Restrictions

These sections list restrictions for Cisco IOS for the Catalyst 6500 series switches:

Restrictions Removed by the PFC3

The PFC3 removes these restrictions that were present with other policy feature cards:

- You can configure features to use up to 3 different flow masks.
- You can configure more than 1 Gateway Load Balancing Protocol (GLBP) group.
- You can configure up to 255 unique HSRP group numbers.
- You can configure a separate MAC address on each interface.
- You can configure Unicast RPF check without reducing the number of available CEF entries.
- You can configure VLAN-based QoS with DFC3s installed.
- You can configure port-based and VLAN-based QoS on a per-port basis on the WS-X6548-RJ-45 and WS-X6548-RJ-21 switching modules.
- You can configure QoS policy maps attached to an EtherChannel formed from interfaces on different DFC-equipped switching modules.

General Limitations and Restrictions

This section describes general limitations and restrictions:

Identifier	Technology	Description
CSCeb06765	—	Request for Global command to enable Syslog
CSCec27709	—	HSRP/VRRP/GLBP switchover delays with switchports on DFC3A
CSCec79733	—	VPN-SM-C2:Ingress traffic on Intf vlan cannot be spanned
CSCed76016	—	traceback message when boot up from sup720
CSCef12458	—	SNMP topology discovery caused Cat6000 high CPU utilization
CSCef19599	—	Service policy marking inbound not working on SUP720
CSCek39668	—	Need to port the fix of CSCsb58066 to cat6000/c7600 platform codes
CSCsb59015	—	L3 packet duplication with vacl capture enabled and WS-X6708 present
CSCse07679	—	Multicast with L2 broadcast address is routed w/o rewriting the mac.
CSCse86399	—	PBR forwards local destined packets which match ACL criteria
CSCsl86574	—	Some commands are not getting converted when converting vs to standalone
CSCsm43234	—	%ISSU_PROCESS-SP-3-SYSTEM: Not enough space in NVRAM;Fail to set bootvar
CSCsw67926	—	C2W2: CFM Draft1 does not support in Whitney2.0.
CSCsx11253	—	Traceback@ip_adm_ha_cf_get_msg_buf during reauthentication with webauth

Identifier	Technology	Description
CSCsy31055	—	Support for patching modular IOS images deferred
CSCtc10254	—	IPv6 raguad fail to drop RA pkt with hop-by-hob and auth ext hdr
CSCte31168	—	CFM:W2.C-->SXI1 with cfm configs, 2MR LCs are powered down
CSCtf08313	—	Authentication doesnt start with traffic on link up.(Multi-auth)
CSCtf18029	—	MPLS ENF-ISSU : Internal VLAN not removed on downgrade
CSCtf32483	—	CSM not loadbalancing when a real is removed/added dynamically
CSCtf41722	—	W2.Clix:Multicast Service Reflect: Translation Entries missing in DFCs
CSCtf48652	—	W2.Clix: service reflection with short masks requires rate-limit tweak
CSCtf48659	—	W.Clix: missing OIF in rmcast Serv Refl with BIDIR in priv network
CSCtf51285	—	W2.Clix: after switchover, some groups miss ServRefl. vlans in OIF-list
CSCtg14082	—	W2.Clix: shutdown of VSL link causes 80 seconds mcast disruption
CSCtg17528	—	c2w2c: fast redirect broken on secondary portchannel id's (still)
CSCtg33528	—	10G SFP+: Some fields in the o/p of 'sh idprom interface' shows unspecif
CSCtg61862	—	Vacl capture stops on reloading legacy card.
CSCtg89884	—	"%FABRIC-SP-6-TIMEOUT_ERR" upon ISSU SXI4-->SXI2a, powering up the LCs
CSCth03423	—	Standby Sup crashes due to bulk sync failure, Issue same as CSCsx46323
CSCti98129	—	rcv queue bandwidth does not programmed in sh qm port for R2D2 rev2
CSCtj16159	—	standby reboots twice and comes up in rpr due to config sync fail
CSCtj27150	—	Module 1. LCC Client UNSOLICITED SCP failed
CSCtj41404	—	Tracebacks at pm_port_want_to_bundle on shut/noshut on range of mlacp po
CSCtj41777	—	multicast traffic drop and duplicate - DR reload
CSCtj46469	—	C2WA1:Ports succesfully budled in sec agg flap on modifying Po VLAN mask
CSCtj49906	—	Marking does not work with CoPP policy
CSCtk76193	—	ES40 QoS : No more than 7 classes accepted on a policy-map (Flat/HQoS)
CSCtl03171	—	c2wa1:Periodic High CPU due to SNMP ENGINE process with NAM in VSS setup
CSCtl06114	—	SGT is not assigned to webauth client even after receiving from ACS
CSCtl69556	—	block qos config on portchannel with same asic as vsl on load/issu
CSCtn45851	—	Fail to achieve more then 5gig through put on 2VTI Tunnel in VSS
CSCtn77625	—	ES40 QoS : cos2exp service-policy command missing in run-confg on reload
CSCto66065	—	NAM3: hw-module shutdown cmd power down NAM3
CSCto85304	—	c2wa1b: ERSPAN monitoring fails when routed src port is on CFC on VSS
CSCtq23239	—	ES40 Interface not coming UP on reload of VSS
CSCtq90786	—	c2wa1b:RG doesn't re-calc sys-id on changing mlacp sys-prio on a PoA
CSCtq95747	—	NAM3:dataport 2 not rcv traffic if same span destined to both dataports
CSCtt16904	—	c2wa1c: memory leak when bennu and nam are reset together.
CSCtt43196	—	L3 Multicast stream over VPLS affected by Link down
CSCtu26761	—	c2wa1c: Bennu states toggles due to autostate down on VSS

Identifier	Technology	Description
CSCtu34830	—	VPLS_TE_FRR_Multicast traffic is affected
CSCtw53269	—	SXJ2:standby reloads twice @%SATVS_IBC-SW1_SP-5-VSL_DOWN_SCP_DROP
CSCty26003	—	LTL index missing post sso+rcfg (no functionality or traffic impact)
CSCub47529	—	Enhanced FlexWAN (WS-X6582-2PA) silently reload during BERT w/o crashinf
CSCdx10583	Cisco IOS	EoMPLS:CWPA: per-vlan shaping does not work when egress is FlexWan
CSCdx39882	Cisco IOS	hubble: SNMP traps delayed 3-10 minutes
CSCea53554	Cisco IOS	const2: C2MSFC3: PBR continues to fwd traffic when RPF check fails
CSCea70203	Cisco IOS	eigrp dflt-network stays prg in h/w when route is removed
CSCeb01860	Cisco IOS	Const2:egress qos policy applied to ingressing rpf failure pkts
CSCec02266	Cisco IOS	Berytos perf drop if sw reboot after NVRAM is cleared
CSCed29392	Cisco IOS	MPLS: Egress ACL applied on tagged traffic
CSCee39455	Cisco IOS	Sporadic packet drops on 6704 linecard under IPv6 Unicast Bidir flows
CSCef08631	Cisco IOS	MVPN: multicast entries in globle table leak into VRF table
CSCef78235	Cisco IOS	Disable egress span of vacl redirected packets
CSCeg29898	Cisco IOS	MMLS: rpdf-cache not updated when using BSR. Bidir flows SW switched
CSCsb25509	Cisco IOS	eibgp load balancing interwork with tag to ip recirculation cause loop
CSCtc91284	Infrastructure	W2C: polling flashMIB shows wrong info with snmp mib flash cache enabled
CSCtc06369	IPServices	NAT source address translation not done for PIM register packets
CSCtd72233	IPServices	Able to disable IPv6 with HSRPv6 configured and the state is active
CSCtl73810	IPServices	c2wa1:ISSU from SXI4a/SXI4 to later image fails due to RF progression
CSCsg08736	LegacyProtocols	@SyncMutexLock_r - Blocked remote registry call from blob to ION pro
CSCee77417	MPLS	TE/FRR:FRR broken when LDP NOT enabled
CSCds22281	platform-cat6xxx	ip multicast rate-limit command inadvertently removed from config
CSCdt72147	QoS	microflow policing not working with same flow on multiple input port
CSCtf14368	Routing	sh runn VRF does not show EIGRP vrf config
CSCtg60341	Routing	IPv6 stalled interface is sending HSRP-Hellos
CSCtk63987	Routing	c2wa1: crash in pdb_flushcache during unconfig
CSCec29255	WAN	Flexwan OIR causes WS-6516-GBIC (no DFC) to reload

- When a redundant supervisor engine is in standby mode, the Ethernet ports on the redundant supervisor engine are always active.
- All Ethernet LAN ports on all modules, including those on a redundant supervisor engine, support EtherChannel (maximum of eight interfaces) with no requirement that the ports be contiguous.
- All Ethernet ports on all modules support 802.1Q VLAN trunking.
- When you add a member port that does not support ISL trunking to an EtherChannel, Cisco IOS software automatically adds a **switchport trunk encapsulation dot1q** command to the port-channel interface to prevent configuration of the EtherChannel as an ISL trunk. The **switchport trunk encapsulation dot1q** command is inactive when the EtherChannel is not a trunk.

- A distributed EtherChannel (DEC) is an EtherChannel with ports on more than one DFC-equipped module or, on a DFC-equipped dual-fabric connection module, with ports that use different fabric connections. (Search for "Dual switch-fabric connections" in this document.)
- To reduce CPU utilization during ACL configuration changes, use named ACLs instead of numbered ACLs whenever possible, because the ACL merge algorithm runs each time you change an ACE in a numbered ACL. With named ACLs, the ACL merge algorithm runs only when you exit the named ACL configuration mode.
- In releases where caveat CSCef78235 is resolved, with any Supervisor Engine 720 hardware revision, local SPAN and RSPAN source ports do not copy VACL-redirected traffic.

  In releases where caveat CSCef78235 is not resolved:

  - With WS-SUP720, hardware revision 3.2 or higher, local SPAN source ports do not copy VACL-redirected traffic.
  - With WS-SUP720 hardware revisions lower than 3.2, local SPAN source ports copy VACL-redirected traffic.
  - With any Supervisor Engine 720 hardware revision, RSPAN source ports copy VACL-redirected traffic.

  Enter the **show module version | include WS-SUP720-BASE** command to display the hardware revision. For example:

```
Router# show module version | include WS-SUP720-BASE
 7    2  WS-SUP720-BASE     SAD075301SZ Hw :3.2
```

- IPsec in software on the MSFC is supported only for administrative connections to Catalyst 6500 series switches and Cisco 7600 series routers.
- PFC QoS does not rewrite the payload ToS byte in tunnel traffic.
- PFC QoS does not rewrite the ToS byte in bridged multicast traffic.
- The PFC3 does not apply egress policing to traffic that is being bridged to the MSFC3.
- The PFC3 does not apply egress policing or egress DSCP mutation to multicast traffic from the MSFC3.
- The MSFC3 supports tunnels configured with egress features on the tunnel interface. Examples of egress features are output Cisco IOS ACLs, NAT (for inside to outside translation), TCP intercept, context-based access control (CBAC), and encryption.
- The PFC3A does not support any PFC QoS features on tunnel interfaces.
- The PFC3BXL does not provide hardware switching for ICMP traffic if you configure NAT.
- The PFC3A does not provide hardware switching for ICMP traffic if you configure NAT or Cisco IOS reflexive ACLs.
- If you have a network device in your network with MAC address reduction enabled, you should also enable MAC address reduction on all other Layer-2 connected network devices to avoid undesirable root bridge election and spanning tree topology issues.

  When MAC address reduction is enabled, the root bridge priority becomes a multiple of 4096 plus the VLAN ID. With MAC address reduction enabled, a switch bridge ID (used by the spanning-tree algorithm to determine the identity of the root bridge, the lowest being preferred) can only be specified as a multiple of 4096. Only the following values are possible: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440.

  If another bridge in the same spanning-tree domain does not run the MAC address reduction feature, it could win root bridge ownership because of the finer granularity in the selection of its bridge ID.

- Cisco IOS software Release 12.2SX does not support:
  - Integrated routing and bridging (IRB)
  - Concurrent routing and bridging (CRB)
  - Remote source-route bridging (RSRB)
- Use bridge groups on VLAN interfaces, sometimes called fall-back bridging, to bridge nonrouted protocols. Bridge groups on VLAN interfaces are supported in software on the MSFC.
- Cisco IOS software Release 12.2SX does not support the IEEE bridging protocol for bridge groups. Configure bridge groups to use the VLAN-bridge or the DEC spanning-tree protocol.
- FlexWAN module interfaces support dNBAR. Do not configure NBAR or dNBAR on other interfaces.
- Ingress IP Packets with TTL=1 that are not addressed to the MSFC and that match QoS filtering parameters might cause overpolicing of other ingress traffic on the same ingress interface.
- When the outgoing interface list for group G traffic transitions to null on a last-hop multicast router, the router sends a (*,G) prune message to the PIM neighbor toward the rendezvous point (RP) to stop the flow of group G traffic (if any) down the shared tree, but does not send an (S,G) prune message to stop the flow of traffic down the shortest path tree (SPT). The transition of the outgoing interface list to null does not trigger an (S,G) prune message. (S,G) prune messages are triggered by the arrival of (S,G) traffic.

  If the last-hop multicast router is a Catalyst 6500 series switch, traffic is forwarded in hardware. In most cases, RPF-MFD is installed for the (S,G) entries. The MSFC does not see the multicast traffic flowing down the SPT and does not send any traffic-triggered (S,G) prunes to stop the flow of traffic down the SPT. This situation does not have any adverse effect on the MSFC because the PFC processes and drops the unwanted (S,G) traffic.
- Cisco IOS software Release 12.2SX does not support network booting.
- The IP HTTP server feature is disabled by default. Enter the **ip http server** command to use the feature.
- For LAN switching modules, the Cisco IOS **show controllers** command generates no output in Cisco IOS software Release 12.2SX. Enter the **show module** command instead.
- To avoid the case where all traffic is out of profile, the burst size specified in a QoS policing rule must be at least as large as the maximum packet size permissible in the traffic to which the rule is applied.
- By default, the MSFC sends Internet Control Message Protocol (ICMP) unreachable messages when a packet is denied by an access group.

  With the **ip unreachables** command enabled (which is the default), the supervisor engine drops most of the denied packets in hardware and sends only a small number of packets (10 packets per second, maximum) to the MSFC to be dropped, which generates ICMP-unreachable messages.

  To eliminate the load imposed on the MSFC CPU by the task of dropping denied packets and generating ICMP-unreachable messages, you can enter the **no ip unreachables** interface configuration command to disable ICMP unreachable messages, which allows all access-group denied packets to be dropped in hardware.
- MAC address-based Cisco IOS ACLs are not supported for packets that are Layer 3 switched in hardware. MAC address-based Cisco IOS ACLs will be applied on software-switched packets.
- If you enable multicast routing globally, then you should also enable multicast routing (using the **ip pim** command) on all Layer 3 interfaces on which you anticipate receiving IP multicast traffic. This command causes the packets to be sent to the process switching level to create the route entry. If you disable multicast routing on the RPF interface, the entry cannot be created and the packet is dropped.

If the source traffic rate exceeds what can be handled by the process level, it can have an undesirable impact on the system. For example, routing protocol packets, such as EIGRP hello packets, might get dropped.

- The in and out ports displayed in Layer 3 table entries are set by the hardware at the time the entry is created. They are not guaranteed to be accurate in case multiple flows use the same entry (for example, if the flow mask is **Dest-only** and some kind of load sharing is active) or if the source or destination of the Layer 3 entry moves in the Layer 2 topology. The port information is not always available when the Layer 3 entry is established. This is the case if the destination port of the rewritten packet is unknown when the shortcut is created.
- For EtherChannels, you can configure the QoS trust state and default CoS directly on the EtherChannel interface with the **mls qos trust** or **mls qos cos** commands, respectively. These two parameters must be the same for all physical interfaces in the channel. No other QoS queueing configuration commands can be applied to EtherChannel interfaces. Other QoS queueing configuration commands can be applied, however, to individual EtherChannel physical interfaces. After the physical interfaces are bundled into an EtherChannel, QoS classification, marking, and policing by the Policy Feature Card (PFC) for the channel packets is determined by the service-policy attached to the EtherChannel interface. The service policies attached to the individual physical interfaces of the EtherChannel do not matter. The same is true for the port-based and VLAN-based QoS state of the EtherChannel interface. You can disable the PFC QoS features using the **no mls qos** interface configuration command on the EtherChannel interface.
- The maximum recommended number of Layer 3 multicast entries is 10,000. The maximum recommended number of multicast entries supported in the Layer 2 forwarding table is 12,000.
- After enabling Protocol Independent Multicast (PIM) on an interface, you need to enter the **ip mroute-cache** command on the interface to enable multicast fast-switching. If you have “no ip mroute-cache” configured, multicast packets that are not hardware switched will go to the process level that increases the load on the router.
- FlexWAN ports do not support SPAN or RSPAN.
- MPLS on the FlexWAN module does not support Virtual Private LAN Service (VPLS).