APPENDIX B

Integrity Checks

The Cisco IP Manager uses checking capabilities developed for the Cisco Netsys network management software. Results of the check are displayed in the **Integrity Check** window.

Feedback is categorized by transport modes (IP, IPX, SRB, SNA STUN, Token Ring, AppleTalk, DECnet, and VINES) and LAN switching, which are then further sub-categorized (possibly subjectively) as High Severity or Warning-Level.

High Severity means the reported problem is believed to be critical, resulting in *major* network problems. These problems *must* be fixed to ensure the network is fully operational and functional.

Warning-Level means the problem is not considered severe, but *should* be fixed to prevent inadvertent side effects and potential network performance degradation.

IP

This section describes the IP integrity checks. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

IP High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following IP high-severity problem entries is placed in the Integrity Checks report and the accompanying text displayed when the **Explanation** button is clicked:

Addresses of an Interface in Different OSPF Areas

The **network area** OSPF commands are sequentially evaluated to determine which area an interface is to be assigned. All addresses on an interface must map to the same area.

Bad Address/Mask on Router Interface

An IP address of an interface is comprised of an address and mask. The mask restricts the network or subnet to which routing updates are broadcast or received. The mask must be compatible with the class (A, B, or C) of the interface address. This check identifies erroneous IP address/mask combinations. For example, an interface address of 198.23.56.0 and a mask of 255.255.0.0 are not compatible since a class B mask is provided for a class C address. This check also verifies the broadcast address for a subnet has not been selected as a **hostid**. This check also verifies that IP subnet-zero is configured when subnet zero is employed.

Bad Masking in Access List

Within the **access-list <access-list-number> {deny|permit}** commands, a setting of 1 within a mask position is an indication to ignore the corresponding bit in the associated address. A 1 in the same position of the address and mask is dubious, since it is an instruction to pay attention as well as ignore the associated bit. For example, the command **access-list 99 deny 131.108.34.0 0.255.255.255** is suspect.

Duplicate Addresses

All IP addresses assigned to router interfaces are checked for duplication. Duplicate addresses can result in serious problems both in a live internetwork and in the Cisco IP Manager simulation of an internetwork.

Encapsulation Mismatch Among Connected Interfaces

Connected interfaces should be configured with compatible encapsulation settings. Otherwise, communication between the involved interfaces will be compromised.

IBGP Neighbors Not Fully Meshed

Internal Border Gateway Protocol (IBGP) is used when an autonomous system (AS) wants to act as a transit system to other ASs. When a BGP speaker receives an update from another BGP in its own AS (via IBGP), the receiving BGP speaker does not redistribute that information to other BGP speakers in its own AS. The receiving BGP speaker only redistributes that information to other external BGP speakers outside of its AS. This is why it is important to maintain a full mesh between the IBGP speakers within an AS (so no information is lost.) This warning is issued when a partial mesh between the IBGP routers is observed.

Indirectly Connected EBGP Neighbor

Configuring a BGP neighbor/peer establishes a neighbor relationship with another router in a different autonomous system. External BGP neighboring relationships should be among *directly* connected routers. This problem report is issued when an indirect External BGP relationship is observed. The only permitted exception to this rule is when an EBGP multihop connection is enabled.

IP OSPF Network Mismatch (No OSPF Communication)

The **ip ospf network {broadcast | non-broadcast | point-to-multipoint}** command allows the OSPF network type to be configured to a type other than the default for the given media. This value should be the same for all interfaces on the same subnet or there will be no OSPF communication among those routers.

IS-IS Duplicate System IDs Found

No two nodes can have NET (network entity title) addresses in which the system ID fields are the same.

MTU Mismatch Among Connected Interfaces

The MTU (maximum transmission unit) size among connected interfaces should agree. When a packet arrives at an interface whose MTU setting is less than the packet's size, the packet is dropped and the connection fails.

Noncontiguous Mask on Router Interface

The net mask component of the IP address of an interface restricts the network or subnet to which routing updates are broadcast or received. While subnet masks in theory can be anything, they are in practice usually required to be contiguous (all ones from the left). For example, the mask 255.7.0.0, when converted to binary, equates to 1111 1111 0000 0111 0000 0000 0000 0000 and thus violates this convention.

Non-utilized Rule in Access List

The **access-list <access-list-number> {deny|permit}** commands specify a collection of rules which can be utilized for route filtering (distribution lists) and packet filtering (access-groups on interfaces). The rules are evaluated sequentially looking for a match. When a match is found, the packet or routing update is either denied/permitted according to the rule. The matching criteria for an extended access list are source address and mask, destination address and mask, destination TCP/UDP port, and protocol. A danger when creating long access lists is preceding rules may subsume (be more general than) subsequent rules. When this happens, the later rules are never utilized (the prior rule shadows them.) When **access-list 80 deny 198.0.0.0 0.255.255.255** preceded the rule **access-list 80 permit 198.65.0.0 0.0.255.255**, the second rule would never be exercised. When the access list action differs between the two rules, as in this case, a (declarative) inconsistency exists which is only resolved by virtue of the sequential (that is, procedural) evaluation of the rules. This is classified as a high severity violation. When, on the other hand, the access list actions agree, the second rule is unnecessary and is identified as a warning. The potential for subsumption also exists in the protocol restriction (ip subsumes tcp, udp, icmp, igrp) and in the port restriction (gt 0 subsumes gt 1024). This check performs a pair-wise comparison between every rule in the same access list.

OSPF Area Does Not Border Area Zero

An OSPF network should be designed such that all areas border area zero. When a connection to the backbone is lost, it can be repaired by establishing a virtual link. Virtual links are defined using the **area virtual link** router subcommand. The current version of the Cisco IP Manager software does *not* model the **virtual link** option, therefore, when you are using this command, you can ignore this check.

Overlapping IP Subnets

Subnets denote a range of host IDs. Subnet host-id ranges should be mutually exclusive. Host addresses must map uniquely to one, and only one, subnet. In general, an encompassing subnet range overlaps with many others. For example, a misconfigured subnet, with an IP address 131.108.1.2 and mask 255.255.0.0, encompasses all legal subnets on this class B network. In the interest of reducing many redundant problem reports, only the first overlap is flagged.

Reflector Clients Not Fully Meshed when Client-to-Client Reflection is Off

When route-reflectors are employed and the command **no bgp client-to-client reflection** is configured on the reflector, a full mesh must be maintained between the clients of this reflector.

Same IP Subnet Restriction

Loopback interfaces on different routers should not be assigned the same IP subnet, as that causes routing ambiguities. Also, two interfaces on the same router cannot be assigned the same IP subnet, unless one of the interfaces serves as a backup to the other. The router will not allow such a configuration.

Same IPX Net Restriction

Two interfaces on the same router cannot be assigned the same IPX network. The router will not allow such a configuration.

Static Route Next Hop is a Shutdown Interface

The **ip route <network> [<mask>] {<address>|<interface>} [<distance>]** command has a parameter to specify a forwarding interface. This check identifies static route definitions that forward to a non-existent or shutdown interface.

(Sub)Net Mask Creates Hostid of Zero

The net mask component of the IP address of an interface restricts the network or subnet to which routing updates are broadcast or received. By convention, the mask should not create a **hostid** of zero. For example, the IP address 158.131.67.17 with mask 255.255.255.240 describes an interface with a **hostid** of 1 on subnet 158.131.67.16. The IP address 158.131.67.16 with mask 255.255.255.240 breaks this convention, as it results in a **hostid** of 0. The Cisco IP Manager software disallows the attachment of this interface to the implied (sub)net.

Subnet Not Consistently Assigned to an OSPF Area

All routers running OSPF should assign shared subnets to the same OSPF area. Otherwise, OSPF routing can become unstable. This check flags cases where the subnet is assigned to different areas from different router perspectives.

Unbalanced BGP Neighbors

Configuring BGP neighbors/peers establishes a neighbor relationship with another router. These neighboring relationships should be balanced (that is, when *router 1* is configured to neighbor *router 2*, router 2 should also be configured to neighbor router 1.) This check also verifies the corresponding **neighbor ... update-source** command is also applied when a loopback address is employed in a neighboring relationship.

Unbalanced Frame Relay Mapping

To use the **frame-relay map <protocol> <protocol-address>** command, the protocol address specified must map to a known router interface. A balancing **frame-relay map** command on that interface, which refers back to the originating interface, must also exist.

Unbalanced SMDS Static-Mapping

To use the **smds static-map <protocol> <protocol-address>** command, the protocol address specified must map to a known router interface. A balancing **smds static-map <protocol>** command on that interface, which refers back to the originating interface, must also exist.

Unbalanced X25 Mapping

To use the **x25 map <protocol> <protocol-address>** command, the protocol address specified must map to a known router interface. A balancing **x25 map <protocol> <protocol-address>** command on that interface, which refers back to the originating interface, must also exist.

Undefined Access List Referenced

The **ip access-group ...** and **distribute-list ...** commands require a reference to an access list. This check identifies references to undefined access lists. Depending on the IOS version, an undefined access list has implicit behavior of denying or allowing all access. In 9.x IOS releases, the behavior is **deny all**. In 10.x IOS releases, the behavior is **allow all**. This check has been extended to also handle the **ip as-path access-list** command.

Undefined Route-Map

A route-map defines the acceptance and denial policies for routing updates. When a route-map is referenced, but not defined, you will not receive a warning from the router. This can lead to unexpected behavior, since no policy will be enforced. Route-maps are referenced within redistribution and neighbor policy definitions.

Unknown Address in SMDS Mapping

To use the **smds static-map <protocol> <protocol-address> <smds-address>** command, both the protocol address and smds address parameters must map to the same known router interface. A high severity violation occurs when the protocol address maps to an interface but the smds address is not assigned to the same interface.

Unknown Address in X25 Mapping

To use the **x25 map <protocol> <protocol-address> <x.121-address>** command, both the protocol address and x.121 address parameters must map to the same known router interface. This warning is issued when the protocol address does not map to an interface. When it does map to a known interface, a high severity violation occurs when the x.121 address has not also been assigned to the same interface.

Unknown IBGP Neighbor

An IP address mentioned in the list of IBGP neighbors/peers should map to a known router interface. This may simply be the result of an incomplete collection of configurations.

Unthrottled Redistribution of Routes

It is good practice to apply a route filter whenever routes are redistributed between dynamic routing protocols. When back doors exist, routes can *leak* back, thus creating potential routing loops. A very dangerous scenario occurs when BGP is redistributed into an Interior Gateway Protocol. BGP tables are normally *huge* and the resultant flood of routes can overwhelm your interior gateway routing protocols.

IP Warning-Level Checks

When a potential problem is uncovered while performing the following checks, one of the following IP warning-level problem entries is placed in the Integrity Checks report and the accompanying text displayed when the **Explanation** button is clicked:

Bad Default Network Specification

The **ip default-network <network-number>** command allows the specification of a catch-all forwarding vehicle to destinations unknown to the local routing processes. This check identifies network addresses not compatible with the class of the network. For example, a default network of 128.0.0.0 is illegal, as this is a reserved class B network. However, a default network address of 90.0.0.0 is legal, as it is a class A network and thus only requires the first octet.

Bad Network Address Specified in Routing Process

The **network <network-number>** router sub-command restricts the networks from which routing updates are acquired and to which local routing updates are advertised. This check identifies network addresses not compatible with the class of the network. For example, a network address of 192.0.0.0 is illegal, as this is a reserved class C network. However, a network address of 90.0.0.0 is legal, as it is a class A network and thus only requires the first octet. Also identified are host or subnetted addresses, as a major network address is required.

Bad Target in Static Route Definition

The mask of the destination address of a static route must be compatible with the class (A, B, or C) of the destination address. This check identifies erroneous address/mask combinations. For example, a destination address of 198.23.56.0 and a mask of 255.255.0.0 are not compatible as a class B mask is provided for a class C address.

Connected IP (Sub)Net not Advertised by RIP/IGRP/EIGRP/OSPF/BGP/ISIS/Static

The **network <network-number>** router sub-command specifies the networks upon which a routing process advertises and accepts routing updates. When a connected network is not mentioned in the **network** command of a defined routing process (such as rip, igrp, eigrp, ospf) routing updates are not advertised to the interfaces of those networks. This check identifies all connected networks and subnets not advertised by any of the local routing processes (including IS-IS). When a connected subnet is mentioned indirectly through a redistributed static route’s next-hop, it satisfies this check and is not reported. When a connected subnet is originated by a BGP autonomous system, it also satisfies this check and is not reported. When a connected subnet is covered by an aggregate entry in the BGP table, it also satisfies this check and is not reported. When one of the dynamic routing algorithms redistributes connected routes, it also satisfies this check and is not reported.

Dead-End Serial Interface

A topology is created by piecing together the information resident in a specified collection of router configuration files. A topology is a collection of *nodes* and *edges* inter-connecting the nodes. The nodes in the topology are comprised from the observed routers and the serial links and LAN segments denoted by the interface *type*, where *type* is one of *ethernet, Token*

Ring, fddi, serial (synchronous), hssi, or *bri*. A LAN segment is created for every observed network or subnet, as defined by the masking information in the IP address for interfaces of type ethernet, fddi, or Token Ring. A serial link is created to connect pairs of serial (synchronous), hssi, or bri interfaces which uniquely map to the same network. *Uniquely* means only two router interfaces map to the net or subnet denoted by the link. If more than two serial interfaces map to the same subnet, a check is made to determine if two of them have consecutive IP addresses (such as 199.35.121.22 and 199.35.121.23). If they do, a link is created to connect the two serial interfaces. If they do not, a link is *not* created. Thus, after topology generation, dangling serial link connections may exist. They may also exist due to any inherent incompleteness in the specified collection of router configuration files, which may only represent a subset of the routers. Dead-end serial links are flagged to allow the missing pieces to be filled in.

Fast Switching Low Speed Serial Interface

Fast switching is configured by default. Routers generally offer better packet transfer performance when fast switching is enabled, with the following exception. On networks using slow serial links (64Kbps or less), disabling fast-switching's per-destination, load-balancing behavior and enabling per-packet, load-sharing is usually the best choice. This can be accomplished via the **no ip route-cache** command.

IGRP/EIGRP Metric Mismatch Between Connected Interfaces

IGRP (Interior Gateway Routing Protocol) and EIGRP (Extended Interior Gateway Routing Protocol) factor several interface metrics (such as bandwidth and delay) into their computation of routing cost. When a metric value is not explicitly stated, default values are used based upon the interface type (such as serial, ethernet, or Token Ring) For example, T1 characteristics are assumed for a **serial** interface. The reality may be the interface is connected to a 56Kbps line, meaning traffic may be non-optimally routed across the link. This check examines all interface interconnections (via both serial links and LAN media types) and checks the bandwidth and delay metric specified or assumed on both sides. When the metrics differ, a warning is issued. A common occurrence is a serial link metric is diligently recorded on one end, but not the other. In the case of a 56Kbps line, the link looks attractive in one direction and is avoided, when possible, in the other direction (leading to asymmetric routing). This may be intentional, but is not the typical case.

IP Subnet Zero Configured with Classful Protocol

Using an IP subnet of zero can cause problems with classful routing protocols (RIP or IGRP). While a classless routing protocol includes a subnet mask with all advertisements and poses no problems, a classful protocol can be confused between advertisements for the zero subnet and a network.

IS-IS Hello Interval Mismatch Between Connected Interfaces

IS-IS Hello intervals among connected interfaces should be the same. Level 1 and Level 2 settings are independently verified.

ISIS Link State Metric Mismatch Between Connected Interfaces

An ISIS metric mismatch among connected interfaces can lead to asymmetric routing. Level 1 and Level 2 settings are independently verified.

Mismatched Keepalives on Serial Interface

Mismatched keepalive settings on connected serial interfaces (not including Frame Relay, X.25, or SMDS) can lead to intermittent interface resets, resulting in routing table recalculations.

Mismatched TCP Header Compression

In order for TCP header compression to occur, both sides of a PPP or HDLC encapsulated link must have the **ip tcp header-compression** command enabled. An X.25 link can be configured for TCP header compression when balanced **x25 map compressed tcp** commands exist on both interfaces. When one side of a link indicates compression and the other does not, this warning is issued.

Missed Opportunity for a Passive Interface

IP routing algorithms can selectively advertise on interfaces. When all connected routers on an interface are *not* running one of the algorithms on the current router, bandwidth is being wasted. Alternative ways to do this include setting up a passive interface, removing a **network** command, or setting up a **deny all** outbound route filter. The bandwidth of the attached media will then not be overloaded with routing updates, which are of no interest to the attached parties. This check identifies this situation on serial interfaces, where wasted bandwidth is especially critical.

When serial links are connected via IP and only one has an **ip router isis [tag]** interface command, this warning is issued.

Network Timers Mismatch Among Connected Routers

The **timers basic <update> <invalid> <holddown> <flush> <sleeptime>** command determines the timing of accepting and announcing routing updates. These settings should agree amongst connected routers, otherwise routing tables can become unsynchronized or have convergence problems.

Non-utilized Rule in Access List

The **access-list <access-list-number> {deny|permit}** commands specify a collection of rules which can be utilized for route filtering (i.e. distribution lists) and packet filtering (i.e. access-groups on interfaces). The rules are evaluated sequentially looking for a match. When a match is found, the packet or routing update is either denied/permitted according to the rule. The matching criteria for an extended access list are source address and mask, destination address and mask, destination TCP/UDP port, and protocol. A danger when creating long access lists is preceding rules may subsume (be more general than) subsequent rules. When this happens, the later rules are never utilized (the prior rule shadows them.) When **access-list 80 deny 198.0.0.0 0.255.255.255** preceded the rule **access-list 80 permit 198.65.0.0 0.0.255.255**, the second rule would never be exercised. When the access list action differs between the two rules, as in this case, a (declarative) inconsistency exists which is only resolved by virtue of the sequential (procedural) evaluation of the rules. This is classified as a high severity violation. When, on the other hand, the access list actions agree, the second rule is unnecessary and is identified as a warning. The potential for subsumption also exists in the protocol restriction (**ip** subsumes **tcp, udp, icmp,** and **igrp**) and in the port restriction (**gt 0** subsumes **gt 1024**). This check performs a pair-wise comparison between every rule in the same access list.

Opportunity for Autonomous/Optimum/Flow Switching

These switching modes enable much faster packet processing in most cases. This is configured on interfaces via the **ip route-cache [cbus | sse | optimum | flow]** command. The **cbus** parameter is only available on the Cisco 7000 series and AGS+ systems. The **optimum** and **flow** options are available on the 7000RSP, 7200, and 7500 family of routers. When the actual router model is known, this check will guess the presence of a 7000+ router by looking for the presence of **slot/unit** parameters (**Ethernet0/1**, for example). When custom or priority queuing is already configured on the interface, no recommendation is made since queuing precludes high-end switching. When an IP access list is configured on the interface, no recommendation is made as this usually limits the switching mode.

There are several documented bugs in the IOS that cause routers to crash when using the **cbus** parameter under certain circumstances. Check your configuration and consult the on-line Cisco CIO documentation.

OSPF Metric Cost Mismatch Between Connected Interfaces

An OSPF metric cost mismatch among connected interfaces can lead to asymmetric routing. When a metric cost value is not explicitly stated, default values are derived based upon configured bandwidth.

Primary and Secondary Addresses Map to Same Subnet

The **network <network-number>** router sub-command specifies the networks upon which a routing process advertises and accepts routing updates. A little known restriction is major network routing updates are only broadcast from the primary address of each interface. Thus, routers interconnected via a secondary address (secondary net or subnet) only receive updates about their subnet and not about the other major networks connected to the other router. Implicit route filtering occurs on interfaces through the secondary address. A good policy is to interconnect routers via primary addresses and to only use secondary addresses to make a network contiguous. Thus, primary addresses should map to (sub)nets accessible via a primary address on other routers. Secondary addresses should map to (sub)nets accessible via asecondary address on other routers. In general, problems occur when primary addresses map to (sub)nets only accessible via secondary addresses on other routers. This check identifies all router interconnections via (sub)nets where the first router is logically connected via a primary address and the second router is logically connected via a secondary address.

Redistribution: Metric Value Missing Where No Default

Metrics for different protocols cannot, in general, be directly compared. A default metric helps solve the problem of redistributing routes with incompatible metrics. Whenever metrics do not convert, use of a default provides a reasonable substitute and enables the redistribution to proceed. IGRP/EIGRP can automatically redistribute IGRP/EIGRP/static, but they do not have a metric conversion for RIP. RIP needs a metric conversion to redistribute IGRP, EIGRP, or OSPF. OSPF always produces a metric conversion when it redistributes, but it is best to provide an explicit metric. This check also looks for the presence of metric assignments within all route-maps which accept routes.

Static Route Next Hop is Indirectly Connected

The next hop address of a static route specifies the address of a *directly* connected router that is one hop away. This check identifies static route definitions that forward to what appears to be an indirectly connected router interface. The problem may stem from an incompleteness in the baseline topology. For example, missing serial links may be the source of the problem.

Static Route Next Hop is an Unresolved Address

The next hop address of a static route specifies the address of a connected router. This check identifies static route definitions that forward to an unknown router interface address. Either an erroneous forwarding address was specified or the problem stems from an incompleteness in the baseline router configuration files. The latter case could occur when the baseline comprised router configuration files representing only a subset of the routers in the internetwork.

Subnet Not Consistently Assigned to an OSPF Area

All routers running OSPF should assign shared subnets to the same OSPF area. Otherwise, OSPF routing can become unstable. This check flags the case where an area assignment exists for one router interface, however the second interface is not running OSPF.

Unbalanced Dialer Mapping

To use the **dialer map <protocol> <protocol-address>** command, the protocol address specified must map to a known router interface. A balancing **dialer map** command on that interface, which refers back to the originating interface, must also exist.

Unconnected Net in Network Command of RIP/IGRP/EIGRP/OSPF

The **network <network>** router sub-command specifies the *connected* networks upon which a routing process advertises and accepts routing updates. When an unconnected network is mentioned in the **network** command of a defined routing process (such as *rip, igrp*, or *eigrp*), the command is essentially ignored. It is not uncommon to see this mistakenly used in an ineffective attempt to filter routes from indirectly connected networks. Therefore, check your assumptions.

Unguaranteed BGP Network Origination

The **network <subnet> mask <mask>** configuration command is used by BGP to indicate its Autonomous System is the origin of the (sub)net. The BGP process will *only* originate the mentioned (sub)net when it is known to the router, whether connected, static, or learned dynamically.

This check verifies the (sub)net is either directly connected or has a static route. It is beyond the scope of this check to look at dynamic routes, so this check might be *noisy*. In any case, it is good practice to employ static routes to guarantee the origination.

Unknown Address in Dialer Mapping

To use the **dialer map <protocol> <protocol-address>** command, the protocol address specified must map to a known router interface. This warning is issued when the protocol address does not map to an interface.

Unknown Address in Frame Relay Mapping

To use the **frame-relay map <protocol> <protocol-address>** command, the protocol address specified must map to a known router interface. When it does not, this warning is issued.

Unknown Address in SMDS Mapping

To use the **smds static-map <protocol> <protocol-address> <smds-address>** command, both the protocol address and smds address parameters must map to the same known router interface. When it does not, this warning is issued.

Unknown Address in X25 Mapping

To use the **x25 map <protocol> <protocol-address> <x.121-address>** command, both the protocol address and x.21 address parameters must map to the same known router interface. This warning is issued when the protocol address does not map to an interface. When it does map to a known interface, a high severity violation occurs when the x.121 address has not also been assigned to the same interface.

Unknown EBGP Neighbor

When an IP address mentioned in the list of IBGP neighbors/peers does not map to a known router interface, this warning is issued if a known router is also in the remote AS. When no known router belongs to the mentioned AS, a warning is not issued, as this could correspond to an Internet connection. This may simply be the result of an incomplete collection of configurations.

Unthrottled Redistribution of Routes

It is good practice to apply a route filter whenever routes are redistributed between dynamic routing protocols. When back doors exist, routes can *leak* back, thus creating potential routing loops.

Unused “distance” Command

The **distance <weight> <ip-source-address> <ip-mask>** commands are sequentially evaluated to assign an administrative distance to the routes received from another router. The first matching distance determines the weight assigned to a learned route. When a preceding command is more general than a subsequent command, the latter command is never exercised. This may not be what you want.

Unused “eigrp summary-address” Command

The **ip summary-address eigrp <autonomous-system>** commands on an interface are sequentially evaluated to determine summarization of advertised EIGRP routes. The first matching rule, with respect to an autonomous system, is used. When a preceding command is more general than another command, the subsequent command is never exercised. This may not be what you want.

Unused OSPF “area range” Command

OSPF ranges are address/mask pairs that let you group subnetted networks residing in the same area. The router generates a single network summary advertisement for the group. The OSPF range in the list is sequentially evaluated to determine how to summarize between areas. The first range net and mask that match, with respect to an area, are used. When a preceding range is more general than a subsequent range, the subsequent range is never exercised. This may not be what you want.

Unused OSPF “network area” Command

The **network area** OSPF commands are sequentially evaluated to determine which area an interface is to be assigned. This check flags area definitions not matching any of the router's interface addresses. This check also flags area definitions made useless by virtue of a preceding area definition. The first matching rule determines the assigned area. Thus, when a preceding command is more general than a subsequent command, the latter command is never exercised. This may not be what you want.

Unused OSPF “summary address” Command

The OSPF **summary-address** commands are sequentially evaluated to determine how to summarize within an area. The first matching rule is used. When a preceding command is more general than subsequent command, the latter command is never exercised. This may not be what you want.

Useless BGP Network Origination

The **network <(sub)net_address> mask <mask>** configuration command is used by BGP to indicate its Autonomous System is the origin of the (sub)net. When the specified mask masks out any bits in the specified (sub)net address, this command is ignored by the router. These commands are also filtered out of the Cisco IP Manager database. For example, *network 131.108.1.0 mask 255.0.0.0* is ignored by the router.

IPX

This section describes the IPX integrity checks performed. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

IPX High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following IPX High Severity problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

Duplicate Address Check

All IPX network/host address pairs assigned to routers are checked for duplication. Each router running IPX, is assigned a host address. This check verifies the IPX network/host-address pairs are unique across all routers.

IPX Logical Topology Out of Synch

When two router interfaces are logically connected via one protocol but do not share any IPX networks in common, the mixed logical topologies are inconsistent. Check your addresses. The IP, IPX, and AppleTalk views should overlay consistently upon one another.

IPX Network Encapsulation Mismatch

Interfaces connected via IPX must agree on the respective Novell encapsulations assigned to each Novell network. Secondary Novell networks can be assigned to an interface, provided different Novell encapsulations are allocated for each Novell network. Each IPX network on an interface must have a unique encapsulation and these network/encapsulation assignments must agree across all connected interfaces.

Non-utilized Rule in Access List

The **access-list <access-list-number> {deny|permit}** commands specify a collection of rules which can be utilized for IPX routing table, SAP, and generic output filtering. The rules are evaluated sequentially looking for the *first* match. When a match is found, the packet or SAP/routing update is either denied/permitted according to the rule. The use of wild-cards (such as -1 network values, 0 socket/protocol values) and masking enables the creation of very general rules. This check locates all preceding rules in access lists which shadow (or are more general than) subsequent rules. When the actions differ (such as permit vs. deny) between the two, the intent for the latter rule is not met (since the rule is never exercised.) This is classified as a high severity problem. When the action flag is the same, a simple redundancy situation exists. This is classified as a warning-level problem.

Same IPX Net Restriction

Two interfaces on the same router cannot be assigned the same IPX network. The router does not allow such a configuration.

Undefined Access List Referenced

The Novell **ipx access-group, ipx input-network-filter, ipx output-network-filter, ipx input-sap-filter,** and **ipx output-sap-filter** commands require a reference to an access list. This check identifies references to undefined access lists. An empty access list has the implicit behavior of denying all access.

IPX Warning-Level Checks

When a potential problem is uncovered while performing the following checks, one of the following IPX Warning Level problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

IPX Delay Mismatch Amongst Connected Interfaces

The IPX delay settings amongst connected interfaces must agree.

IPX Update Interval Mismatch Amongst Connected Interfaces

The IPX update interval settings amongst connected interfaces must agree.

Non-utilized Rule in Access List

The **access-list <access-list-number> {deny|permit}** commands specify a collection of rules which can be utilized for IPX routing table, SAP, and generic output filtering. The rules are evaluated sequentially looking for the *first* match. When a match is found, the packet, SAP update, or routing update is either denied/permitted according to the rule. The use of wild-cards (such as -1 network values, 0 socket/protocol values, and masking) enables the creation of very general rules. This check locates all preceding rules in access lists which shadow (or are more general than) subsequent rules. When the actions differ (such as permit vs. deny) between the two, the intent for the latter rule is not met (since the rule is never exercised.) This is classified as a high severity problem. When the action flag is the same, a simple redundancy situation exists. This is classified as a warning-level problem.

IPX SAP Update Interval Mismatch Amongst Connected Interfaces

The IPX SAP update interval settings amongst connected interfaces must agree.

Remote Source Route Bridging (SRB)

This section describes the Remote Source Route Bridging (RSRB) integrity checks performed. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

SRB High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following SRB High Severity problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

Local SRB Peer Definition Problem

Each router in a remote source-bridged ring group must specify a local interface address as its stop on the virtual ring. For Fast Sequenced Transport (FST) encapsulation, the **source-bridge fst-peername <local-interface-address>** command must be used to define the local address. For Transport Control Protocol (TCP) encapsulation, a **source-bridge remote-peer <ring-group> tcp <ip-address>** command must exist to define the local address.

Multiple SRB Remote Peer References to a Router

The **source-bridge remote-peer <ring-group> {fst|tcp} <ip-address> [version n]** command requires a reference to a router interface address which denotes a stop on the virtual ring. Only one remote peer relationship, per ring group, per router is allowed.

Redundant Source Route Bridging Interface

The **source-bridge <local-ring> <bridge> <target-ring>** command is used to configure an interface for source route bridging. It is illegal for two separate interfaces to be configured using the same combination of local-ring, bridge, and target-ring settings.

Referenced Remote Peer is Not the Local Peer

The **source-bridge remote-peer <ring-group> {fst|tcp} <ip-address> [version <n>]** command requires a reference to a router interface address which denotes a stop on the virtual ring. When the specified address resolves to a known router interface address, but that address is *not* the local peer on that router, a problem exists. Note this became a requirement for IOS 10.x releases, but is not a problem for IOS 9.x releases.

SRB Remote Peers Encapsulation Mismatch

The **source-bridge remote-peer <ring-group> {fst|tcp} <ip-address> [version <n>]** command requires an encapsulation type (fst or tcp) and optionally, a version tag. When the specified address resolves to a known router interface address, but the remote peer does not agree on the encapsulation attribute, this constitutes a potential high severity problem.

Unbalanced SRB Remote Peers

The "**source-bridge remote-peer <ring-group> {fst|tcp} <ip-address> [version <n>]** command requires a reference to a router interface address which denotes a stop on the virtual ring. When the specified address resolves to a known router interface address, but the remote router does not have a balancing **source-bridge remote-peer** command pointing back to the first router's peer address, an unbalanced remote peer relationship exists.

Unresolved DLSw Local Peer Address Referenced

The **dlsw <local-peer> <peer-id> <IP-address>** command requires a reference to a local router interface address. When the specified address does *not* resolve to a known interface address on the router, this problem report is issued.

SRB Warning-Level Checks

When a potential problem is uncovered while performing the following checks, one of the following SRB Warning-Level problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

Unresolved DLSw Remote Peer Address Referenced

The **dlsw <remote-peer> {tcp|fst} <IP-address>** command requires a reference to a remote router interface address. When the specified address does *not* resolve to a known interface address on the router, this warning is issued.

Unresolved SRB Remote Peer Address Referenced

The **source-bridge remote-peer <ring-group> {fst|tcp} <ip-address> [version <n>]** command requires a reference to a router interface address which denotes a stop on the virtual ring. When the specified address does *not* resolve to a known router interface address, this warning is issued.

SNA STUN

This section describes the SNA STUN integrity checks performed. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

SNA STUN High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following SNA STUN High Severity problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

STUN Route Does Not Reference STUN Peername

When a STUN connection uses TCP/IP encapsulation, the IP addresses of the **stun route address** and **stun route all** interface configuration commands must match the IP addresses of the complementary **stun peer-name** global configuration commands.

STUN Needs Local Peername

The **stun peer-name** command must reference one of the IP addresses on the router.

SNA STUN Warning-Level Check

When a potential problem is uncovered while performing the following checks, one of the following SNA STUN Warning-Level problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

STUN Route References Unknown Address

The peer IP address specified in a **stun route address** or **stun route all** interface configuration command does not match any known IP address.

Token Ring

This section describes the Token Ring integrity checks performed. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

Token Ring High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following Token Ring High Severity problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

Ring Numbering Causes Addressing Inconsistencies

When two interfaces have been assigned the same ring number, they should also be connected via IP, IPX, or AppleTalk when those protocols are configured on the two interfaces. When this is not the case, there is a high probability a duplicate ring number exists within the topology.

Ring Speed Mismatch Between Connected Interfaces

The ring speed of all connected Token Ring interfaces must match. When they do not, this constitutes a high severity problem.

AppleTalk

This section describes the AppleTalk integrity checks performed. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

AppleTalk High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following AppleTalk High Severity problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

AppleTalk Zone Misconfiguration

All router interfaces connected to the same AppleTalk segment (cable-range) must have their AppleTalk zones in agreement. The primary zones must match on all router interfaces connected to a given cable-rage. All secondary zones must match on all interfaces on a given cable-range.

AppleTalk Logical Topology Out of Synch

When two router interfaces are logically connected via one protocol but do not have common AppleTalk cable-ranges, the mixed logical topologies are inconsistent. Check your addresses. The IP, IPX, and AppleTalk views should overlay consistently upon one another.

Overlapping AppleTalk Cable Ranges

A cable range denotes a range of AppleTalk nets connected to an interface. The cable-ranges assigned to two different interfaces must match exactly and must not overlap.

AppleTalk Warning-Level Checks

When a potential problem is uncovered while performing the following check, an AppleTalk Warning-Level problem report is generated and placed in the Integrity Checks report.

Duplicate Address Check

All AppleTalk network/host address pairs assigned to routers are checked for duplication. This check verifies the AppleTalk network/host address pairs are unique across all routers. The router can tolerate redundant AppleTalk addresses in most cases, but it will be a problem in situations where static mappings are employed (such as Frame Relay, SMDS, X.25, or dialer maps).

DECnet

This section describes the DECnet integrity checks performed. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

DECnet High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following DECnet High Severity problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

DECnet Cost Mismatch Among Connected Interfaces

The DECnet cost statement enables DECnet routing over an interface. All router interfaces on the same cable must share the same value.

Duplicate Address Check

All DECnet area/node address pairs assigned to routers are checked for duplication. This check verifies these addresses are unique across all routers.

VINES

This section describes the VINES integrity checks performed. High severity checks are discussed first, followed by warning level checks. Individual checks are alphabetized within categories.

VINES High Severity Checks

When a potential problem is uncovered while performing the following checks, one of the following VINES High Severity problem entries is placed in the **Integrity Checks** report and the accompanying text displayed when the **Explanation** button is clicked:

VINES Metric Mismatch Among Connected Interfaces

The VINES delay metric settings amongst connected interfaces must agree. This command enables VINES routing on an interface and affects the delay of routing updates.

VINES Serverless Mismatch Among Connected Interfaces

The **vines serverless** command is used to indicate the presence of a Banyan VINES network without a server. This command can be used on several routes to build a path to a network containing servers. All connected interfaces must agree with respect to this setting.