Cisco IOS Early Deployment Release 12.0(21)SZ

Introduction

This Product Bulletin describes the content and delivery information concerning Cisco IOS™ software release 12.0(21)SZ. It should be used in conjunction with Product Bulletin titles, Cisco IOS Software Release 12.0(21)SZ. 12.0(21)SZ is a short-lived Early Deployment (ED) release parented to 12.0(21)S and supporting Cisco 12000 platform. Support for 12.0(21)SZ will be discontinued three months after the 12.0(23)S release appears on CCO.

Migration Guide

Diagram 1 below displays Cisco 12.0(21)SZ release functionality relative to the 12.0S release. This diagram also identifies the recommended migration path.

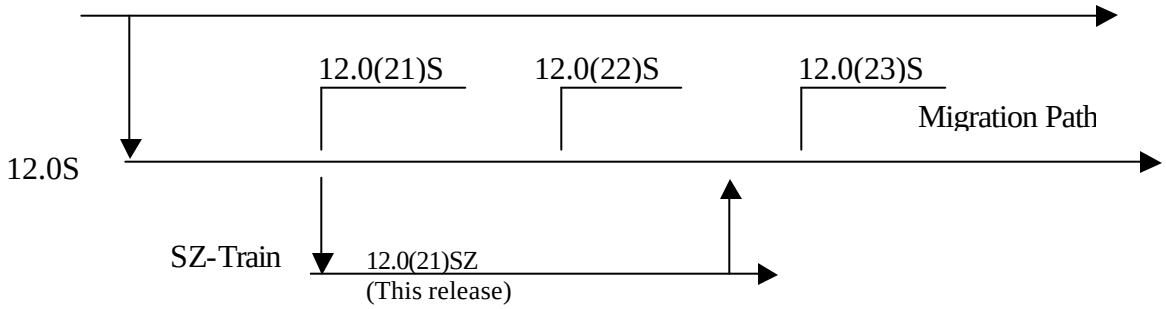

New Features in **12.0(21)SZ**

New features in 12.0(21)SZ include software support for the following new line cards:

- Modular Gigabit Ethernet Card (1 Port base card with modularity to 10 ports)
- 1-PORT 10GE Line Card
- 8-Port OC3 STM-1 ATM Line Card

New Product Support	Cisco 12000
Modular Gigabit Ethernet Line Card • Basic MPLS switching with LDP • MPLS VPN PE • MPLS Switching over 802.1q VLANs • MPLS Specific QoS • Source/Destination MAC accounting • BGP Policy Accounting • Sampled Netflow • WRED/MDRR • Input Security Extended ACLs • Output Security Extended ACLs • CAR • Output Traffic Shaping • VLAN CoS Bit Copying • HSRP/VRRP • Policy Based Routing	X
1-Port 10-Gigabit Ethernet Line Card • Basic MPLS switching with LDP • MPLS VPN PE • MPLS Switching over 802.1q VLANs • MPLS Specific QoS • Source/Destination MAC accounting • BGP Policy Accounting • Sampled Netflow • WRED/MDRR • Input Security Extended ACLs • Output Security Extended ACLs • CAR • Output Traffic Shaping • VLAN CoS Bit Copying • HSRP/VRRP • Policy Based Routing	X
8-Port OC3 STM-1 ATM Line Card • Multiprotocol encapsulation over AAL5 (RFC 1483) with support for LLC/SNAP and IP Mux • UNI 3.0/3.1 (including LMI) • ATM Service Classes (VBR-nrt, UBR with optional PCR) • Per-VC and per-VP traffic shaping • Configurable Queue Depth • Per-VC WRED • Per-VC MDRR (including LLQ) • F4 & F5 OAM, AIS/RDI alarm, alarms and loopback cell • Basic MPLS switching with LDP • MPLS VPN PE • MPLS CoS (Per-VC queuing with EXP bit support with WRED/MDRR) • Modular QoS command line interface (MQC)	X

Table 1. Cisco IOS Release 12.0(21)SZ Features

Basic MPLS switching with LDP

- Unicast switching, with specific support for (i) up to 6 label push operations, (ii) 1 label pop operation (2 label pop operation in case of Explicit Null), or (iii) 1 label swap with up to 5 label push operations, at each MPLS switch node.
- Packet fragmentation supported, where MPLS packets larger than the configured MTU size are fragmented and forwarded by line card CPU. While length of IP packet and Layer 2 header are both included in the comparison against MTU, length of appended MPLS labels is not included.
- Time-To-Live (TTL) Propagate: Support for both (i) "Propagate" mode, where the IP TTL value is copied to the MPLS TTL during label imposition, or (ii) "No-Propagate" mode, where a constant value of 255 is copied. In both cases, the TTL value is decremented at each MPLS node. For a pop2label operation, MPLS TTL value is propagated down to the next label. For pop2ip operation, MPLS TTL value is propagated to the IP TTL if it is less than the current IP TTL.
- Support for Explicit Null label to preserve CoS information when forwarding packets from P to PE routers.
- Support for Implicit Null label to request that penultimate hop router forward IP packets without labels to the router at the end of the label switched path (LSP).
- Number of LSPs supported is complicated function of line card memory structure and route processor resource requirements. Cisco IOS team will have tested up to 100k labels per system by FCS.

MPLS VPN PE

- Number of MPLS VPN VRF tables is limited to 512 in software. The maximum number needed could be 1 per VLAN in some applications, or potentially larger if, for each VLAN, Policy Based Routing (PBR) is used to select a different VRF based on parameters such as IP source address.
- Number of MPLS VPN routes limited to 200,000 per line card.
- 802.1q VLAN to VPN mapping, where the VLAN ID can be used to select the appropriate VRF table (i.e. forwarding decision made based on [VLAN ID, IP destination address] pairs). Simultaneous support on the same physical interface for VLAN to VPN mapping, VLAN to EoMPLS LSP/TE tunnel mapping, and basic IP forwarding.
- Support for aggregate and non-aggregate labels. PE routers use aggregate labels in an effort to conserve MPLS label usage. An aggregate label effectively represents an aggregate or summary of VPN routes to individual hosts or small subnets, and avoids the need for a larger number of individual VPN routes to/from each of these destinations/sources. The modular GE card and 1x10GE line card support both imposition and disposition of aggregate labels.
- Support for these VRF aware commands: HSRP, DHCP, PING. Support is enabled on a per VLAN basis given The modular GE card/1x10GE line card's support for VLAN to VPN mapping. As an example of this feature, by supporting VRF-aware HSRP, two Modular GE card interfaces configured in the same HSRP group can both assign a given VRF table to a virtual HSRP MAC address, and support fail-over from the working to protect interface if necessary. No support is available for these VRF aware commands: TFTP, NAT, WCCP.

MPLS Switching over 802.1q VLANs

- Unusual case here, given that ISPs typically run IP within VLANs before traffic from an end customer reaches an edge router (where VLANs are mapped to MPLS VPNs or EoMPLS LSPs). However, this feature is useful when a Layer 2 switch sits between edge and core routers that both run MPLS, and the service provider wants to run VLANs between these two routers for purpose of traffic segmentation (e.g. traffic on one VLAN goes to core router A, traffic on another VLAN goes to core router B, etc.) This architecture reduces the number of core router interfaces required for some POP topologies.

MPLS Specific QoS

- Support E-LSP mode, in which Exp bits determines the appropriate queuing services
- Support for MDRR and WRED
- Support for input and output CAR, rate limiting and packet coloring
- Bit Copying:
    - Copying IP precedence to/from MPLS Exp bits in support of "E-LSP Mode" (i.e. MPLS CoS translation). No support for creating an arbitrary mapping between IP precedence and MPLS Exp bits.
    - Copying 802.1p bits to MPLS Exp bits in support of "E-LSP Mode". No support for creating an arbitrary mapping between IP precedence and MPLS Exp bits.
    - No support exists for copying MPLS Exp bits to 802.1p bits. Although hardware is capable of supporting this feature, a software implementation would decrease by factor of 8 the number of supported adjacency entries.
- MPLS QoS transparency is supported, in which MPLS Exp bits may be modified but IP precedence bits are not modified.

Source/Destination MAC accounting

- Byte or packet counting per port for each source MAC address. Whether the user selects byte or packet counting, the selection must be the same as for VLAN Accounting below.
- 512 counters per modular 1xGE or 10GE port.
- Counting number of bytes and packets transmitted to each destination MAC address.
- At least 512 counters per port; adjacency memory provides relevant memory, so no practical limit on number of entries. The number of counters is independent of the source MAC accounting and VLAN counters.

BGP Policy Accounting

- Supported on a per physical interface basis, not on a per VLAN basis.
- This feature enables assignment of up to eight (8) 64-bit packet and byte counters per line card for matches on destination Autonomous System (AS) numbers.
- Examples of this feature include counting packets or bytes that (i) are destined to a given AS number, (ii) pass through a given AS number, or (iii) follow a path through a specific series of AS numbers.
- The need for BGP-PA on GSR Ethernet line cards has become increasingly important as these line cards are being used to a greater extent for peering and customer access.

Sampled Netflow

- Support for v5 and v8. Please reference this document for further information on this feature: http://wwwin.cisco.com/irg/ipsbu/sales/faqs/netflow_gsr_faq.doc

WRED/MDRR

- Supported per physical port, not per VLAN.
- Both Random Early Detection (RED) and Weighted RED (WRED) are supported as schemes to mitigate congestion at the line card by randomly dropping packets based on RED/WRED algorithms. RED/WRED offers a programmable minimum discard threshold and maximum discard threshold for each IP precedence value in the ingress and egress Virtual Output Queues (VOQ) to allow CoS-differentiated congestion control. The algorithm is based on an average queue depth, as is required by industry standards.
- RED allows the user to configure one drop profile (minimum and maximum discard thresholds) for all CoS queues on a port. Whereas in WRED, different drop profiles can be applied to different queues on a single port to differentiate treatment of the classes of service within a single port.

- Modified Deficit Round Robin (MDRR) is supported as a scheme to de-queue packets buffered in the line card's ingress and egress Virtual Output Queues (VOQ). MDRR is configurable to operate in one of two modes: (i) Strict Priority Mode, where all packets in the High priority/Low Latency (HP/LL) queue are serviced whenever the queue is non-empty), with a round robin scheme for the remaining 7 CoS queues), and (ii) Alternating Priority Mode, where there is an alternating priority between the HP/LL queue and the remaining 7 CoS queues that rotate in a round robin scheme.

Input Security Extended ACLs

- Configurable on either physical port or VLAN basis.
- Feature is based on Destiny ASIC and TCAM technology for performance that is independent of xACL size.
- Number of xACL rules that can be configured:
  - Actual number depends heavily on the specific xACLs that are configured. Review EDCS-188883 for details if a given service provider is concerned about approaching the limits listed below.
  - When configured on sub-interfaces (VLANs), ACL rules may not be shared among VLANs. When configured on interfaces, ACL rules may be shared among interfaces.
  - Practically speaking, The modular GE card supports greater than 6,000 input extended ACL rules per line card. This corresponds to 600 rules per GE port.
  - Practically speaking, The 1x10GE line card supports greater than 3,000 input extended ACL rules per line card.
  - When configured, the ACL "merge" algorithm will combine all configured ACLs on a given line card into a merged ACL that optimizes for a minimum total number of configured ACL rule entries. This has implications for ACL match counters (see Accounting Section above).
- This feature applies equally to multicast traffic as it does for unicast traffic.
- This feature is not supported for MPLS packets, only IP.

Output Security Extended ACLs

- Configurable on a physical port basis, not per VLAN.
- Up to 1,024 output extended ACL rules may be configured. Each rule involving a range (e.g. drop packets sent from IP source addresses between A and B) counts as two rules against this limit.
- When configured, the ACL "merge" algorithm will combine all configured ACLs on a given line card into a merged ACL that optimizes for a minimum total number of configured ACL rule entries. This has implications for ACL match counters (see Accounting Section above).
- Regarding per VLAN output xACLs, IOS will not support use of destination IP subnet as a proxy for VLAN ID. Although this approach is used for input/output CAR (see below), it is not supported for output xACLs because service providers typically configure an ACL independent of the VLAN, and then apply the ACL to one or more VLANs. The end user would not normally include the destination IP subnet that is needed to define the VLAN in the ACL, before later applying it to that VLAN. In such a case, a security hole could result. While a similar problem might arise when configuring CAR rules, the downside is not as significant as it would be with security xACLs.
- Filtering is performed on the egress Modular GE card or the 1x10GE line card in the chassis, and does not follow the original Engine 0, 1, or 2 model in which output ACL filtering was performed on ingress line cards. As such, feature characteristics are determined by capabilities in TX+ ASIC: (i) some performance impact for long ACLs, so performance is not independent of ACL size, (ii) TX+ does not have access to Ethernet frame to get VLAN ID…only TCP/IP header.
- This feature applies equally to multicast traffic as it does for unicast traffic.
- This feature is not supported for MPLS packets, only IP.

CAR

- Configurable on a per physical port basis in both input and output directions.
- This token bucket implementation allows packets to be checked against up to 32 CAR rules per port to determine if packets conform to or exceed a configurable traffic rate. (For VLANs on 1GE or 10GE ports,

limit is 32 divided by number of VLANs configured.) The result of this check will be to transmit the packet, drop the packet, or color (i.e. modify the IP precedence or MPLS experimental value) and then transmit the packet.

- Classification of packets is based on input extended ACL rule matches, as described above. Therefore, CAR rules may be configured to rate limit not only ICMP traffic, but also specific types of ICMP traffic.
- Rules can be applied with less than 1 Mbps granularity.
- This feature applies equally to multicast traffic as it does to unicast traffic.
- This feature is not supported for MPLS packets, only IP packets.

Output Traffic Shaping

- Output Traffic Shaping is supported in the egress direction, but with two options: (i) on a per IPprec value basis, or (ii) on a per physical port basis. Both shaping features are completely independent and can be configured together or separately.
- This implementation allows rate shaping per output queue (8 unicast queues per port) using a programmable token bucket mechanism that allows bursts of traffic to be smoothed out at user-configured rate.
- Rules may be applied with less than 1 Mbps granularity.

VLAN CoS Bit Copying

- Ingress: Support copying 802.1p bits to IP precedence bits. There is no support for creating an arbitrary mapping from 802.1p bits to IP precedence bits. There are two reasons for supporting this feature: (i) It allows an ISP with a substantial Layer 2 access network, where 802.1p bits are extensively used, to map the 802.1p bits to IP precedence, and (ii) Cisco 7600 supports the feature, so it is useful from a feature consistency standpoint.
- Not supported when CAR is enabled because the CAR mapping feature will overrule the VLAN mapping feature. This limitation should not be a problem as CAR is typically configured to rate limit and adjust precedence, so a user would not need to use both CAR and the VLAN CoS bits for copying.
- No support exists for mapping from IP precedence bits to 802.1p bits. Although hardware is capable of supporting this feature, a software implementation would decrease by factor of 8 the number of supported adjacency entries.

HSRP/VRRP

- Support for 256 HSRP/VRRP groups per port.
- HSRP (Hot Standby Routing Protocol) allows one router/line card to automatically assume the function of a second router/line card within 20 seconds (could be much less depending on time out configurations) if the second router fails. Convergence of routing protocols, such as EIGRP or OSPF, requires seconds, and even then, a host's proxy ARP would have to timeout before hosts on the network are updated.
- HSRP allows two or more HSRP configured routers to use the MAC address and IP network address of a virtual router. Hosts forward packets to this virtual router address instead of to the direct address of an HSRP configured router. A single router in this network is 'active,' responding to messages sent to the virtual router's address. Upon failure of the 'active' router (determined by a lack of multicast hello messages from the active router), a 'backup' router will emerge as the new 'active' router. This transition is completely transparent to all hosts on the network.
- No specific per VLAN functionality is required, so this feature is supported when VLANs are configured.
- Load sharing can be achieved with HSRP, by having some hosts configured for one virtual router and the rest configured for another.
- VRRP is a standardized version of Cisco's proprietary HSRP.
- Automatic Protection Switching not supported given lack of SONET framing information. HSRP/VRRP substitutes for APS.

Policy Based Routing

- Packets classified based on input xACL rule match. Upon a match, forwarding decision is based on use of a PBR-specific lookup table.
- Support for 1,024 rules (aggregate with WCCPv2 and input CAR support).
- Support for VRF Selection is under investigation. If supported, this feature would be available post FCS. This feature involves using the Policy Based Routing feature to select a VRF table for packet forwarding.

Detailed Information for ATM line card features (8-Port OC3 STM-1)

Multiprotocol Encapsulation over AAL5 (RFC 1483)

- When using permanent virtual circuits (PVCs), a user has two ways of carrying multiple protocols over Asynchronous Transfer Mode (ATM).
    - Virtual Circuit (VC) multiplexing - The user defines one PVC per protocol. This method uses more VCs than LLC encapsulation, but reduces overhead, because a header is not necessary.
    - LLC/SNAP Encapsulation - The user multiplexes multiple protocols over a single ATM VC. The protocol of a carried protocol data unit (PDU) is identified by prefixing the PDU with a Logical Link Control (LLC)/Subnetwork Access Protocol (SNAP) header.

UNI 3.0/3.1 (including LMI)

- ATM interfaces create SVCs by sending call signaling messages to establish virtual circuits on demand. All devices on the end-to-end path need a special control permanent virtual circuit (PVC) with QSAAL encapsulation. This PVC carries the call signaling messages. Cisco router ATM interfaces support the ATM Forum's UNI 3.0, 3.1 and 4.0 signaling protocol specifications.

ATM Service Classes

- UBR is the QoS class defined by the ATM Forum for ATM networks. UBR allows any amount of data up to a specified maximum to be sent across the network but there are no guarantees in terms of cell loss rate and delay.
- VBR is subdivided into a real time (RT) class and non-real time (NRT) class. VBR (NRT) is used for connections in which there is no fixed timing relationship between samples but that still need a guaranteed QoS.

Per-VC and per-VP traffic shaping

- Traffic shaping is a function typically provided on ATM edge devices to ensure that bursty traffic conforms to a predetermined "contract." To implement traffic shaping, the 8-Port OC-3 ATM Line Card supports per-VC and per-VP shaping, including unmatched n x 64-kbps granularity, allowing flexibility and control over every VC and VP configured. Traffic shaping is supported on both nrt-VBR and UBR (via peak cellrate [PCR]) ATM service classes.

IP/MPLS QoS features (Per-VC WRED/MDRR)

- The 8-Port OC-3 ATM Line Card supports per-VC IP QoS features, allowing customers to apply advanced queuing and bandwidth management using WRED and MDRR, including LLQ, to individual VCs to avoid congestion and delay.

- Similar support is provided with the MPLS experimental bits.

Basic MPLS Switching with LDP

- Unicast switching, with specific support for label push operations and label pop
- Packet fragmentation supported, where MPLS packets larger than the configured MTU size are fragmented and forwarded by line card CPU. While length of IP packet and Layer 2 header are both included in the comparison against MTU, length of appended MPLS labels is not included.
- Time-To-Live (TTL) Propagate: Support for both (i) "Propagate" mode, where the IP TTL value is copied to the MPLS TTL during label imposition, or (ii) "No-Propagate" mode, where a constant value of 255 is copied. In both cases, the TTL value is decremented at each MPLS node
- Number of LSPs supported is complicated function of line card memory structure and route processor resource requirements.

MPLS VPN PE

- Number of MPLS VPN VRF tables is limited to 256 in software.
- Number of MPLS VPN routes limited to 120,000 per line card.
- Support for aggregate and non-aggregate labels. PE routers use aggregate labels in an effort to conserve MPLS label usage. An aggregate label effectively represents an aggregate or summary of VPN routes to individual hosts or small subnets, and avoids the need for a larger number of individual VPN routes to/from each of these destinations/sources. The 8-Port OC3 ATM line card supports both imposition and disposition of aggregate labels.

Modular QoS Command Line Interface (MQC)

- The Modular QoS CLI allows users to specify a traffic class independently of QoS policies.

- The Modular QoS CLI is used to configure QoS. The Modular QoS CLI contains the following three steps:

    - Defining a traffic class with the class-map command

    - Creating a service policy by associating the traffic class with one or more QoS policies (using the policy-map command)

    - Attaching the service policy to the interface with the service-policy command

For more detailed information about the platforms and features being delivered in 12.0(21)SZ, please reference the following documents:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/relnote/12000ser/index.htm
http://www.cisco.com/univercd/cc/td/doc/product/software/

Support

Cisco IOS software release 12.0(21)SZ follows the standard Cisco support policy as indicated in the following link:

http://www.cisco.com/warp/public/437/27.html

Product Numbers

Cisco IOS Release 12.0(21)SZ Feature Sets, Images, and Memory Recommendations

Platform	Software Feature Set	Product Code	Image	Flash	DRAM
12000	Cisco 12000 Series IOS SERVICE PROVIDER	S120Z-12.0.21SZ	gsr-p-mz	20MB	128MB
12000	Cisco 12000 Series IOS SERVICE PROVIDER/SECURED SHELL 56	S120K3Z-12.0.21SZ	gsr-k3p-mz	20MB	128MB
12000	Cisco 12000 Series IOS SERVICE PROVIDER/SECURED SHELL 3DES	S120K5Z-12.0.21SZ	gsr-k4p-mz	20MB	128MB

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

European Headquarters
Cisco Systems Europe
11 Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe