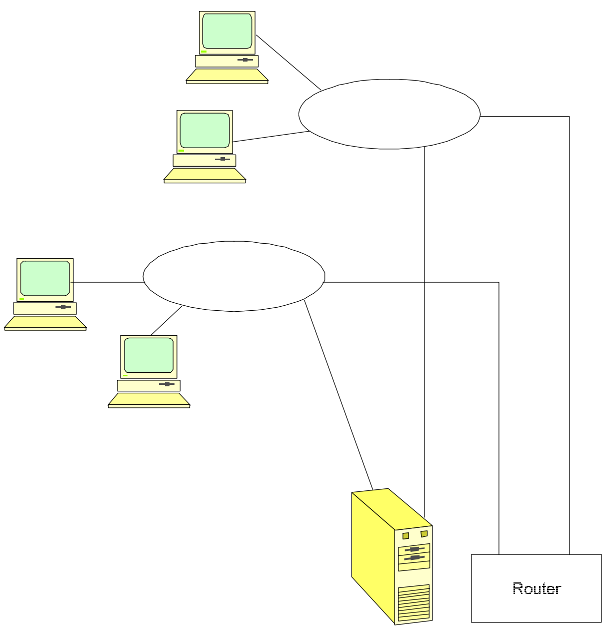

Figure 12: The old way: routing between IPX networks

The example is a common configuration. A file server has independent connections serving workstation rings on the network, a resilient path is provided by a LAN interconnect device linking the two rings together. The traditional NetWare model demands that the interconnect device be an IPX router. This simple model illustrates a bigger problem; it suggests that any network model which includes resilient connections into one or more servers requires that multiple IPX networks and routers must be a part of the network solution.

This has been a key objection to implementing a wholly switched collapsed backbone, but with the rise in popularity in LAN switching and the performance benefits switching offers specifically in a NetWare environment, Novell themselves have responded with new NetWare 4.1 drivers and upgraded 3.1 drivers to avoid routing in the in building LAN.

IPX Solutions

The switch solution for NetWare environments is to avoid local IPX routing altogether, creating a flat single IPX network number for the whole of the building/campus LAN.

Standard software from Novell is now specifically targeted to allow this operation in switched networks. From NetWare 4.10 onwards, Novell have updated the IPX router in the server to allow an “IPX routing off” configuration option . This allows the same IPX number to be bound to two or more adapters. Each adapter can be plugged into either the same ring or switch connected rings with the server load-sharing between them. The obvious benefit is resilience in a switched network: if one adapter (or the network connection) fails, IPX on a workstation will simply reconnect to the other. But a secondary benefit is bandwidth: the load sharing capability and opportunity to dedicate multiple 16Mbit token ring (or 32Mbit full duplex token ring) both contribute to increased capacity in and out of the server.

The new functionality is implemented in the NLM “IPXRTR” which can be loaded with the switch option ROUTING=NONE The program is fully documented in NetWare 4.10 manuals.

Novell’s commitment to switching has encouraged them to release an upgrade for NetWare 3 too, which is now available. For NetWare 3.11, 3.12 and 4.0x, Novell have provided an upgrade “IPX Upgrade for NetWare Servers 1.0”.

These upgrades allow new high performance and highly resilient network solutions for centralized server applications, taking advantage of Ringswitch low latency and ability to directly attach servers with full duplex token ring.

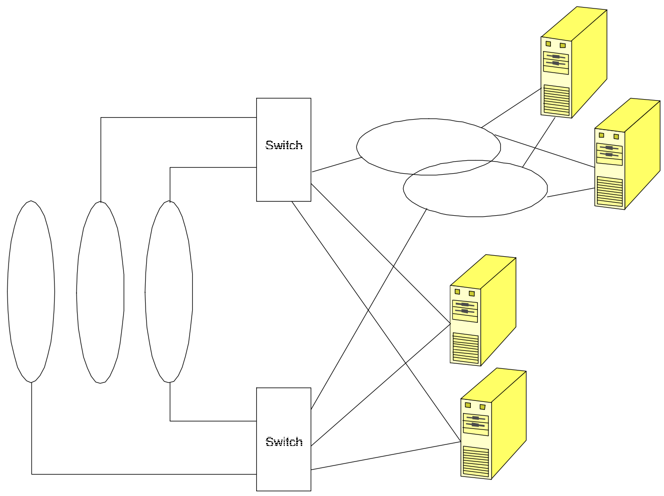

Figure 13: The new way: A resilient, switched backbone for Novell based client/server networking

The above dual switch example has duplicate paths into all servers with key central servers running full duplex Token Ring connections. The workstation rings on the far left hand side are shown singularly connected, with connections interleaved between the two switches. Critical workstation rings can be resiliently connected into a port on each switch. Ringswitch uses a combination of spanning tree and source routing to resolve network loops.

NetWare SAP broadcast control

IPX SAP (Service Advertisement Protocol) broadcasts are central to discovering and maintaining visibility of servers and services in a NetWare environment. A significant point about the NetWare centralized server architecture illustrated above is that the network now segregates into either workstation or server rings.

This is a strong and administratively simplistic basis on which to control broadcast in a NetWare collapsed backbone network; the key is that workstation rings need not receive the hum of SAP broadcasts from local services (such as printers) on every other workstation ring. Removing this volumous and pointless broadcast information from workstation rings cuts right into a broadcast problem that otherwise can become troublesome as the network expands.

Source Route and NetBIOS Broadcast control

Broadcast control is a key requirement in large switched backbone networks. Ringswitch processes all broadcast frames form all switch ports and forwards them according to pre-set configurations. The concept is known as *active broadcast control.* Active broadcast control is splits into two forms;

Ring type broadcast control and virtual LANs

> Automatic broadcast control strategies intended to significantly reduce broadcast traffic with minimal configuration.

Advanced broadcast control

> Highly targeted broadcast control for specific protocols (IPX, IP and NetBIOS) offering a greater degree of control and configuration.

Ring Type Broadcast Control and Virtual LANs

Ring type broadcast control

Ring type broadcast control reduces unnecessary broadcast traffic from the knowledge of whether a ring carries workstations only or whether a ring has servers or other central resources attached to it. This filtering is targeted specifically at NetBIOS, IPX and token ring source route, but not the well controlled TCP/IP broadcasts. This type of broadcast control is highly effective for IPX printer SAP broadcast control, workstation generated NetBIOS Names[4] broadcasts and workstation originated source route broadcasts, none of which need to be re-generated on any other workstation designated ring.

Virtual LANs

Virtual LANs enable broadcast control strategies in two forms:

- The ability to build separate logical LANs where targeted and broadcast frames are not allowed to cross pre-configured port boundaries. This strategy controls broadcast by logically segmenting one part of the building LAN from another, using a gateway or router as an external communication between the isolated segments.

- The ability to define overlapping Virtual LANs when common resources are shared between two or more workgroups. This concept is similar to the ring type broadcast control. Broadcast frames generated in the non overlapping part of each virtual LAN are contained, broadcast frames generated in the overlapping space are forwarded to all participating virtual LANs.

Advanced Broadcast Control

Advanced broadcast control software provides solutions for very large NetBIOS and IPX networks. For NetBIOS, names learning and caching offers a very tightly controlled broadcast filtering regime. For IPX, advanced SAP filtering and additional workstation broadcast control strategies are configurable. Specific details on these filtering options are beyond the scope of this paper. They are covered in the Madge guide “Broadcast Control Strategies for Switched Token Ring LANs” due for publishing in the second half of 1995.

Conclusions

This paper has outlined some complex system issues to be considered for a strategic approach to networking in corporate HQ and large local Token Ring LANs. The whole thrust of the paper has been one of simplifying LAN administration and increasing LAN performance through the use of Token Ring switching. While local routing in some cases allows a stay of execution on LAN simplification by allowing sophisticated configurations to overcome the address limitations of IP or the hostile operation of IPX NetWare servers without the latest upgrades for switching, it is at the expense of a general lower cost of ownership. In the end it is clear that the role of routing in the local LAN has been usurped by switching.

Router vendors have recognized the change in terms of simplicity and performance that switching as a technology brings to HQ LANs, which is why there has been such a flurry of activity an this front. At the same time, application vendors, such as Novell and Microsoft have been quick to adapt core client/server solutions to embrace LAN switching, and IP standards activity has been pushed in the same direction.

In the near future, we will see new IEEE 802.1 (local bridging) initiatives to include a standard way of controlling broadcast within local networks. This will be the final ratification for switching as the way forward.

In the meantime, Madge Networks intends to remain as an authoritative vendor in switching solutions and will continue to lead and implement new standards in Token Ring and ATM switching as they emerge in the last half of the 1990's.

Further Reading

- K. Washburn &J.T. Evans discuss practical strategies for IP addressing schemes in chapters 3 on their book "TCP/IP - Running a Successful Network" (ISBN 0-201-62765-5).

- The full details behind the choice of private Vs public space and recommendations for private addressing schemes are also covered in this book.

Amsterdam
Atlanta
Bangkok
Berlin
Boston
Brussels
Cape Town
Chicago
Cologne
Copenhagen
Dallas
Denver
Detroit
Frankfurt
Hamburg
Hong Kong
Jakarta
Johannesburg
Kuala Lumpur
London
Los Angeles
Madrid
Minneapolis
Munich
Nashville
New York
Ontario
Oslo
Paris
Philadelphia
San Jose
Seattle
Singapore
Stockholm
Sydney
Tokyo
Toronto
Vancouver
Washington DC

Madge Networks

Europe, Middle East, Africa
Knaves Beech Business Pk
Loudwater, High Wycombe
Bucks HP10 9QZ
England
Tel: 44 1628 858000
Fax: 44 1628 858011

Americas
2310 North First Street
San Jose
Calif. 95131-1011
United States
Tel: 1 408 955 0700
Fax: 1 408 955 0970

Asia
64-01 Central Plaza
18 Harbour Road
Wanchai
Hong Kong
Tel: 852 2593 9888
Fax: 852 2519 8022

Japan
Believe Mita
43-16 Shiba 3-chome
Minato-ku, Tokyo 105
Japan
Tel: 81 3 5232 3281
Fax: 81 3 5232 3208