

Cisco Catalyst 5000 Series Release 4.1 Supervisor Software

Overview

Release 4.1 of the Catalyst® 5000 series supervisor software adds support for high-performance application-specific integrated circuit (ASIC)-assisted, Layer 3 switching; Gigabit Ethernet modules; enhanced Fast EtherChannel® backbone modules and offers many significant new software features. Release 4.1 runs on all Catalyst 5000 series Supervisors with at least 16 MB of DRAM and on the Catalyst 2926.

Features at a Glance

New Hardware Support

- NetFlow Feature Card (NFFC) on Supervisor Engine III
- 3-port Gigabit Ethernet module (WS-X5403)
- 2-port 1000BaseSX uplink module for Supervisor-III (WS-U5534-GESX)
- 24-port 10/100BaseTX backbone Fast Ethernet switching module, RJ-45 (WS-X5225R)
- 12-port 100BaseFX backbone Fast Ethernet switching module, SC (WS-X5201R)

New Software Features

- Layer 3 IP switching¹
- NetFlow data export¹
- Protocol filtering
- IGMP snooping
- IP TraceRoute
- Backbone fast convergence–STP enhancement
- Switch TopN ports report
- Active uplinks (enable uplink ports on standby Supervisor)
- IEEE 802.1Q VLAN trunking
- Multiple default IP gateways
- Configuration change syslog message and SNMP trap
- RMON2 configuration group
- CISCO-SWITCH-ENGINE-MIB

1. Use of this feature requires the purchase of an enhanced feature set software license (see ordering information).

- RSM support for Token Ring VLANs
- CISCO-MEMORY-POOL-MIB
- ENTITY-MIB
- Several new and enhanced SNMP traps

Table 1 Release 4.1 Software Feature-Function-Benefit Table

Feature	Function	Benefit
Enhanced Feature Set		
Layer 3 IP Switching	<ul style="list-style-type: none"> • Hardware-accelerated IP switching/forwarding with a NetFlow Feature Card • Learns routed IP flows • Maintains a Layer 3 flow cache • Gathers detailed traffic statistics 	<ul style="list-style-type: none"> • Next-generation Layer 3 switching performance for campus networks
NetFlow Data Export and NetFlow Cache Summary Report	<ul style="list-style-type: none"> • NetFlow Data Export sends summary statistics on all routed traffic flows to a SwitchProbe,® which then proxies this data to RMON2 for analysis from CWSI/TrafficDirector™ • The NetFlow Cache Summary Report provides detailed Layer 3 flow statistics from the console/Telnet command line interface (CLI) 	<ul style="list-style-type: none"> • NetFlow Data Export and the NetFlow Cache Summary Report provide extensive Layer 3 traffic statistics on network flows
Standard Feature Set Console/Telnet		
Protocol Filtering	<ul style="list-style-type: none"> • Wire-speed, hardware-based protocol filtering • Requires a NetFlow feature card 	<ul style="list-style-type: none"> • Eliminates unnecessary traffic on user ports based on protocol stack • Feature can be used to block IPX SAP and RIP packets from consuming bandwidth and host CPU cycles for IP-only clients
IGMP Snooping	<ul style="list-style-type: none"> • Enables the Catalyst 5000 to act directly upon IGMPv2 join and leave messages for fast and efficient pruning of multicast traffic • Requires a NetFlow feature card 	<ul style="list-style-type: none"> • High-performance standards based multicast traffic pruning • Directs multicast traffic only where needed for optimal network bandwidth utilization
IP TraceRoute	<ul style="list-style-type: none"> • Provides support for the standard ICMP-based IP trace route facility from the system console • Displays the IP routers traversed in order to reach a destination IP host 	<ul style="list-style-type: none"> • Provides a common and useful IP connectivity troubleshooting tool directly from the Catalyst system console
BackboneFast Convergence	<ul style="list-style-type: none"> • BackboneFast is Cisco's latest enhancement to spanning tree, which can reduce convergence time for certain types of network faults by up to 20 seconds • (This is done by advancing the standard spanning tree, Max-Age timer to zero once root path stability is confirmed) 	<ul style="list-style-type: none"> • Maximizes network availability • No manual tuning required • Extends Cisco's existing PortFast and UplinkFast spanning tree convergence time enhancements
Switch TopN Ports Report	<ul style="list-style-type: none"> • Provides a sorted list of the most active switch ports by any of the following metrics: utilization, packets, bytes, multicasts, broadcasts, errors, or buffer overflows 	<ul style="list-style-type: none"> • Helps the network administrator quickly identify which ports on the system have the highest utilization, broadcast rates, error rates, and so forth
Active Uplinks	<ul style="list-style-type: none"> • Enables the ports on the standby supervisor to be used just like any other line card ports on the system 	<ul style="list-style-type: none"> • Maximizes utility of uplink ports on a standby Supervisor module
IEEE 802.1Q VLAN Trunking	<ul style="list-style-type: none"> • Provides IEEE 802.1Q-compliant VLAN frame encapsulation 	<ul style="list-style-type: none"> • IEEE standards-based VLAN encapsulation and trunking for Ethernet VLANs only • Enables interoperability with third-party equipment
Multiple Default IP Gateways	<ul style="list-style-type: none"> • Allows configuration of several fallback default IP gateways for outbound system console Telnet, TFTP and SNMP traffic 	<ul style="list-style-type: none"> • Provides more reliable and robust communications from the Catalyst system console and SNMP agent
Configuration Change Syslog Message and SNMP Trap	<ul style="list-style-type: none"> • Sends notification to configured Syslog receivers and/or SNMP Trap receivers when the system's NVRAM configuration is changed from either SNMP or the console/Telnet CLI 	<ul style="list-style-type: none"> • Alerts the network administrator and network management system that the switch's configuration has been updated and provides information on who made the change (TACACS user name, DNS name, IP address, or console port)
RMON2 Agent/Probe Configuration Group	<ul style="list-style-type: none"> • Adds support for the RFC 2021 RMON2 Agent/Probe Configuration group. Supports the following MIB objects: <ul style="list-style-type: none"> – trapDestination table – probeCapabilities – probeSoftwareRev – probeHardwareRev 	<ul style="list-style-type: none"> • Enables use of an IETF standard MIB for configuration of SNMP trap receivers • Provides information on which RMON and RMON2 groups the embedded agent supports • Provides software and hardware revision information
SNMP Trap Enhancements	<ul style="list-style-type: none"> • New warmStart Trap sent on redundant supervisor switchover • ModuleUp/Down Trap now includes moduleType and moduleName • LinkUp/Down traps now include information from the port security feature • New Spanning Tree Topology Change Trap (only sent from current spanning tree root switch) 	<ul style="list-style-type: none"> • The system will now send a standard SNMP warm start trap when a standby supervisor becomes active • Several MIB varbinds (variable bindings) have been added to existing SNMP traps to provide more information to the network administrator
SNMP MIB Additions and Enhancements	<ul style="list-style-type: none"> • CISCO-SWITCH-ENGINE-MIB • CISCO-MEMORY-POOL-MIB • ENTITY-MIB 	<ul style="list-style-type: none"> • The switch engine MIB enables configuration and/or monitoring of the following: Layer 3 switching engine cache aging mechanisms, active Layer 3 flows, Layer 3 switching performance, NetFlow Data Export, and protocol filtering • The memory pool MIB provides information on the amount of used and free DRAM on the active system Supervisor module • The entity MIB is an emerging standard MIB to provide information on system elements



Multilayer Switching (MLS)

With Release 4.1 software and the NetFlow feature card (NFFC) the Catalyst 5000 family of LAN switches now provides Layer 3 switching at performance levels previously only available for Layer 2 switching. The actual Layer 3 switching function now resides in silicon on the NFFC. The NFFC first identifies flows by using both network Layer and transport layer information and then rewrites and forwards packets between subnets using Cisco's advanced switching ASICs. The RSM performs the route processing on the Catalyst switch with a NFFC while providing central configuration and control of all Layer 3 services. Routing services can also be provided by an externally attached Cisco 7500, 7200, 4700, or 4500 series router.

A NFFC populates its Layer 3/Layer 4 switching cache dynamically by observing/learning the flow of a traditionally routed packet. The NFFC performs flow classification by parsing each packet (in hardware) as far as the transport layer header. In order to perform Layer 3 switching, the NFFC must see the original packet destined for the router (a candidate) and the "routed packet" (enabler) returned from the router. When the Catalyst switch with an NFFC is switching IP packets, it is performing complete rewrites of the VLAN index, Layer 2 source and destination addresses, TTL and TOS in the IP packet header, and recalculating and rewriting the IP header checksum and Layer 2 frame checksum just as a traditional router would.

Cisco IOS[®] software running on the route switch module (RSM) has the ability to instruct the NFFC hardware via a lightweight control protocol called the Multilayer Switching Protocol (MLSP), to flush cache entries in the event of topology change or modification of access control lists. This enables the NFFC to enforce access control lists based on IP addresses as well as transport-layer information.

Components of Cisco's Multilayer Switching (MLS)

- MLS Switching Engine (MLS-SE)—a Layer 3 switching/forwarding entity (for example, the NFFC)
- MLS route processor (MLS-RP)—a Layer 3 route processing entity (for example, an RSM)
- MLS protocol (MLSP)—the protocol that runs between the MLS-RP and MLS-SE

MLSP provides a mechanism for the Route Processor to:

- Invalidate/purge cache entries in the NFFC when access lists or routes change
- Configure various parameters such as the flow mask in the switching engine
- Discover which switching engine has a cache entry and is forwarding traffic for a particular IP address pair
- Manually install flow cache entries

MLS prerequisites:


- NetFlow feature card
- Release 4.1 Enhanced Feature Set license
- The appropriate IOS image running on an RSM or a directly attached Cisco 7500, 7200, 4700, or 4500 series router

NetFlow Data Export

With NetFlow Data Export, the performance management capabilities of the Catalyst switches have now been extended to provide comprehensive monitoring of all inter-subnet flows passing through the NFFC and the RSM. A new NetFlow Data Export version 7 format is supported from the Catalyst 5000 with a NFFC when it is performing Layer 3 switching. This complements the embedded mini-RMON capabilities already available on the Catalyst switch, which provide visibility into all port traffic at Layer 2.

The NetFlow Data Export mechanism captures Layer 3 traffic statistics as each NFFC cache entry expires. It then bundles several of these statistics records into a User Datagram Protocol (UDP) datagram and exports it to a NetFlow Data Collector such as a Cisco SwitchProbe. The SwitchProbe will proxy or map the NetFlow statistics to the appropriate standards-based RMON2 groups for analysis from CWSI/TrafficDirector or any RMON2-compliant application.

Figure 1 NetFlow Data Export to a Cisco SwitchProbe



The flow statistics gathered in the data collection phase enable several key customer applications and benefits:

- **User Monitoring and Profiling:** NetFlow data enables network managers to gain a detailed understanding of user utilization of network and application resources. This information can then be used to efficiently plan and allocate access, backbone, and application resources as well as to detect and resolve potential security and policy violations.
- **Application Monitoring and Profiling:** NetFlow data enables network managers to gain a detailed view of application traffic patterns over the network. Content providers can use this information to plan and allocate network and application resources (for example, Web server sizing and location) to responsively meet customer demands.
- **Network Monitoring:** NetFlow data enables advanced and comprehensive network monitoring capabilities. RMON, RMON-2, and flow-based analysis techniques can be used to visualize traffic patterns associated with individual switches as well as on a network-wide basis and provides proactive problem detection, efficient troubleshooting, and rapid problem resolution.

Additionally, the NetFlow Cache Summary Report provides Catalyst system console access similar to the Cisco IOS `show ip cache flow` command to display a report of the available NetFlow traffic statistics.

Note: For most complete statistics, MLS must be in “full flow mask” mode.

Protocol Filtering

Catalyst 5000 systems with NetFlow feature cards can now perform hardware-based protocol filtering to allow further segmentation within a VLAN by protocol.


Protocol filtering enables per-port filtering of traffic by any of the three following protocol groups:

- IP
- IPX
- AppleTalk, DECnet, and VINES

These groups can be individually enabled or disabled on each switch port. By default the IP group is set to ON and the other two groups are set to AUTO. When a port becomes a part of the group, it starts receiving broadcasts for that group. So, for example, a port attached to a client system that will only be running IP can be enabled for IP and disabled for IPX and AppleTalk/VINES/DEC so it won't receive broadcasts from those “chatty” protocols. Additionally, this feature can be used to provide security by protocol so a port can be disabled for IP traffic; for example, if a network administrator only wants to allow Novell users on a given network.

Additionally, an AUTO mode allows the switch to automatically enable the port for a given protocol group if a frame of that protocol type is received on the port from the client. For example, a switch port may have IP enabled and IPX set to AUTO, so initially the port will only forward IP traffic to the IP-only client. At a later date, if an IPX stack is loaded onto the client and the client sends an IPX packet to the switch, the switch will automatically enable IPX on this port.

Figure 2 Protocol Filtering on Catalyst 5000 series switch ports

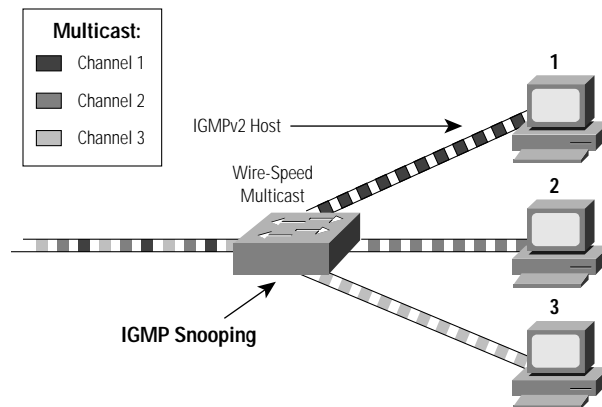


IGMP Snooping

Catalyst 5000 systems with a NetFlow feature card are now capable of looking into every frame deeply enough to detect and act upon the IGMP version 2 Join and Leave requests from switch-attached workstations.

IGMP Snooping is a completely standards-based mechanism and operates on IGMP version 2 Join and Leave messages. It provides faster Join and Leave processing than other approaches and is especially valuable when bursts of joins/leaves are received by the switch. The NFFC-based IGMP Snooping feature handles IGMP leaves that come on the GDA as well as the ALL-ROUTERS address.

Figure 3 IGMP Snooping on the Catalyst 5000 Series



IP TraceRoute

TraceRoute is helpful in locating routing issues on a distant gateway. If the local routing tables are correct on routers locally administered, distant routers under another administration may be the cause of misrouted or dropped data packets. TraceRoute will give a general location of where the packet is getting lost. Remember, once a packet is handed to the next gateway, it is that gateway's responsibility to see to it that the packet continues correctly to the final destination.

TraceRoute is used to display the IP gateways traversed to reach a destination host. TraceRoute uses UDP data packets with incrementing TTL, or time-to-live, values and an invalid port number to build the TraceRoute list of gateways.

The Catalyst switch originating the TraceRoute begins by sending a UDP data packet with the IP destination address of the remote host and a TTL value of 1. The first gateway that receives the packet will decrement the TTL value by 1. Since the TTL value is now 0, the gateway sends back an ICMP “Timeout Exceeded” message to the originating host. The originating host then adds that gateway, to the TraceRoute list and sends out another data packet with a TTL value of 2. The first gateway will receive the packet and decrement the TTL value by 1 and forward the packet to the next gateway, which will decrement the TTL value to 0, send an ICMP Timeout Exceeded message, and so forth. This process continues until the destination host receives the data packet, it sees that the port number is invalid, and sends back an ICMP “Unreachable Port” message back to the originating host. When the originating host receives this message, it knows that the host has been reached and ends the TraceRoute.

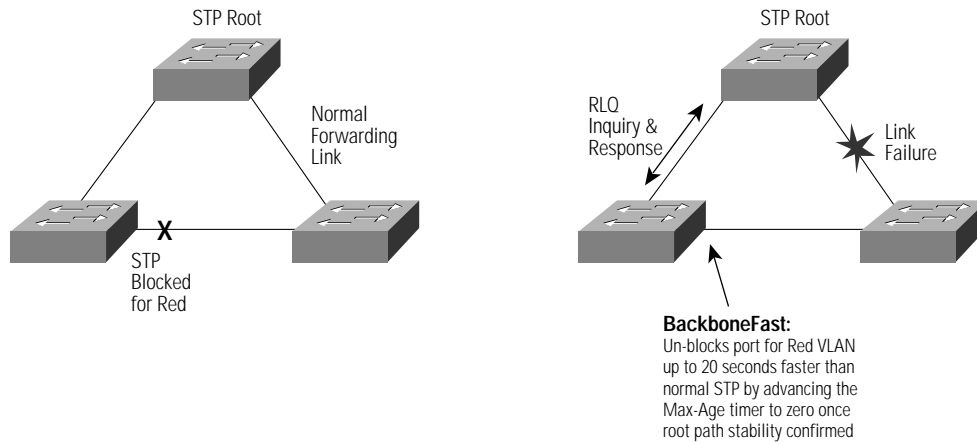
BackboneFast Convergence

BackboneFast is Cisco’s latest enhancement to spanning tree; it can reduce convergence time for certain types of network faults by up to 20 seconds. This is achieved by reducing the normal 20-second Max-age timer period to a few seconds on blocked ports for indirect failures in a spanning tree.

Normally when a switch receives an inferior BPDU on a blocked port it waits through the entire max-age interval (default is 20 seconds) before it transitions to the listen and learn states and eventually makes that port a designated port. Generally when a switch receives an inferior BPDU on any of its blocked ports (except self-looped ports) and the inferior BPDU is sourced by the designated switch for that port/segment, this indicates an in-direct network failure, meaning that the other switch has lost its path to the spanning tree root.

In many cases waiting through the entire max-age interval is unnecessary and only serves to delay spanning tree convergence. Cisco has optimized this state by first confirming root path stability by: a) sending a new protocol packet called a Root Link Query (RLQ) from its root port, and then upon receiving an affirmative response, b) advancing the Max-Age timer to zero on ports receiving the inferior BPDU.

Figure 4 BackboneFast Convergence



BackboneFast extends Cisco’s existing PortFast and UplinkFast spanning tree convergence time enhancements.

Switch TopN Ports Report

This feature will help network administrators quickly identify which switch ports have the highest utilization, broadcast rate, or error rates. The Switch TopN utility will collect the following data for each physical port over a user-specified interval:

- Port utilization
- Bytes (transmitted and received)
- Packets (transmitted and received)
- Broadcast packets (transmitted and received)
- Multicast packets (transmitted and received)
- Errors (received)
- Number of buffer-overflow errors

At the end of the sample interval, a delta value is calculated for each item above (final value less initial value). The list is then sorted by the user-selected metric, and the first N ports are displayed to the user. Additionally, the user may specify the type of port for the report. Options are: All (default), Ethernet, Fast Ethernet, Gigabit Ethernet, Token Ring, or FDDI ports. This feature also supports an option to run it in the background and notify the user when a report becomes available. This is useful when running reports over longer intervals such as 15 minutes.

802.1Q VLANs

IEEE 802.1Q Will Initially Be Supported On The Following Modules:

- 3-port Gigabit Ethernet module (WS-X5403)
- 2-port 1000BaseSX uplink module for Supervisor III (WS-U5534-GESX)
- 24-port 10/100BaseTX backbone Fast Ethernet switching module, RJ-45 (WS-X5225R)
- 12-port 100BaseFX backbone Fast Ethernet switching module, SC (WS-X5201R)

The 802.1Q encapsulation is compliant with the IEEE 802.1Q draft standard. 802.1Q is an “internal tagging” or one level tagging scheme, whereas Cisco’s ISL is an “external tagging” or two level tagging scheme. Currently 802.1Q VLANs are only defined for Ethernet frames. See figures below on 802.1Q tagging.

Cisco’s PVST (Per-VLAN Spanning Tree) has been modified to support the CST (Common Spanning Tree) environment required by IEEE 802.1Q. These modifications, simply referred to as PVST+, map 802.1Q’s CST to the PVST spanning tree for one VLAN (normally VLAN #1). All the other PVST spanning trees are tunneled through the CST. These tunneled BPDUs are sent to a different multicast address, so that the 802.1Q switches see them as regular multicast packets.

Note: The Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) will be supported in the Boulder release which is currently targeted for Q2 - CY '99.

Figure 5 802.1Q Internal VLAN Tagging

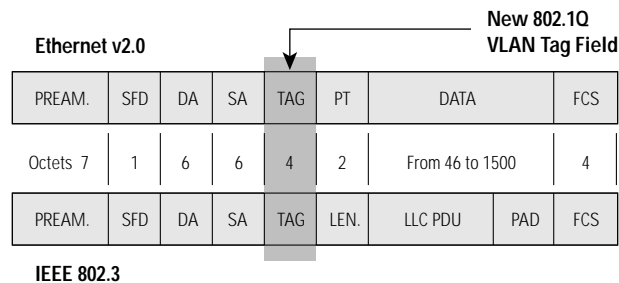
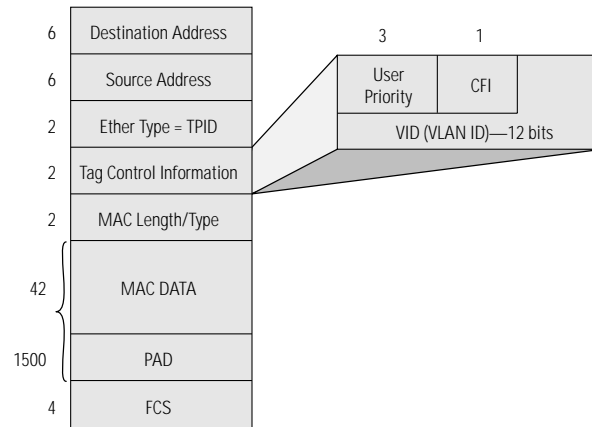


Figure 6 802.1Q VLAN Tag Field



SNMP Traps

To further enhance the manageability of the Catalyst 5000 series, several new MIBs and SNMP traps have been added and/or extended.

The following new SNMP traps have been added in release 4.1:

- Configuration Change Trap: New trap to indicate that the system configuration in NVRAM has changed
- WarmStart Trap: To be sent when a standby supervisor becomes active

The following SNMP traps have been enhanced to include additional varbinds:

- Spanning Tree Topology Change Trap: Now includes ifName of transitioning port
- ModuleUp/Down Trap: Now includes Cisco Stack MIB.moduleType and Cisco Stack MIB.moduleName
- LinkUp/Down Trap: Now includes an indication if linkDown was caused by a port security violation and to include LastSrcAddr
- Chassis alarm trap: Now adds the following MIB objects to the trap varbind list:
 - Cisco Stack MIB.chassisPs1Status
 - Cisco Stack MIB.chassisPs1TestResult
 - Cisco Stack MIB.chassisPs2Status
 - Cisco Stack MIB.chassisPs2TestResult
 - Cisco Stack MIB.chassisFanStatus
 - Cisco Stack MIB.chassisFanTestResult
 - Cisco Stack MIB.moduleStatus SupervisorModule (for active supervisor)
 - Cisco Stack MIB.moduleTestResult SupervisorModule (for active supervisor)

Useful References

- **Catalyst 5000 Series Documentation**
<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/index.htm>
- **Catalyst 5000 Series Release Notes**
<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/c5krn/index.htm>

System Requirements for Release 4.1

Release 4.1 Software runs on all Cisco Catalyst 2926 Switches and Catalyst 5000 Series Supervisors with at least 16 MB of DRAM:

- **Supervisor Engine I modules with 20 MB DRAM**
- **All Supervisor Engine II modules (16 MB default DRAM)**
- **All Supervisor Engine III modules (32 MB default DRAM)**

Ordering Information

The standard and enhanced feature sets are combined in a single image for the 4.1 release. Use of Layer 3 switching on the NetFlow Feature Card and NetFlow Data Export require purchase of the Enhanced Feature Set license, which also includes a license for mini-RMON. If only mini-RMON is to be used, the mini-RMON license must be purchased. Only one enhanced feature set software license or RMON agent license is required per chassis (that is, a chassis with redundant supervisors does NOT require two licenses). Licensing is based on the honor system.

The Enhanced Feature Set License covers use of the following features:

- **Layer 3 Switching**
- **NetFlow Data Export**
- **Mini-RMON Agent**

Table 2 Product Ordering Information

Product Number	Description
Release 4.1 Enhanced Feature Set License	
SWC5K-BEFS-4X	Catalyst 5000 Rel. 4.x SW License, Enhanced Feature Set
SWC5K-BEFS-4X=	Catalyst 5000 Rel. 4.x SW License, Enhanced Feature Set (spare)
SWC5K-BEFS-4X-UG ¹	Catalyst 5000 Rel. 4.x SW License, Enhanced Feature Set, upgrade
Mini-RMON Agent Licenses	
WS-C5002-EMS-LIC	Catalyst 5002 RMON Agent License
WS-C5K-EMS-LIC	Catalyst 5000 RMON Agent License
WS-C5505-EMS-LIC	Catalyst 5505 RMON Agent License
WS-C5500-EMS-LIC	Catalyst 5500 RMON Agent License

1. Customers who have previously purchased RMON agent licenses for the Catalyst 5000 series chassis where the NFFC will be installed are eligible for the upgrade price, which acknowledges prior purchase of the RMON agent license.

Customers can download Release 4.1 Supervisor software from Cisco Connection Online (CCO) in the Software Image Library. Customers who are unable to download the files electronically can order a software kit with 3.5-inch IBM formatted floppy disks by contacting Cisco at 408 526-4000 or, in North America, call 800 553-NETS (6387).

If you have specific questions about this product bulletin or input regarding Catalyst 5000 series Supervisor software enhancements, please email bdebolle@cisco.com.

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<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
Villebon, BP 706
91961 Courtaboeuf Cedex
France
<http://www-europe.cisco.com>
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
<http://www.cisco.com>
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters

Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
<http://www.cisco.com>
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 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
England • France • Germany • Greece • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Russia • Saudi Arabia • Scotland • Singapore