Cisco Express Forwarding (CEF)

Cisco's *Express Forwarding* (CEF) technology for IP is a scalable, distributed, layer 3 switching solution designed to meet the future performance requirements of the Internet and Enterprise networks. It represents the latest advance in Cisco IOS™ switching capabilities that includes *NetFlow™ Switching* and *Distributed Switching.* CEF is also a key component of Cisco's *Tag Switching* architecture.

Express Forwarding evolved to best accommodate the changing network dynamics and traffic characteristics resulting from increasing numbers of short duration flows typically associated with Web-based applications and interactive type sessions. Existing layer 3 switching paradigms use a route-cache model to maintain a fast lookup table for destination network prefixes. The route-cache entries are traffic-driven in that the first packet to a new destination is routed via routing table information and as part of that forwarding operation, a route-cache entry for that destination is then added. This allows subsequent packets flows to that same destination network to be switched based on an efficient route-cache match. These entries are periodically aged out to keep the route cache current and can be immediately invalidated if the network topology changes. This 'demand-caching' scheme — maintaining a very fast access subset of the routing topology information — is optimized for scenarios whereby the majority of traffic flows are associated with a subset of destinations. However, given that traffic profiles at the core of the Internet (and potentially within some large Enterprise networks) are no longer resembling this model, a new switching paradigm was required that would eliminate the increasing cache maintenance resulting from growing numbers of topologically dispersed destinations and dynamic network changes.

CEF avoids the potential overhead of continuous cache churn by instead using a Forwarding Information Base (FIB) for the destination switching decision which mirrors the entire contents of the IP routing table. i.e. there is a one-to-one correspondence between FIB table entries and routing table prefixes; therefore no need to maintain a route-cache.

This offers significant benefits in terms of performance, scalability, network resilience and functionality, particularly in large complex networks with dynamic traffic patterns.

CEF Performance

CEF implements a Cisco patent-pending expedited IP look-up and forwarding algorithm to deliver maximum layer 3 switching performance. Additionally *Express Forwarding* is less CPU intensive than route-caching therefore it allows more CPU horsepower to be dedicated to packet forwarding.

CEF Scalability - *dCEF*

Cisco's *Express Forwarding* technology is optimized for information distribution allowing it to take advantage of the distributed architecture of the high end Cisco IOS(tm) routers such as the C7500. Thus Distributed CEF *(dCEF)* delivers scalable switching capacity by providing each of the C7500 Versatile Interface Processors (VIPs) with an identical on-card copy of the FIB database enabling them to autonomously perform *Express Forwarding* and therefore significantly increase aggregate throughput. CEF also uses Adjacency Tables to hold the layer 2 next hop addresses for all FIB entries so that the associated prepend can be added locally, thus minimizing latency, before switching between linecards. Therefore in the case of *dCEF* on the C7500 platform, the Route Switch Processor (RSP) is relieved of any switching operation and so has significantly more CPU power available to perform routing functions, management, network services etc.

C7500 Distributed Architecture

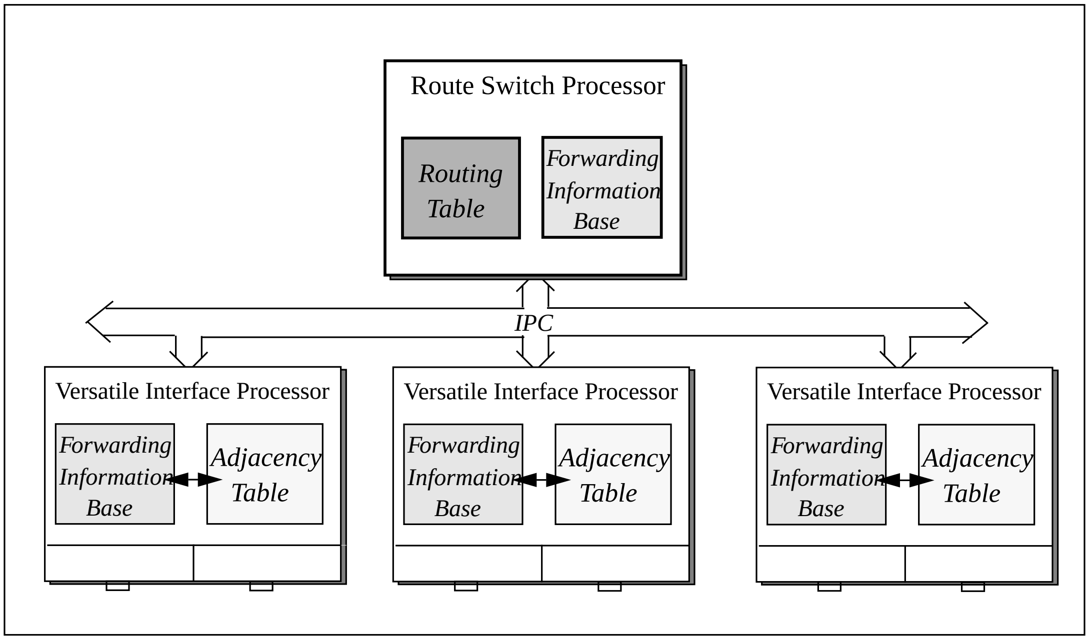

dCEF employs a reliable Inter Process Communications (IPC) mechanism to ensure guaranteed and synchronized FIB information and because of the correlation between FIB entries and the routing table, updates are only sent to the VIP linecards when changes in the routing topology occur.

Cisco's upcoming *Gigabit Switch Router (GSR)* will also employ *dCEF* to deliver full topology forwarding at performance rates of millions of packets per second per line card.

CEF Resilience

Express Forwarding offers an unprecedented level of switching consistency and stability in large, dynamic networks.

This results from the fact that the FIB look-up table contains all known routes therefore eliminating the potential for "cache-misses" which occur with demand caching schemes. i.e. if a route is not found in the forwarding cache, the first packet(s) then looks to the routing table to determine the outbound interface and then a cache entry is added for that destination. Since the cache information is derived from the routing table, routing changes cause existing cache entries to be invalidated and then reestablished to reflect any topology changes. In networking environments which frequently experience significant routing activity (such as the Internet backbone) this can cause traffic to be forwarded via the routing table (which represents Cisco IOS(tm) *Process Level Switching*) as opposed to via the route-cache (Cisco IOS(tm) *Fast Switching*). During major network convergence or flux, performance can thus be sub-optimal.

Express Forwarding obviates this *Process Switching/ Fast Switching* scenario and since the FIB is topology-driven rather than traffic-driven, CEF's switching performance is largely independent of and unaffected by network size or dynamics.

CEF Functionality

Express Forwarding is feature-rich in that it can be used in conjunction with leading-edge Cisco IOS(tm) capabilities. Even when deployed in *dCEF* mode, there is no compromise between performance and functionality, hence backbone networks can realize the benefits of *Express Forwarding* while still having the flexibility to offer Cisco IOS(tm) network services.

These capabilities include:

- Quality of Service – IP Precedence based Class of Service and Weighted Random Early Detection (wRED) support.
- Load Balancing – Per destination (the default) and per packet over equal/unequal cost links for as many paths as known in the routing topology
- Traffic Statistics – Byte and packet counts at a granularity of per-prefix, per-neighbour etc.
- Media Independence – CEF currently supports Packet over Sonet, ATM/AAL5, Frame Relay, Ethernet, FDDI, HDLC and mPPP
- Tunneling – Generic Route Encapsulation (GRE)
- Cisco IOS(tm) subinterface support – allowing for the flexibility of per subinterface configurations e.g. MTU

Express Forwarding and Netflow(tm) Switching

NetFlow Switching manages traffic in terms of flows between network layer source-destination address pairs and unlike conventional destination-prefix based forwarding, it also provides visibility to the application layer. By caching this level of information on a per-session basis, *NetFlow Switching* is designed to support complex packet filtering security without incurring any per-packet overhead. i.e. only the initial packet is passed through the Access Control Lists (ACLs) and if allowed, used to build a NetFlow entry containing address, port and application information such that subsequent packets belonging to that flow match this entry and therefore do not need to also be compared against the ACLs. *NetFlow Switching's* traffic profiling capabilities, when used in conjunction with *NetFlow Data Export*, also facilitate sophisticated network billing and accounting applications. Traffic flows are now invariably short-lived and to maximize performance entries are removed as soon as the particular flow finishes. In environments with very many concurrent flows, this can mean significant NetFlow demand-cache maintenance, with the initial packet from new flows used in establishing the entries having to be *Process Switched*. Now, by using the Forwarding Information Base for flow creation, *NetFlow Switching* avoids this overhead.

When deployed along with CEF, *NetFlow Switching* is typically used at the periphery of the Layer 3 backbone network and *Express Forwarding* used in the core since it is most useful and efficient to perform any security, QoS or accounting operations early in the data path and at the network boundary as traffic is forwarded to and from the backbone:

Express Forwarding and *Tag Switching*

Tag Switching combines the price/performance and traffic management capabilities of Layer 2 cell or frame switching with the scalability and functionality of Layer 3 routing. This multilayer integration is achieved by having routers and switches interconnect as peer networking devices, an architecture that scales well beyond existing models. Essentially the technology associates a small tag value with destination prefixes such that packets going to a particular destination now have a tag prepended as they are switched across the network. This therefore allows *Tag Switching* devices to perform very efficient and fast forwarding decisions (using a technique known as 'label-swapping') and so deliver the switching speeds that will be required with future line rates.

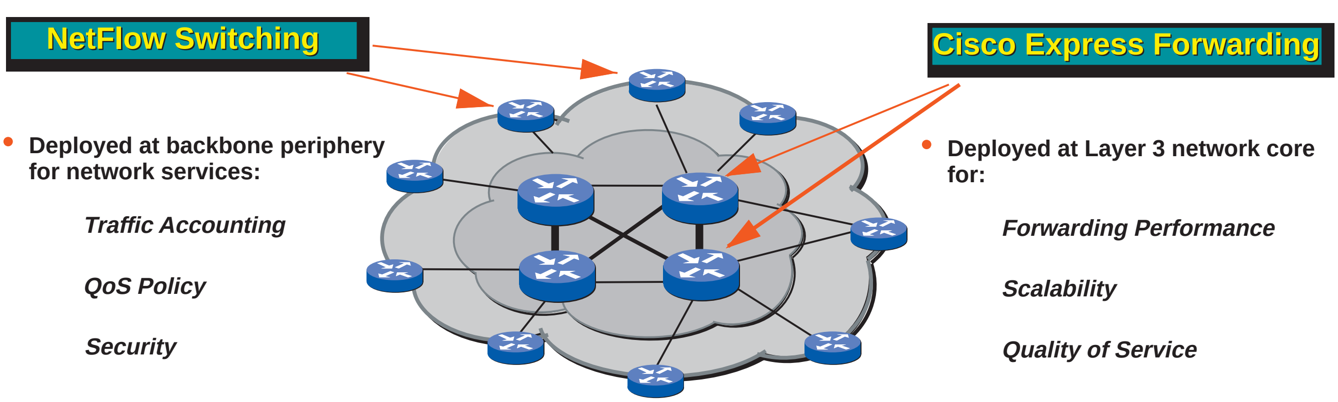

Cisco's implementation of tag switching on the Cisco 7500 leverages the advantages of *Express Forwarding* and both can be used together. Consider a network consisting of Cisco 7500 routers at the edge and either Cisco 7500 or Cisco Gigabit Switch Routers in the backbone. One option is to enable *dCEF* for IP traffic, and utilize tag switching for its traffic engineering capabilities. In this scenario, the edge routers will use *Express Forwarding* to decide how to route the packets, and if the packet is to travel over a tag switching traffic engineered tunnel, they will apply tags when forwarding packets. The FIB used by *Express Forwarding* interfaces to Tag Switching's Tag Information Base (TIB), thus enabling *Express Forwarding* to apply tags for appropriate interfaces. A second scenario is to switch all traffic via tag switching, possibly to take advantage of the benefits of hierarchical tagging (e.g., isolation of exterior routes from interior routers). In this case the edge routers use *Express Forwarding* as above to determine how to route the packet and how to tag it. The backbone tag switching routers leverage *Express Forwarding* by switching tagged packets based on the TIB, which is distributed to the VIP linecards just like the FIB tables are distributed with *Express Forwarding*, providing the same benefits of scalable performance. A third case involves ATM switches as the backbone, with routers on the periphery. In this scenario, the ATM switches and routers would support tag switching functions. The edge routers would, as above, use *Express Forwarding* to forward packets coming from and going to non-tag intelligent devices, and would apply tags when forwarding to the ATM tag switches.

Considerations

1. *Express Forwarding* is available as a regular Cisco IOS(tm) upgrade initially in software release 11.1CA for the Cisco C7500 platform with support on the C7200 to follow.
2. The recommended minimum memory requirements for CEF when deployed in platforms carrying full Internet routing information is 64 Meg for the centralized RSP and 32 Meg per VIP linecard in distributed mode.
3. *Express Forwarding* can co-exist with other Cisco IOS switching modes, except in the case of VIP-distributed Fast Switching which should not be used simultaneously with dCEF.
4. Features or encapsulations currently not supported by CEF will default to the other switching modes, such as *Optimum Switching* or *Fast Switching*, as configured.

Conclusion

Cisco's *Express Forwarding* is an advanced layer 3 switching technology that offers the performance and scalability that will be required by the Internet backbone and expanding Enterprise networks.

CISCO SYSTEMS

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
World Wide Web URL:
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic-Batiment L1/L2
16, Avenue du Quebec
BP 706-Villebon
91961 Courtaboeuf Cedex
France
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Americas
Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
Tel: 408 526-7660
Fax: 408 526-4646

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
Tel: 81 3 5219 6000
Fax: 81 3 5219 6010

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the

Cisco Connection Online Web site at http://www.cisco.com.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark
Finland • France • Germany • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Philippines • Poland • Portugal • Russia • Singapore • South Africa • Spain • Sweden
Switzerland • Taiwan, ROC • Thailand • United Arab Emirates • United Kingdom • Venezuela