



www.archive.org
415.561.6767
415.840-0391 e-fax

Internet Archive
300 Funston Avenue
San Francisco, CA 94118

AFFIDAVIT OF CHRISTOPHER BUTLER

1. I am the Office Manager at the Internet Archive, located in San Francisco, California. I make this declaration of my own personal knowledge.

2. The Internet Archive is a website that provides access to a digital library of Internet sites and other cultural artifacts in digital form. Like a paper library, we provide free access to researchers, historians, scholars, and the general public. The Internet Archive has partnered with and receives support from various institutions, including the Library of Congress.

3. The Internet Archive has created a service known as the Wayback Machine. The Wayback Machine makes it possible to surf more than 450 billion pages stored in the Internet Archive's web archive. Visitors to the Wayback Machine can search archives by URL (i.e., a website address). If archived records for a URL are available, the visitor will be presented with a list of available dates. The visitor may select one of those dates, and then begin surfing on an archived version of the Web. The links on the archived files, when served by the Wayback Machine, point to other archived files (whether HTML pages or images). If a visitor clicks on a link on an archived page, the Wayback Machine will serve the archived file with the closest available date to the page upon which the link appeared and was clicked.

4. The archived data made viewable and browseable by the Wayback Machine is compiled using software programs known as crawlers, which surf the Web and automatically store copies of web files, preserving these files as they exist at the point of time of capture.

5. The Internet Archive assigns a URL on its site to the archived files in the format `http://web.archive.org/web/[Year in yyyy][Month in mm][Day in dd][Time code in hh:mm:ss]/[Archived URL]`. Thus, the Internet Archive URL `http://web.archive.org/web/19970126045828/http://www.archive.org/` would be the URL for the record of the Internet Archive home page HTML file (`http://www.archive.org/`) archived on January 26, 1997 at 4:58 a.m. and 28 seconds (1997/01/26 at 04:58:28). A web browser may be set such that a printout from it will display the URL of a web page in the printout's footer. The date assigned by the Internet Archive applies to the HTML file but not to image files linked therein. Thus images that appear on a page may not have been archived on the same date as the HTML file. Likewise, if a website is designed with "frames," the date assigned by the Internet Archive applies to the frameset as a whole, and not the individual pages within each frame.

6. Attached hereto as Exhibit A are true and accurate copies of printouts of the Internet Archive's records of the HTML files for the URLs and the dates specified in the footer of the printout.

7. I declare under penalty of perjury that the foregoing is true and correct.

DATE: 1/30/17

Christopher Butler

Exhibit A



The 3Com Corebuilder Switch product line has been placed on [End of Life Notice](#).

Scaling Performance and Managing Growth with the CoreBuilder 3500 MultiLayer Switch

As networks become more complex and performance requirements more demanding, organizations need greater ability to manage the evolution of these infrastructures while making them increasingly transparent to users. Network managers require the ability to build flexible, scalable, and easily managed network solutions. Users also want all their applications, however sophisticated, to be instantly available and extremely responsive. With the CoreBuilder 3500 MultiLayer Switch and Transcend® network management software, these goals are not only attainable, but you can build a network that will have the right tools to scale performance and manage growth while providing policy services for the first time.

Evolution of Switched Networks

The demands of today's networks are changing from a model where 80 percent of traffic remained local with the rest going through a backbone to either centralized servers or to some other part of an enterprise network. Today, 50 percent of traffic remains local with the balance going across workgroups or to the backbone. Also, shared media desktop connections are rapidly changing to switched connections to support increased traffic loads. The immediate impact of this evolution is primarily felt in the traditional backbone routers that have been deployed to segment LANs. With so much more traffic crossing IP subnet and IPX network boundaries, these routers become bottlenecks in traffic flows. The reason is that traditional routers were optimized to handle a diverse set of media types and speeds, where translation and buffering ability were of paramount importance, not wire-speed throughput or low latencies. While some recent entries in the router space are much improved in performance--some reaching 1 million pps--they have done so at a very high price point.


The CoreBuilder 3500 MultiLayer Switch--Layer 3 Wire-Speed, Multiprotocol Routing

Today's LAN infrastructures require a new solution to the problem of forwarding traffic across network boundaries. The answer is a new type of networking device. This new device combines the functions of current Layer 2 switches and LAN routers while adding features beyond what legacy LAN routers can support. These new devices have been referred to as Layer 3 switches, but that label does not fully describe the ability of the CoreBuilder 3500 MultiLayer Switch. This product can route packets at wire speeds across 10/100 Mbps Ethernet ports, scaling up to 4 million pps. The CoreBuilder 3500 MultiLayer Switch also has low latencies associated with Layer 2 switches, typically 15 to 30 microseconds. Plus, the CoreBuilder 3500 MultiLayer Switch supports sophisticated Class of Service (CoS) and Quality of Service (QoS) features that allow different types of traffic to have different levels of service--without affecting performance.

Networks are not all based on Ethernet. FDDI and ATM are popular LAN backbone technologies. ATM was once considered the only backbone technology that could be deployed if multimedia applications were to be used. Newer technologies such as Gigabit Ethernet have the potential to become the choice of tomorrow's LAN backbone infrastructures. Some Gigabit Ethernet switches becoming available support CoS and QoS and, combined with massive bandwidth, have the potential to provide service levels adequate for video, voice, and multimedia over Ethernet LANs. While not as robust as ATM QoS, it will be available at lower cost and complexity. Given this reality, any networking device in consideration for use in backbones must provide a way to interconnect these technologies and provide a migration path for those networks that will be changing backbone technologies. This is an area where the CoreBuilder 3500 MultiLayer Switch plays very strongly.

CoreBuilder 3500 Switching Architecture


What makes the CoreBuilder 3500 product the industry-leading multilayer switch is its unique architecture. The switch is built around a third-generation architecture whose key element is the FIRE ASIC. This generation not only builds upon advancements in Layer 2, but provides extensive capabilities and wire-speed performance levels for Layer 3 routing, multicast forwarding, and user-selectable policy services as well. There is no longer a disparity between Layer 2 and Layer 3 performance. True Layer 3 switched topologies can now be built with the control of Layer 3 and the performance of Layer 2 with absolutely no compromises.



The CoreBuilder 3500 switch used as a LAN backbone router.

Flexible Intelligent Routing Engine (FIRE)

3Com has extended the FIRE programmable ASIC capabilities even further by embedding a processor into the silicon. This gives FIRE the performance of an ASIC and makes it extensible as well. Future enhancements, such as IP version 6, can be incorporated without a product upgrade or a sacrifice in performance. This future proofing is a key consideration when evaluating next-generation switching solutions. See the figures on this page and on page 4 for examples of FIRE architecture.

 400347c.gif

Fire Architecture

FIRE introduces Distributed Packet Processing (DPP). With DPP, multiple distributed forwarding engines rapidly and independently transfer packets through the system. As a packet travels through the pipeline, the system:

- Verifies frame integrity
- Captures statistics for relevant MIBs, including RMON
- Determines VLANs
- Distinguishes bridged and routed frames
- Classifies specific data flows
- Polices data rates on flows and reservations
- Applies filters
- Modifies packet headers if routing or tagging
- Applies priorities
- Transmits the packet

Within an individual pipeline, several ASICs simultaneously handle multiple frames. This parallelism and pipelining bring forwarding performance at Layer 2 and Layer 3 to new levels: wire speed on all ports for unicast, multicast, and broadcast traffic.

All buffering is done on the output because all the necessary checking and updates are complete by the time a packet arrives at the output stage. This avoids accessing a packet in memory, one of the most costly operations in terms of performance.

Dynamically Scalable Memory

The performance potential of advanced LAN switches is built on an intelligently designed memory subsystem. With FIRE, a portion of the buffer memory is associated directly with the forwarding engines. As interface modules are added with their respective forwarding engines, memory scales accordingly. Memory is not statically coupled with the forwarding engines, but is made available to all forwarding engines in the system. This physically distributed, yet globally shared memory provides a dramatic increase in performance, especially for multicast forwarding. FIRE allocates buffers to the size of received frames and does so dynamically, so large packets get large buffers and small packets get small buffers. This uses memory efficiently and increases the system's ability to handle large traffic bursts without losing a packet. These dynamically constructed buffers are then allocated in a two-level hierarchy, on a per-port basis and a common pool. This arrangement allows FIRE to guarantee a fixed amount of buffering for each interface yet handle large bursts of data without depleting resources. The result is efficient utilization with outstanding performance.

Advanced Queuing Mechanisms

Traditional LAN switches buffer output traffic on a single queue. Traffic was served on a first-in, first-out manner, with excess traffic discarded when the queue became full. The larger queue allowance added to latency. These characteristics made deploying real-time and multimedia applications very difficult. In response to this, 3Com introduced PACE™ technology, which allows for different classes of service over the same Ethernet LAN with controls for latency and jitter. The FIRE architecture builds on PACE technology and introduces four levels of output queuing for even greater levels of traffic discrimination. Packets associated with real-time flows and multimedia are placed on higher-priority queues. Weighted Fair Queuing (WFQ) services the higher-priority queues more frequently, while not starving the lower-priority queues.



Performance scales with system expansion.

Automatic Flow Classification

FIRE can instruct its packet pipeline to discriminate among user-specified traffic flows. These network manager-defined priority settings are implemented in silicon. This results in decreased latencies, high-priority transmissions, and congestion avoidance. AutoClass instructs the pipeline to classify data flows and assign them a queue-based priority. The process is media independent, working over all Ethernet, FDDI, and ATM media. AutoClass can recognize Data Link Encapsulation such as SNAP types and LLC, as well as protocol types. Unicast, multicast, and broadcast traffic can be distinguished, as well as IP UDP, TCP source, destination addresses, and well-known ports. Mappings of 802.1p and 802.1Q services are provided as well as default classifiers. To simplify the task of traffic classification setup, Transcend network management applications offer simplified configuration of classifiers and queuing.

One example of using AutoClass is prioritizing FTP traffic during large file backups to ensure timely completion. Another is wire-speed firewalls, where certain flows are allocated zero bandwidth.

CoreBuilder 3500 Features

- **VLANs** provide enhanced network performance by isolating broadcast domains. Three basic types of VLANs are supported:
Port-based VLANs are an arbitrary group of ports within a bridge group.
Protocol-based VLANs are a collection of ports designated as a VLAN interface for packets belonging to a specific Layer 3 protocol family.
Network-based VLANs are a collection of ports designated as a VLAN interface for packets belonging to a specified Layer 3 network address.

IEEE 802.1Q tagging is also supported for interoperability with switches in a multivendor environment. 3Com proprietary tagging and Layer 3 address tagging are also supported.

- **Trunking** is a method of aggregating traffic over a group of ports either to form a big pipe between switches or a fault-tolerant pipe to another switch. The trunking group is treated as a single Layer 2 pipe with respect to spanning tree, filtering, VLAN membership, and so on.
- **Traffic control mechanisms** are supported to prevent the propagation of certain traffic. Packet filtering allows the user to define filters that cause blocking of traffic to or from the designated source or destination address or many other criterion. Protocol filtering prevents certain protocols from being processed on particular ports. Route filtering defines the prevention of routes from being advertised on selected interfaces. This type of filtering is mainly a security feature.
- **Multicast packet firewalls** limit the rate at which multicast packets are forwarded.

Internetworking

As large, flat Layer 2 networks are built for performance and simplicity reasons, it is becoming obvious that Layer 3 control is necessary. Users are coming full circle again as they realize the issues of Layer 2-only designs: broadcast storms, spanning tree topologies making inefficient use of data paths, limited dissimilar network connectivity, security, and so on, which forced the use of routers years ago. The reason users tried to reduce the amount of routing in networks was their greater latency and slower performance. Today, the CoreBuilder 3500 MultiLayer Switch is the solution, providing Layer 3 forwarding at the performance levels of Layer 2 switches. Networks can still benefit from Layer 2 switches, but a proper network design must balance Layer 2 with the control of Layer 3 devices. The CoreBuilder 3500 MultiLayer Switch supports the most common Layer 3 protocols in use today:

- **IP** is the protocol most networks are using. Full IP support is incorporated in the CoreBuilder 3500 Layer 3 switch, including TCP, UDP, RIP, OSPF, DVMRP, and all RFCs defined for routing of IP. Both unicast and multicast traffic types are supported. IP version 6 will be supported as well.
- **IPX**, which is used in NetWare environments, is an important protocol as most networks today have some amount of NetWare usage. Although Novell's future direction is toward IP, support of IPX is critical in migrating today's networks toward an IP-only model. The CoreBuilder 3500 MultiLayer Switch supports IPX RIP, SAP, and all IPX Ethernet encapsulation types.
- **AppleTalk** version 2 remains ubiquitous in large networks and is critical in supporting migration to the IP-only direction Apple Computer is taking. The CoreBuilder 3500 MultiLayer Switch supports AppleTalk SNAP, AARP, DDP, RTMP, ZIP, AEP, and NBP.
- **Autocast VLANs** filter IP multicasts based on IGMP and IP multicast routing packets that are received. Multicast traffic will be flooded to ports where IGMP packets have been observed. Multicast packets will also be flooded to DVMRP ports to other multicast routers. This is a benefit in applications that stream

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.