3Com

LANplex High-Function Switches

Delivering Performance and Services for the Network Core

LANplex High-Function Switches
Delivering Performance and Services for the Network Core

Contents

LANplex High-Function Switches

Delivering Performance and Services for the Network Core

By Bob Gohn and Brendon Howe

Bob Gohn is a product line manager for the LANplex switching platforms in 3Com's switching division. Before coming to 3Com, Bob worked in various development and marketing roles for SMC, Raycom Systems, Netways, and Hazeltine.

Bob holds an M.S.E.E. from Polytechnic University and a B.E.E.E. from the State University of New York at Stony Brook.

Brendon Howe is a product line manager for LANplex system software and network management in 3Com's switching division. Before joining 3Com, Brendon worked as a senior product manager for Proteon, and held both marketing and engineering positions at Analog Devices.

He holds an M.B.A. from New York University and a B.S.E.E. from Boston University.

Network administrators worldwide are facing an unprecedented demand for bandwidth on local area networks (LANs), and this demand is virtually certain to increase over time. Bandwidth needs, driven by increasingly powerful desktops, servers, client/server and groupware applications, and the internet/intranet phenomena, are placing an enormous burden on traditional LANs based on conventional hubs, switches, and routers. Networks are now more complex than at any time in the past, and network administrators are challenged to flexibly expand their networks and manage operations and services while keeping ownership costs in check. Administrators therefore require solutions that ensure reliability, enhance performance, and enable effective management of network growth.

Many planners have added LAN switching on an ad hoc basis to address these exploding bandwidth requirements. However, continued network expansion, emerging technologies, and new multimedia applications require that LAN switching be strategically designed into the network rather than deployed as a tactical, local bandwidth booster. This is particularly relevant in collapsed backbone router networks where traffic control is essential.

This paper will demonstrate how a new breed of switching solutions called High-Function Switches satisfies all of these criteria, and how 3Com's LANplex® High-Function Switches provide features, flexibility, and performance to build switch-based networks that deliver a strategic, competitive advantage for the core of today's networks.

The Evolution of Scaling Network Performance

Networks have rapidly evolved from completely flat topologies that relied on bridges for basic LAN connectivity to hierarchical structures created by the deployment of routers. Routers were needed for traffic control and security as LAN connectivity grew, and as the demands for controlling broadcast propagation (akin to bridges) also grew. Routers also provided interconnection of multiple user subnets. But as high-powered servers were centralized at the core of the network, routers began to have a negative impact on network performance. The "80/20 rule" of traffic patterns changed—where once 80 percent of the traffic remained local to a workgroup, 80 percent of the traffic was now traversing the backbone (through the router), destined for centralized server farms. Latency increased and network performance became sluggish as routers simply couldn't handle the high traffic loads. In client/server environments, the complexity of router administration also became a major resource drain on network managers.

At this point, switches emerged on the market, and rapidly became the technology of

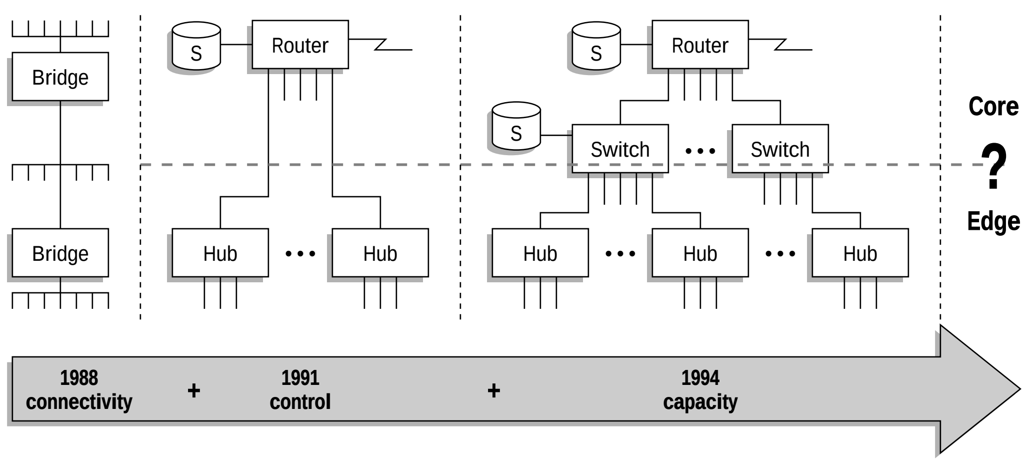

Figure 1. *Evolution of Scaling Network Performance*

choice for creating tactical boosts in available bandwidth and network performance. The operative word in this stage of the evolution became capacity, since more capacity was needed to satisfy the demands of these new client/server traffic patterns.

Switching technology also yielded new price/performance ratios unattainable by traditional solutions. During this period, switches successfully aggregated hub-wired workgroups and improved connectivity to data centers or high-speed backbones. As a result, routers gradually found their way to the periphery of many networks, providing WAN connectivity. Routers continue to be used today in LANs that demand very high levels of control and security.

Networks worldwide are in various stages of evolution. Many LANs remain based on shared Ethernet or Token Ring backbones utilizing two-port bridges for connectivity. Others are using first-generation switching solutions, while yet others have migrated to ATM switch-based infrastructures. The rate at which a network grows and changes depends on many factors. Clearly, as expansion continues in application bandwidth requirements, desktop processing power, and internet/intranet usage, network managers must consider how to move to the next step in the evolution of their LANs' performance (Figure 1).

Traffic Meltdown at the Core and Edge

Since the need for services has evolved faster than the network's ability to support them, conventional routers and switches can fall short of satisfying the most demanding network needs, creating sluggish, inefficient performance or, worse, network meltdown. At the edge of the LAN (typically the workgroup or desktop), a shortage of network capacity, coupled with the proliferation of broadcasts and multicasts propagating throughout the network, is creating a significant network challenge. As we address this challenge, we find a new challenge at the network core (typically the data center). When the demands imposed by the edge exceed the bandwidth capacity at the core, buffer overruns create capacity overload and lost packets, reducing the availability and reliability of the network (Figure 2). Many of today's transitional core switches, while they may be able to provide capacity, cannot deliver the benefits of broadcast containment and security originally associated with the classic collapsed backbone router. As a consequence, networks and their users are suffering. Congestion, suboptimal server access, and slow response times

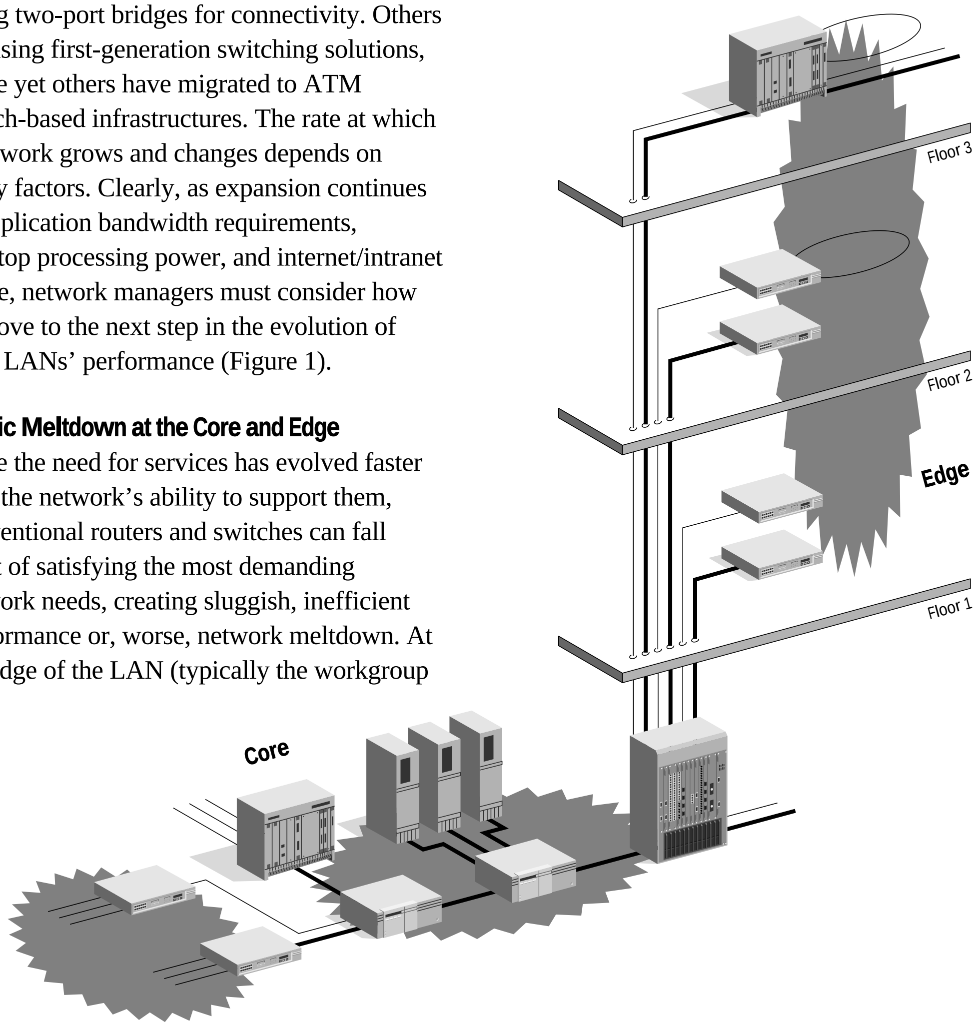

Figure 2. *Network Meltdown Can Occur at the Network Core or Edge*

Acronyms and Abbreviations

ARP
Address Resolution Protocol

ASIC
Application-specific integrated circuit

ATM
Asynchronous Transfer Mode

DVMRP
Distance Vector Multicast Routing Protocol

ELAN
Emulated LAN

FDDI
Fiber Distributed Data Interface

FFAST
Fast Ethernet, FDDI, ATM Switching Technology

HSI
High-Speed Interconnect

IGMP
Internet Group Management Protocol

IP
Internet Protocol

IPX
Internetwork Packet Exchange

ISE
Intelligent Switching Engine

LANE
LAN Emulation

RAP
Roving Analysis Port

RISC
Reduced instruction set computing

RMON
Remote monitoring

VLAN
Virtual LAN

WAN
Wide area network

translate directly into inefficient company operation and potential network meltdown.

New solutions need to meet the requirements of both the core and edge by combining connectivity, capacity, and control. The ultimate goal for delivering these services is to distribute simplicity to the edge and centralize complexity at the core. Edge services require connectivity that is both economical and flexible enough to handle future growth. Control should be provided through quality of service and simple, plug-and-play functionality. Sufficient desktop bandwidth and low latency should deliver the necessary capacity for edge operation.

The core of the network provides services to an enterprise-wide community of computing interest, so its requirements are much different than those at the edge. Connectivity at the core must provide high port densities to scale network growth, and must be extremely resilient. Because the core is also a control point for virtual LANs (VLANs), network segmentation, and network management, core services must be scalable and able to efficiently isolate traffic and define broadcast domains. High capacity levels must support bandwidth to enable users at the edge to access resources such as server farms in the core, even during periods of peak demand, without compromising the network's performance. Otherwise, communications will be impaired.

The Next Step: High-Function and Boundary Switching

To preserve their investments in legacy and client/server LANs, administrators must leverage new technologies to resolve this performance dilemma triggered by all the new applications, services, and uses. Conventional routers and switches, however, were not designed to meet the demands of today's LANs, and they lack the capabilities administrators need to manage the bandwidth explosion and to guard against "network meltdown."

3Com has developed a new generation of advanced solutions to meet these evolving performance demands by delivering core and edge services through High-Function and Boundary

Beyond the Traditional Solutions

Routers, long the traditional devices delivering core services for LANs, have not kept pace with networking needs. Routers are too complex for many applications, especially as network segmentation increases, and too costly for most core requirements, and they cannot scale performance as bandwidth demands increase. Routers are optimized for WAN connectivity where their cost and complexity are justified by improved leverage of expensive and limited WAN bandwidth.

Conventional switches used in today's network core are also inadequate at providing the services required to manage network bandwidth. Pure ASIC-based switches cannot support these services, and processor-only switches deliver unacceptable performance when core services are invoked. This severely handicaps efforts to manage bandwidth or scale performance to sustain data-intensive core services. Neither conventional routers nor conventional, pure layer 2 switches were designed to meet the changing demands of today's networks (Figure 3).

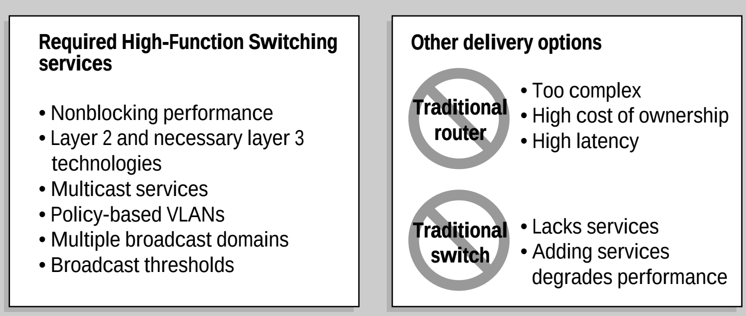

Figure 3. *Service Requirements at the Network Core*

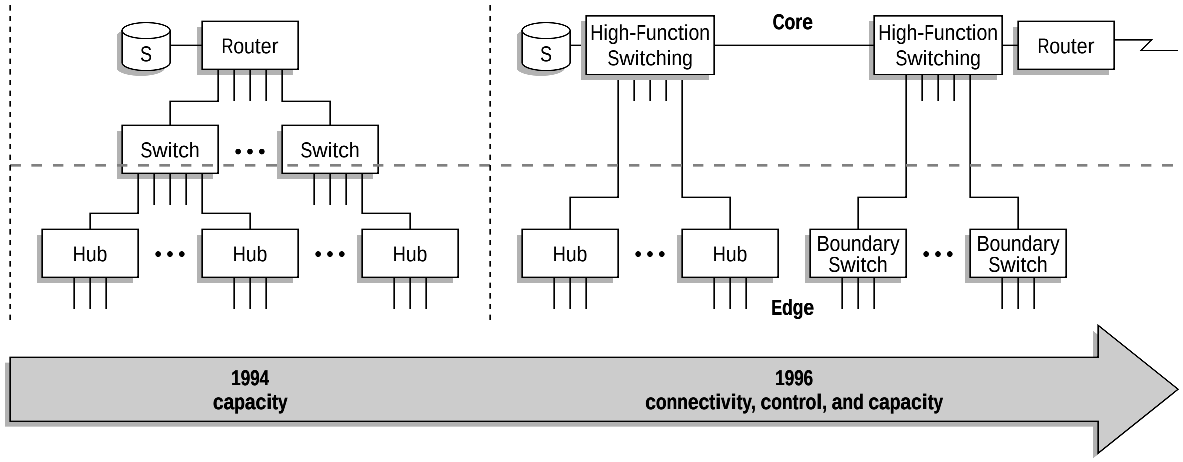

Figure 4. *The Next Step in Scaling Network Performance*

Switching (Figure 4). The switching solutions are part of a larger framework that 3Com calls Transcend® Networking, which is a blueprint for designing, building, and managing enterprise networks from the user's desktop to remote sites. These purpose-built switching solutions provide an advanced combination of connectivity, control, and capacity not found in other devices today. High-Function and Boundary Switches provide the best pairing of functionality and cost. They are designed around ASICs and ASIC-plus-RISC architectures, and incorporate a suite of innovative features that provide elevated levels of performance and functionalities never before available in switching technologies. They allow administrators to cost-effectively manage and control bandwidth at the core, while distributing simplicity and performance to the edge.

3Com's High-Function Switches are delivered in a variety of form factors to meet a wide range of performance, functionality, and packaging needs. The LANplex family of packet-optimized High-Function Switches is ideal for migrating collapsed backbone router networks to switch-centric networks. The ONcore® Integrated System is a high-performance, multifunction, modular family of products that supports both Boundary and High-Function Switching, optimized for migrating shared LANs to switched internetworks. The CELLplex™ family consists of market-leading, cell-optimized High-Function Switches that offer the robustness, flexibility, and rich feature set to support ATM-based backbones today.

Staying a Generation Ahead: LANplex High-Function Switches

As the established market leader for many years, the LANplex family of High-Function Switches provides unprecedented performance and functionality for high-end core applications. They are available in two platforms: the LANplex 2500 and the LANplex 6000. The LANplex family supports and eases the router-based to switch-based migration in mission-critical LANs by combining performance migration, bandwidth management, and distributed switch management.

The LANplex 6000 is a no-compromise backbone aggregation device with a powerful multi-gigabit backplane to support a variety of technologies from switched Ethernet and Token Ring to high-density switched Fast Ethernet and FDDI, along with ATM. The LANplex 6000's modular design, high-speed backplane, and distributed ASIC-plus-RISC processor architecture delivers scalable performance and technology migration, comprehensive management, and fault tolerance.

Each LANplex 6000 module has its own ASIC-plus-RISC engine, so performance and intelligence scales as modules are added. Managers can quickly grow their networks by simply introducing new modules rather than installing, integrating, and configuring new systems. In addition, should one module fail, the others will continue to perform unaffected. This level of scalable performance allows network administrators to meet present networking needs and strategically plan their LANs much more effectively than is possible with conventional switches and routers. Further, the LANplex 6000 specifically contrasts with many competitive solutions that neither scale nor offer resiliency because they rely on some form of single, central processing engine.

The LANplex 2500 is unmatched in delivering on network migration needs. Dubbed the "ultimate migration machine," the LANplex 2500 delivers robust, market-leading

Transcend Networking

Networking has become a pervasive part of our world, but managing today's growing networks can be a daunting challenge. 3Com is committed to making the complexities of networks invisible to the user, and to making network management flexible, simple, and unconstrained. This means providing technology that makes the network fundamentally easier to design, install, maintain, and evolve.

The cornerstone of this commitment is Transcend® Networking, a unique, global network framework that enables network managers to harness the power of their networks. Transcend Networking takes a three-vectored approach to evolving networks (Figure 5). Each vector delivers technologies, innovations, and solutions that can be tailored to each customer's state of network evolution and ever-changing network needs. The three parts of Transcend Networking are:

- *Scaling performance.* Based on the most advanced switched infrastructure
- *Extending the reach.* Connecting remote sites and users through router and remote access solutions
- *Managing the growth.* Managing networks through the most comprehensive distributed management solutions

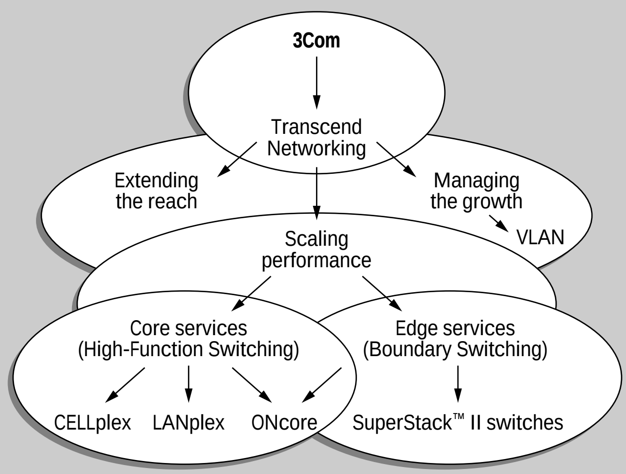

Figure 5. *Transcend Networking Framework*

The driving principles of Transcend Networking include:

- Centralize complexity and distribute simplicity
- Remain technology- and platform-neutral
- Protect network investments and provide an incremental evolution by leveraging a standards-based, multivendor architecture
- Ensure comprehensive systems-level management—not simply device management

High-Function and Boundary Switching are an integral part of scaling performance within the Transcend Networking framework. The strategy of scaling network performance distributes bandwidth and simplicity to workgroups by using high-performance, low-cost Boundary Switches. At the core, where traffic aggregates, the strategy is to deploy modular High-Function Switches that provide bandwidth control to optimize core services and maintain high performance. By providing connectivity, control, and capacity throughout the enterprise, this paradigm offers the best long-term solution. It also enables users to best match network costs to their applications and provides a broad migration path to new services.

performance and multitechnology backbone support including switched Ethernet, FDDI, Fast Ethernet, and flexible ATM-legacy LAN integration. Based on a reliable ASIC-plus-RISC design, the LANplex 2500 is one of the few solutions—perhaps the only one—that can integrate multiple technologies in one platform at its price/performance and functionality levels. The benefit to the customer is tremendous flexibility in the choice of backbone and server technology—today and tomorrow.

Performance Migration

As High-Function Switches, the LANplex family is characterized by advanced functionality, data-center class performance, modular packaging, fault tolerance, and technology migration, and boasts a host of performance migration features.

LANplex Switching Architecture: ASIC plus RISC

What separates the LANplex from conventional solutions is the power of the LANplex Switching Architecture. In contrast to conventional switches, which rely on ASICs to process traffic, and other switches, which rely solely on RISC processors, each LANplex tightly couples an ASIC and an intelligent RISC processor (Figure 6). The startling advantage of this unique architecture is its ability to apply intelligence to the device, such as adding software-defined bandwidth management techniques, without slowing performance of the switch. 3Com's Intelligent Switching Engine (ISE-chip) ASIC is found in both the LANplex 6000 and LANplex 2500. The LANplex Fast Ethernet, FDDI, ATM Switching Technology (FFAST-chip) is a second-generation ISE ASIC that provides the basis for 100 Mbps-plus switching in the LANplex 6000 series. Using the ISE and FFAST ASICs, managers are assured of non-blocking wire-speed performance.

Switching performance is further enhanced by packet assistance schemes that implement packet preprocessing. The ASIC provides fast lookup of all source and destination addresses as well as packet decoding, while the RISC applies intelligent forwarding decisions; the result is faster LAN transactions. No other switch can offer this.

The open LANplex Switching Architecture allows new functionality and standards to be added. For example, as VLAN standards such as IEEE 802.1Q emerge, the LANplex architecture can easily support them through a software revision. (Competitive technologies would require a complete respin of their architecture.) This means that the LANplex can bring new functionality and standards to market much more quickly without a forklift upgrade.

High-Speed Interconnect Bus

The LANplex 6000 adds a High-Speed Interconnect (HSI) bus and three internal FDDI rings, providing a total of 19.5 Gbps of backplane capacity, to the ASIC-plus-RISC architecture found on each of the LANplex switching modules (Figure 7). The total throughput capacity of a fully configured LANplex 6012 is an amazing 6.2 million packets per second—exceptionally robust performance for data center applications. The HSI's "spatial reuse" lets more than one module use the backplane simultaneously.

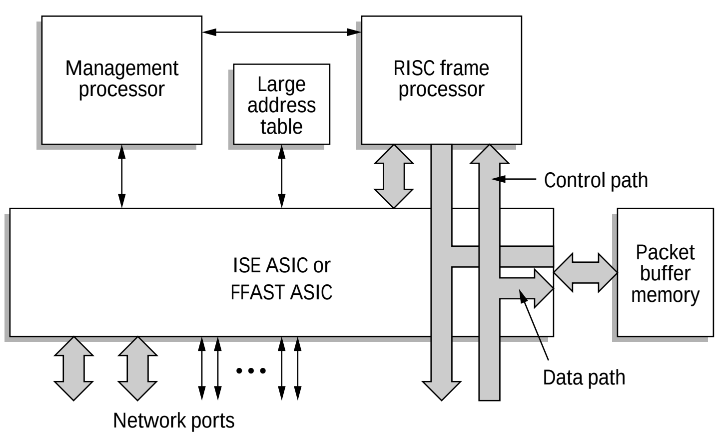

Figure 6. *LANplex ASIC-plus-RISC Architecture*

ATM and High-Function Switching

With the introduction of an ATM module for the LANplex 2500, 3Com is taking High-Function Switching to new heights of flexibility and performance. The ATM module is based on industry standards for interoperability and provides a single 155 Mbps OC-3c interface. It can be used together with FDDI or Fast Ethernet uplinks, providing a convenient migration mechanism.

The ATM module supports LAN Emulation (LANE) 1.0 for powerful mapping between High-Function Switching, VLANs, and Emulated LANs (ELANs), and for extension of VLANs across the ATM backbone. This mapping is implicit and does not require any type of frame tagging. In the same way, users can map subnets to ELANs and apply LANplex routing to IP subnets across ATM. For IP users, the LANplex 2500 ATM module also provides Classical IP over ATM (RFC 1577), allowing extension of IP networks across ATM.

The combination of LANplex High-Function Switching and ATM support makes for a powerful solution in ATM-based networks.

This means that, unlike conventional shared contention-based backplanes, the LANplex allows multiple concurrent communications paths, dynamically created, providing high-performance switching. This LANplex 6000 design allows for the migration from 10/100 switching to full 100 Mbps Fast Ethernet and FDDI switching and, in the near future, ATM support. This level of scalable performance lets network administrators meet current networking needs and strategically plan their LANs much more effectively than with conventional switches and routers.

Buffer Management with Elastic Packet Buffering

LAN traffic is bursty by nature, and peak demand can soar. For example, when users log into the network's centralized servers at the start of the work day, or when overnight network backup services are invoked, heavy data streams can overwhelm the buffers of conventional switches. The switch architecture must be able to absorb these peaks to avoid delays and lost communications through dropped packets—keeping the network available, users happy, and businesses up and running.

Many conventional devices provide an unmanaged buffer system that relies solely on static buffers, which allocate fixed amounts of buffer memory to every port (Figure 8a). The problem here is that no port can access more than its allocated amount of buffer space, even for an instant. As a result, the switch may drop packets on ports that are out of buffer space, even though other ports are not using their buffers at all. The result is dropped communications.

With dynamic packet buffering, on the other hand, buffers are set aside in a shared pool, which can be divided dynamically as congested points come and go in the switch (Figure 8b). However, there are disadvantages to this buffering scheme, since it takes system overhead to manage buffer allocation, which can degrade performance. Furthermore, one exceptionally busy port can consume all available buffers, leaving none for the other ports. Again, the result is dropped communications.

The LANplex Switching Architecture delivers a much more sophisticated buffer management scheme called Elastic Packet Buffering, which provides high levels of reliability during periods of peak packet demand (Figure 8c). The LANplex Elastic Packet Buffering feature provides a unique and powerful combination of static and dynamic buffering. The static pool of memory (64 KB per Ethernet port, 512 KB per high-speed port) ensures fairness to each port, while the dynamic pool (2 MB per switching engine) acts as a reserve to absorb bandwidth peaks. Elastic Packet Buffering ensures that ports always have sufficient memory available, even during extreme traffic loads. Because Elastic

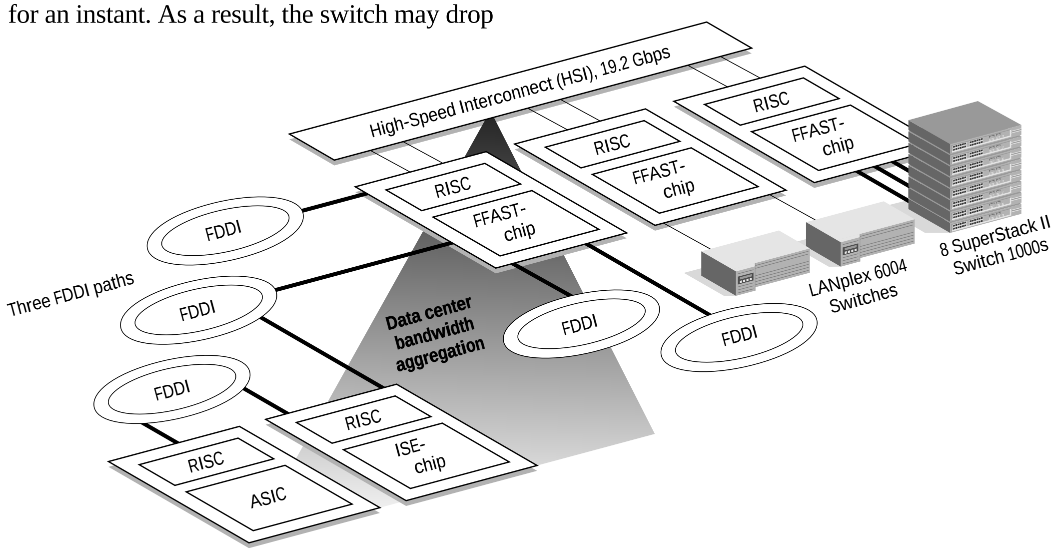

Figure 7. *LANplex 6000 Architecture*

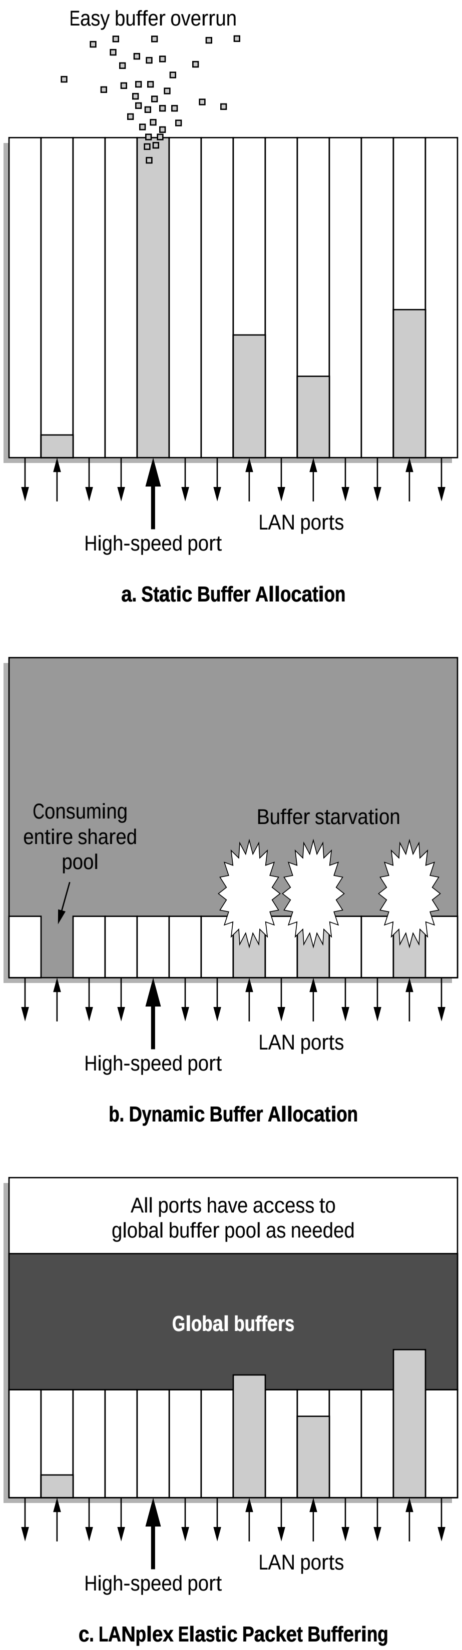

Figure 8. *Buffer Management*

Packet Buffering is implemented in silicon and is an inherent part of the LANplex Switching Architecture, there is no sacrifice in switch performance—even at peak demand. This capability reduces the probability of dropped packets and delayed communications in the network core, where 80 percent of client/server traffic aggregates. As a result, the network remains available, users are happier, and businesses can sustain their operations and services at the highest possible level. Though it may sound simple, the Elastic Packet Buffering approach is unique in the industry, and provides significant advantages over other approaches.

Bandwidth Management

Bandwidth management provides the necessary traffic control at the core of the network, allowing network administrators to improve network performance while enforcing traffic flow and other network policies. Bandwidth management also enables the network administrator to manage the growth of the network by providing mechanisms to ease the migration from router-based network cores, while preserving the existing logical network infrastructure or simplifying the migration to a flatter network environment.

Demand for bandwidth due to new applications, new traffic patterns, and the abundance of bandwidth funneling in from the edge of the network are overwhelming conventional network designs. Traditional switches are unable to efficiently distribute and control bandwidth across the LAN because they are either based entirely on ASIC-only technology and lack the necessary functionality, or they are RISC processor–based and slow in performance when high functionality is invoked. Conventional routers, which may have performed some of these high functions in the past, do not scale well in today's large switched infrastructures. The results are lack of control over network traffic and overextended network resources, both of which impede business-critical communications.

To alleviate these conditions, 3Com's High-Function Switches are purpose-built to deliver the benefits not achievable by either

Performance Migration Issues

What are the key performance migration issues facing today's network administrators?

- Bandwidth capacity for the core and edge. *To ensure a smooth migration in performance, network managers need to look at solutions that have the capacity to handle the bandwidth explosion at the core and edge of the network. LANplex High-Function Switches deliver this with the LANplex Switching Architecture, a unique coupling of ASIC-plus-RISC technology.*
- Investment protection of current infrastructure. *LANplex High-Function Switching lets organizations evolve their infrastructures one step at a time with proven technology.*
- Long-term viability of additional investments. *The scalable, modular architecture of the LANplex lets it grow and change with the network.*
- Integration of technology. *The open platform of the LANplex system offers flexible interfaces and functions and is easily adaptable to a variety of current and future network configurations. Because the LANplex architecture was designed to support future needs, its total cost of ownership is far lower than competitive solutions.*

Bandwidth Management Issues

What are the key bandwidth management needs for which today's network administrators must find solutions? And what solutions are provided by 3Com's LANplex?

- Logical user associations for policy/bandwidth enforcement. *LANplex solution: Policy-based VLANs defined by port, MAC address, and/or protocol.*
- Support emerging multimedia applications. *LANplex solution: AutoCast VLANs and IP multicast routing.*
- Preservation of address structure. *LANplex solution: Intranetwork routing of mainstream protocols.*
- Secure resources and control access. *LANplex solution: User-defined packet filtering.*
- Fault protection. *LANplex solution: Broadcast firewalls.*

conventional switches or conventional routers. LANplex switches allow administrators to create, control, conserve, allocate, and absorb bandwidth quickly and easily.

The LANplex's software architecture adds sophisticated traffic control mechanisms to the tightly coupled ASIC-plus-RISC architecture. The LANplex does this by providing a number of flexible mechanisms for establishing logical views of the network (i.e., grouping) through a unique routing-over-switching model in which the integrated routing and switching engines work together to optimize bandwidth management and other mechanisms for controlling traffic flow.

Logical grouping (often referred to as virtual LANs or VLANs) can be performed based on a very flexible set of policies, and is used not only to provide simpler network adds/moves/ changes, but also to deliver strong performance tuning by providing broadcast containment. This eases network administration by allowing logical groupings to be defined by management domains, and provides security mechanisms by allowing policies to be defined.

The LANplex's routing-over-switching model is designed to support a "switch where you can, route where you must" model. This model is different from the conventional multiprotocol router's side-by-side routing/bridging model. With the LANplex's routing-over-switching model, whenever traffic can be switched within the LANplex, it is switched. When the traffic must traverse a routing entity, the LANplex's architecture allows the traffic to be preprocessed by the ASIC (offloading the processor and bypassing the switching engine to minimize hops), and to be directly processed by the built-in routing engine. This streamlined processing optimizes performance. The LANplex routing-over-switching model also provides very flexible, overlapping capabilities, supporting multinetting and per-port routing. In contrast, a conventional multiprotocol router is designed to route all packets for its routable protocols, causing lower performance, more complex administration, and less flexibility.

Building upon the LANplex's unique hardware and software architectural attributes, the LANplex provides a powerful set of features to deliver the bandwidth management benefits required in the core of the network. Some of these features include: policy-based VLANs, AutoCast VLANs and IP multicast routing, intranetwork routing, broadcast firewalls, user-defined packet filters, and performance and flexibility.

Policy-Based VLANs

The LANplex system provides sophisticated mechanisms for grouping users for the purpose of creating logical views of the network. The benefits delivered by virtual LANs, or VLANs, vary depending on the extent of the switch's VLAN capabilities. For example, some switches offer only a limited set of VLAN membership criteria, such as MAC address or port grouping only techniques. The LANplex allows network administrators to create policy-based VLANs. While VLANs in general are a growing strategic bandwidth management and change management resource for today's businesses, policy-based LANplex VLANs go one step further by providing the flexibility to define VLAN membership based on a wider range of criteria.

In contrast, classic router-based LANs, which may deliver broadcast containment and security benefits to some degree, are able to provide these benefits only to the extent that the group membership corresponded to the users' physical network locations. With the rise of conventional switches, users have to some degree lost many of these benefits.

The advent of High-Function Switching VLANs restores these broadcast and security features, and is augmenting today's switched infrastructures by logically grouping users independent of their physical network location. In doing so, VLANs allow users to regain many of the broadcast containment and security benefits once associated with traditional routers, while improving upon the deficiencies associated with those same classical routers. These deficiencies include slow performance and the burden of complex and costly administration associated with adds/moves/changes.

LANplex policy-based VLANs allow administrators to group users based on a wide range of criteria, including not only layer 1 and layer 2–oriented criteria such as physical port groups and MAC addresses, but also layer 3–oriented criteria such as subnet address, network protocol, and IP multicast group (Figure 9).

The LANplex's policy-based VLANs are extremely flexible, allowing traffic to be contained to specified domains based on a set of rules appropriate for each particular network environment. Network administrators benefit by applying the right type of VLAN to achieve particular network goals, rather than deploying a "one-size-fits-all" VLAN structure.

For example, using subnet-based VLANs, network administrators can deploy switching while still capturing broadcast containment benefits—such as bandwidth optimization and security—previously provided with classical routers. VLANs based on subnets allow preexisting subnet infrastructures to be maintained while reducing or delaying desktop reconfiguration.

As a first step in VLAN configuration, a LANplex front-ends a collapsed backbone router for increased performance—and uses protocol-based VLANs to preserve the mapping of subnets over switched segments. The router is still used to connect subnets. In the second step, a network administrator can replace a classic collapsed backbone router with the routing-enabled LANplex without touching any of the preexisting subnet infrastructure. By enabling routing in the LANplex, the administrator can eliminate the dependence on and cost of an external router to interconnect subnets. Immediately, the broadcast and security domains established by the original router are maintained.

The third step may vary, depending on the network goals. If the goal is to continue along a performance maximization path, the network administrator can begin to flatten out the network by slowly combining or reducing the number of existing routing interfaces on the LANplex. This provides a minimal set of router interfaces to administer and a reduced amount of routing delay. If the goal is to pace the rate of desktop reconfiguration, the network administrator can similarly flatten out the network, but at a pace that is consistent with the size of the network. If, on the other hand, the goal is to maintain separate routed domains, the third step may be nothing more than occasionally defining a new subnet on a LANplex port as users move around (known as multinetting).

Another case in which policy-based VLANs are particularly beneficial is in environments containing nonroutable protocols. Network administrators may wish to prevent traffic flooding (broadcast/multicast/unicast unknown) associated with nonroutable protocols from degrading performance on other parts of the network. For example, the network administrator may wish to isolate NetBIOS-generated flood traffic to a particular set of ports to optimize the available bandwidth, and thus optimize access to file servers, on the remainder of the network (perhaps IP-based).

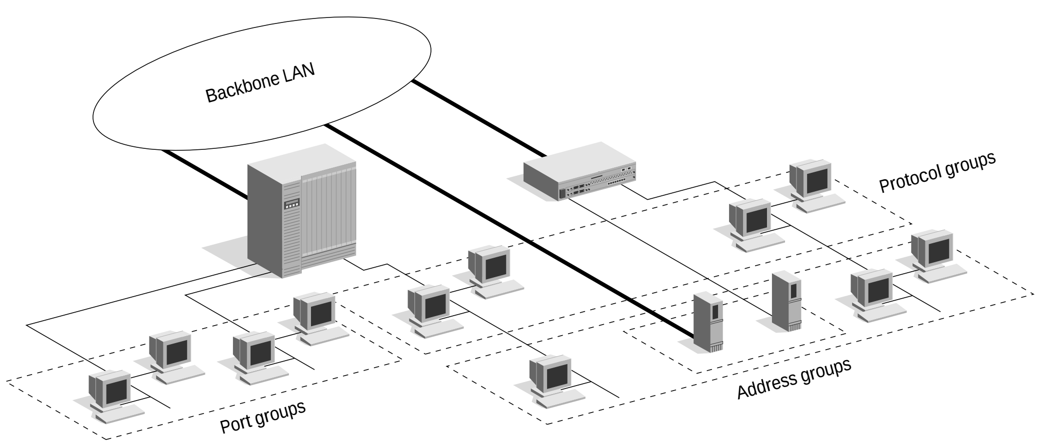

Figure 9. *Policy-Based Virtual LANs*

To do this, the network administrator need only define a NetBIOS VLAN by choosing the ports expected to contain NetBIOS traffic. The result is that all flood traffic originating from any of the defined NetBIOS ports is prevented from reaching any other ports.

Policy-based VLANs are a powerful and flexible bandwidth management tool that sets the LANplex apart from competing technologies. The flexibility lets administrators meet two daunting challenges in today's complex networks—reducing administrative tasks and maximizing performance. Policy-based VLANs enable administrators to precisely meet the communication needs of users and simplify essential management tasks, helping make the LANplex High-Function Switches the ideal migration path for customers with existing collapsed backbone networks.

IP Multicast Support: AutoCast VLANs and IP Multicast Routing

IP multicast support is becoming more frequently used to support network-based audio/video applications. IP multicast provides a standard, efficient, and dynamic method of grouping users through an exchange of Internet Group Management Protocol (IGMP) packets between clients and the IP multicast server. Clients request membership in a particular IP multicast session by periodically transmitting IGMP packets to the server. The IP multicast server and interim IP multicast–supporting devices respond by transmitting data from the appropriate multicast session to those clients requesting membership. Because these applications typically generate heavy broadcast traffic, potentially flooding the network, High-Function Switches are designed to alleviate this bandwidth drain on the network.

The LANplex supports and handles IP multicast traffic in two ways. First, the platform supports AutoCast VLANs in non-routed, or flat, portions of the network. Second, it supports IP multicast routing across subnetted portions of the network.

AutoCast VLANs. AutoCast VLANs are an intelligent switching mechanism that allows the switch to listen to the IGMP packets. By listening, the switch automatically creates VLANs, allowing the forwarding of the IP multicast packets only to registered users, rather than flooding the traffic to all switch ports. This process preserves bandwidth on segments devoid of registered users, and is generally sufficient in completely flat, non-routed networks. The reason for special handling of IP multicast traffic is that all the data packets that are part of an IP multicast session are layer 2 multicast packets and are therefore forwarded out every port of the switch, even if only one client on one segment is participating in an IP multicast session. These sessions are bandwidth intensive, so controlling this traffic is essential in flat portions of the network (Figure 10).

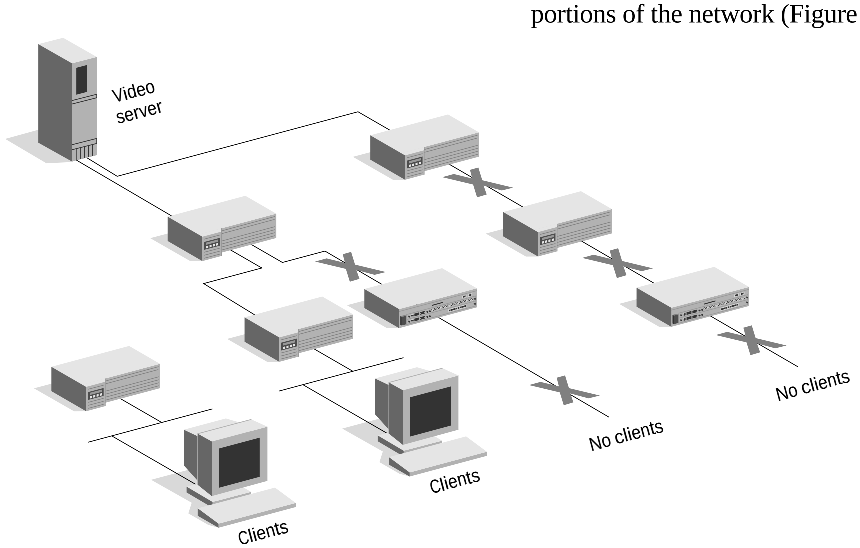

Figure 10. *AutoCast VLANs for Multimedia*

IP Multicast Routing. When IP multicast is deployed in a routed network, IP multicast routers communicate with other IP multicast routers, typically using the Distance Vector Multicast Routing Protocol (DVMRP), to define delivery paths for forwarding multicast packets. IP multicast routing allows the IP router to forward the multicast traffic across subnets normally handled by the router. It also dynamically prunes back unused subnets by periodically checking the status of the registered clients. In a routed environment that does not completely support IP multicast routing, an IP multicast router can actually create a tunnel to cut through the network.

The LANplex support of flexible yet fully integrated IP multicast routing eliminates the dependency on a dedicated router and provides administrators with critical options for meeting the needs of a wide range of multimedia-rich networking environments. Administrators using the LANplex system can conserve bandwidth at both the core and the edge by keeping multicasts off designated segments, thus reducing the stress placed on the core by users' increasing demands for multimedia applications. LANplex is the leader in this application since it supports AutoCast VLANs and IP multicast routing, both important in today's networks, which typically have some degree of routing and switching needs.

Intranetwork Routing

While today's philosophy of "switch where you can, route where you must" reduces the role of routing, some degree of routing functionality will continue to be necessary in local area networks. For example, as LAN infrastructures migrate from the classic collapsed backbone router design toward flatter, switch-based networks, administrators will want to preserve their logical (layer 3) view of the network and the control, fault isolation, and security that routing provides.

The LANplex integrates IP, IP multicast, IPX™, and AppleTalk® routing within the tightly coupled ASIC-plus-RISC architecture. The coupling of 3Com's unique switching and routing engines results in the following key benefits for network administrators:

- The LANplex's support for multinetting and per-port routing allows administrators to vary the routing functionality required when migrating to flatter infrastructures. This capability lets network administrators pace the rate of desktop reconfiguration inevitably required upon flattening the network.
- The LANplex's ability to internally route between interfaces defined by distinct VLANs minimizes the number of routing interfaces and associated administration.
- The LANplex's tightly coupled ASIC-plus-RISC architecture maximizes routing performance by allowing the ASIC to perform the routine preprocessing tasks, offloading the RISC processor and minimizing the amount of code executed by the processor.
- The handoff of the packet from the ASIC to the RISC processor bypasses the layer 2 switching function and goes directly to the routing portion of the code, minimizing the number of forwarding hops. This streamlined forwarding function is not possible with an external router. In addition, all intra-subnet traffic is switched (without routing intervention) at full line rate.
- The internal switch/router interaction of the LANplex outperforms the speed of would-be interface links to one or more external routers.
- The routing function of the LANplex leverages preexisting hardware and avoids the need for external interface links, while minimizing cost and complexity.

Broadcast Firewalls

Yet another form of broadcast containment comes with the LANplex's ability to support broadcast firewalls. Broadcast or multicast storms—commonly caused by faulty protocol implementations, undetected network loops, and faulty network equipment—may cause significant disruption to the network. By configuring a threshold of a reasonable quantity of broadcast packets per unit time, the network administrator can ensure that broadcast levels will not exceed the preestablished threshold. If the threshold is reached, additional broadcasts are filtered until the storm elapses. This has the benefit of diminishing the negative effect of broadcast storms on network performance while still allowing the broadcast or multicast propagation required for some applications.

User-Defined Packet Filters

The LANplex also supports the most basic criteria of VLAN definition: filters defining layer 2 or layer 3 rules. Should a network administrator seek to create rules based on address groups allowed to access certain resources, MAC addresses can be established to define the groups and filters can be applied to each group to enhance network security. Alternatively, the administrator may define port groups for clear-cut traffic barriers for firewalling purposes. In these cases, all traffic—including flood and unicast traffic—is restricted to the defined domains.

Additionally, network administrators may define a set of filters restricting traffic flow based on layer 3–oriented criteria (up to 64 bytes into the packet). The network administrator, for example, may choose to prevent telnet traffic from entering a particular port, or prevent stations with particular IP addresses from reaching certain devices. To achieve this, the administrator will set up a filter to isolate traffic meeting these restrictions to the appropriate ports.

Performance and Flexibility

In all of the above cases, the LANplex utilizes its embedded ASIC-plus-RISC technology to perform ASIC-based preprocessing, followed up with RISC-based processing where needed. This powerful combination not only maximizes performance while controlling flood-oriented traffic, but also leaves room for addressing new, as yet undefined standards emerging in the areas of VLAN, multimedia, and ATM technology. This ASIC-plus-RISC combination shelters network investments from obsolescence and delivers full functionality without compromising performance.

Distributed Switch Management

Although switches often provide a simple and scalable approach to increasing network bandwidth, managing a switched network infrastructure can prove to be difficult for a number of reasons. As network usage expands and becomes more sophisticated, LANs are becoming increasingly distributed and complex. The complexity can increase with the migration from router-based to switch-centric networks.

Traditional management approaches tend to rely on cumbersome polling methods to collect data needed to monitor network segments, using central management stations to collect, store, and analyze the data. As switches are added and more network segmentation occurs (in order to increase bandwidth), the amount of "raw" data that must be collected may increase exponentially and can often become overwhelming. In addition, a lot of network bandwidth as well as switch memory can be consumed in administering the numerous data collections that occur.

Another challenge is that the management characteristics of a switched network are often different from traditional router-based networks. These network management systems most often monitored port status and device characteristics (buffer loads, error states, etc.), whereas switched networks must often be managed by utilization levels, traffic flow patterns, and other network—*not device*—information. This further increases the amount of information that must be polled within the switched environment.

The ultimate goal of network management is to make the switched network easier to manage via self-management tools that let users isolate problems and then invoke specialized capabilities for quick troubleshooting—all without adversely impacting switch or network performance.

The only practical way to manage the switching infrastructure is to offload the central management station's task of collecting, archiving, and analyzing network information to the agents within the switches themselves. This approach provides three important benefits:

- It allows the switch to collect and summarize data for the management station (rather than depending on the central station).

- It provides the capability for the switch to take local action on network events without depending on human intervention (through a central management station).
- It provides the capability for the switch to automatically solicit (“trigger”) help from a specialized resource (such as an RMON probe), independent of the management application that is monitoring the network.

LANplex AutoFocus SmartAgent

As High-Function Switches, the LANplex systems meet these needs by distributing switch network management through an array of innovative management tools. Key to simplifying the switch, and ultimately switched network management, is 3Com’s AutoFocus SmartAgent™ feature that provides a unique combination of capabilities (Figure 11). This LANplex SmartAgent automatically gathers and correlates information from all LAN segments attached to the device, and can take corrective actions locally as problems are encountered. It also has the ability to trigger an external probe for more comprehensive management analysis. The result is around-the-clock, exception-based management that is distributed throughout the network, and a reduction in the amount of polling required by other management tools. Bandwidth is thereby conserved, management control is maintained, and network performance is enhanced.

The LANplex AutoFocus SmartAgent features include:

- *Embedded RMON agent.* The embedded RMON agent in the LANplex provides four groups of RMON: statistics, events, alarms, and history, all on a per-port basis. The LANplex provides these four management groups while maintaining high performance (full wire speed). The LANplex also provides additional statistics beyond the embedded RMON, including buffer management information, transmit traffic statistics, and other switch/router status information.
- *Roving Analysis Port (RAP).* The LANplex’s RAP provides the industry’s most powerful roving port capability. RAP allows an external probe such as 3Com/AXON’s FlexiProbe™ to collect more detailed traffic information (all nine groups of RMON, RMON2) on any switched segment in the LANplex without being physically attached to the segment. This feature allows sophisticated traffic analysis techniques to be applied to the network without burdening the performance (or capacity) of the switch. The RAP works locally (within the same switch) or across multiple LANplex systems with any type of backbone connection (FDDI, 100BASE-T Fast Ethernet, or ATM). For example, a segment at a remote location can be monitored at a central network control center. When the RAP is coupled with the LANplex’s embedded RMON agent, administrators can continuously monitor and

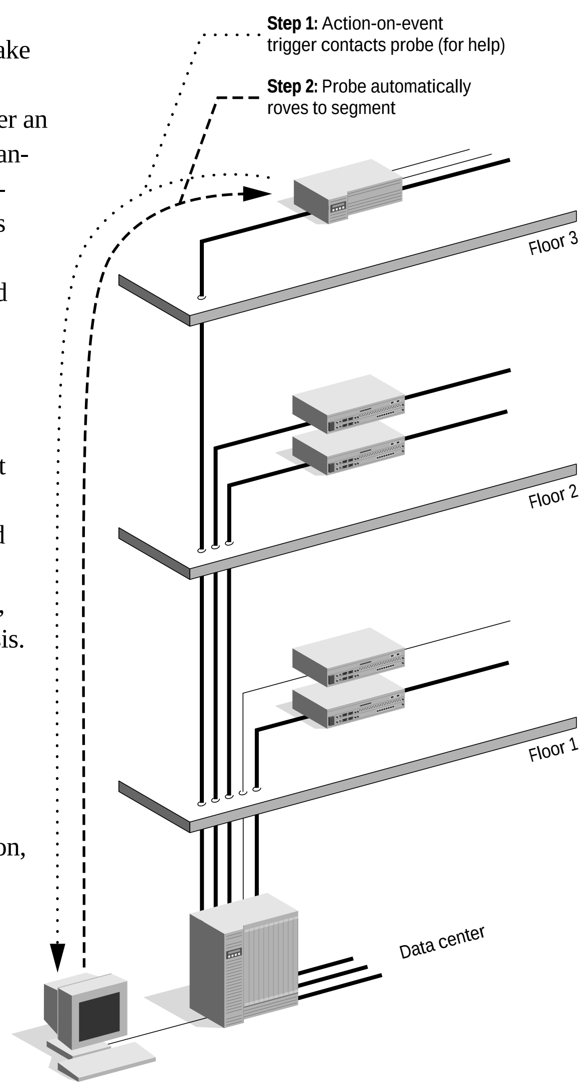

Figure 11. *LANplex AutoFocus SmartAgent*

Distributed Switch Management Issues

What are the key distributed switch management needs for which today’s network administrators must find solutions? And what solutions are provided by 3Com’s LANplex?

- Self-diagnosing, self-healing networks. *LANplex solution: Logical, not physical, deployment of specialized tools and personnel.*
- Fast, efficient utilization of management expertise and equipment. *LANplex solution: Diagnostic and analysis tools built into switched infrastructure.*
- Accurate, timely data for capacity planning and accounting. *LANplex solution: Automatic activation of diagnostic/recovery tools and functions.*
- Around-the-clock, 24 hours a day, seven days a week operation for mission-critical services. *LANplex solution: Designed-in reliability.*

maintain the health of all network segments and quickly perform more sophisticated network analysis if trouble is found on any network segment. This minimizes downtime and data loss across the LAN.

The LANplex's RAP capability is further enhanced by 3Com's LANsentry® and Traf*fix* network monitoring applications. This agent-to-application integration automates the configuration of the RAP capability, and when coupled with 3Com's distributed RMON agents, delivers self-managed networks that can recognize when and where a problem occurs and dispatch the appropriate analysis tools to that part of the network.

- *LANplex Administration Console.* The LANplex Administration Console is an easy-to-use, menu-driven system that allows each LANplex device to be managed directly by the user. The administration console session can be accessed by a direct serial port connection to the device or through an in-band telnet or rlogin session. Each console session is secured through password protection.

Management Applications Are Still Key

All LANplex systems are managed through 3Com's Transcend Enterprise Manager, which provides features such as configuration management, performance monitoring, and standards-based SNMP management. Transcend Enterprise Manager allows users to manage the LANplex through industry-standard platforms, with a graphical user interface to simplify underlying network complexity.

Available on Windows®, UNIX®, and Novell's NMS™, among other popular platforms, Transcend's graphical user interface allows reduced administrative training and overhead. From this interface, administrators can quickly configure, troubleshoot, expand, and secure enterprise-wide network resources, including network interface cards, hubs, switches, bridges/routers, and remote access servers.

Networking Solutions for Today, Tomorrow, and Beyond

The confluence of advanced desktop functionality, new data-intensive applications, and internet/intranet services has rendered traditional router-based LANs obsolete. With these new modes of sharing knowledge that is mission-critical for virtually every medium- to large-scale organization, network administrators can no longer afford to rely on traditional internetworking solutions to handle increasing demands for bandwidth and management of network growth and capacity.

For firms to gain a competitive edge by bringing products and services to market more quickly, they need high-performing, scalable, and manageable networks that can keep pace with rapidly evolving needs, uses, and technologies. The most effective and economical way to meet this challenge is through the combination of High-Function and Boundary Switching. These new switching technologies offer a combination of fault tolerance that also provides high performance levels and smooth migration for the core of the network and fast, low-cost, easy-to-use solutions for the edge of the network.

3Com's LANplex 2500 and LANplex 6000 High-Function Switches offer a rich feature set that provides performance migration, bandwidth management, and distributed switch management across the full range of network backbone designs. LANplex solutions greatly simplify and strengthen administrative control over LANs based on servers or legacy systems, and deliver flexible performance migration and long-term price/performance benefits. When LANplex High-Function Switches are integrated into 3Com's Transcend Networking framework, which includes Boundary Switches and comprehensive network management tools, administrators can create network infrastructures that fully exploit new applications and technologies now and into the twenty-first century. That, above all, is what businesses require to remain efficient and competitive, and what administrators must provide if they are to satisfy the mission-critical needs of their users and organizations. ◘

3Com Corporation
P.O. Box 58145
5400 Bayfront Plaza
Santa Clara, CA 95052-8145
Phone: 800-NET-3Com
or 408-764-5000
World Wide Web: http://www.3com.com

3Com ANZA
ANZA East
Phone: 61 2 9937 5000
Fax: 61 2 9956 6247

ANZA West
Phone: 61 3 9866 8022
Fax: 61 3 9866 8219

3Com Asia Limited
Beijing
Phone: 8610 8492568
Fax: 8610 8492789

Shanghai
Phone: 86 21 63740220 x6115
Fax: 86 21 63552079

Hong Kong
Phone: 852 2501 1111
Fax: 852 2537 1149

Indonesia
Phone: 6221 523 9181
Fax: 6221 523 9156

Korea
Phone: 82 2 319 4711
Fax: 82 2 319 4710

Malaysia
Phone: 60 3 732 7910
Fax: 60 3 732 7912

Singapore
Phone: 65 538 9368
Fax: 65 538 9369

Taiwan
Phone: 886 2 377 5850
Fax: 886 2 377 5860

Thailand
Phone: 662 231 8151 4
Fax: 662 231 8153

3Com Japan
Phone: 81 3 3345 7251
Fax: 81 3 3345 7261

3Com Benelux B.V.
Belgium
Phone: 32 2 7250202
Fax: 32 2 7201211

Netherlands
Phone: 31 30 6029700
Fax: 31 30 6029777

3Com Canada Inc.
Phone: 416-498-3266
Fax: 416-498-1262

3Com European HQ
Phone: 44 1628 897000
Fax: 44 1628 897041

3Com France
Phone: 33 1 69 86 68 00
Fax: 33 1 69 07 11 54

3Com GmbH (Germany)
Phone: 49 89 627320
Fax: 49 89 62732168

Hungary
Phone: 36 20 394 695
Fax: 36 1 274 4281

Poland
Phone: 0048 22 6451351
Fax: 0048 22 6451352

Czech Republic
Phone: 42 2 21845 800
Fax: 42 2 21845 811

Switzerland
Phone: 0041 31 9961414
Fax: 0041 31 9961410

3Com Ireland
Phone: 353 1 820 7077
Fax: 353 1 820 7107

3Com Latin America
U.S. Headquarters
Phone: 408-764-6075
Fax: 408-764-5730

Argentina
Phone: 541 814 1391
Fax: 541 815 7165

Brazil
Phone: 55 11 546 0869
Fax: 55 11 246 6813

Chile
Phone: 562 633 9242
Fax: 562 633 4338

Columbia
Phone: 571 616 2884
Fax: 571 616 4534

Mexico
Phone: 525 520 7841
Fax: 525 520 7837

Northern Latin America
Phone: 305-261-3266
Fax: 305-261-4901

Venezuela
Phone: 582 953 8122
Fax: 582 953 9686

3Com Mediterraneo
Milano, Italy
Phone: 39 2 253011
Fax: 39 2 27304244

Rome, Italy
Phone: 39 6 5917756
Fax: 39 6 5918969

Spain
Phone: 34 1 3831700
Fax: 34 1 3831703

3Com Middle East
United Arab Emirates
Phone: 971 4 349049
Fax: 971 4 349803

3Com Nordic AB
Phone: 45 39 27 85 00
Fax: 45 39 27 08 44

3Com South Africa
Phone: 27 11 807 4397
Fax: 27 11 803 7405

3Com UK Ltd.
Buckinghamshire
Phone: 44 1628 897000
Fax: 44 1628 897041

Manchester
Phone: 44 161 873 7717
Fax: 44 161 873 8053

Scotland
Phone: 0131 220 8228
Fax: 0131 226 1410

Printed in U.S.A.

500612-001 7/96