

Title: Use of the HIERARCHICAL RECORDING Attribute of IDRP
Source: IBM
Reference: SC6 N6387 (IDRP)

One of the features of IDRP is the ability to control the transitivity of externally generated NPDUs through a routeing domain. SC6 N6387 provides this functionality through the DIST_LIST_INCL and DIST_LIST_EXCL path attributes. These attributes explicitly enumerate RDs that can (can not) receive the routeing information contained in a given UPDATE PDU. Since NPDUs flow in a direction opposite to the flow of UPDATE PDUs, this provides a means for a given RD to control use of its resources by other RDs. That is, by controlling the advertisement of its resources, a given RD indirectly controls the flow of user traffic.

IBM believes that it is possible to also use the HIERARCHICAL RECORDING attribute as a second means by which transitivity through a routeing domain can be controlled, and offers the following outline for discussion. The proposed method is offered as an additional way for controlling transitivity. If the transitivity permitted by the proposed methods is not suitable for a particular RD, that RD still has the option of using the the currently defined methods (DIST_LISTs).

Proposed Additional Method:

If the Routeing Domain Confederations are supported, then the associated hierarchical relationships provide an opportunity for defining another mechanism for controlling transitivity. The mechanism will impose transitivity constraints on traffic flow within a given RDC based on the value of the HIERARCHICAL RECORDING attribute and the knowledge of whether it is necessary to enter or exit a confederation in order to reach an adjacent RD (See SC6 N6387, clause 7.17.4).

Proposed New Attribute Usage Rules:

The HIERARCHICAL RECORDING attribute is used to control the destinations to which an UPDATE PDU may be propagated. This attribute is redefined to be a single octet in length, and its allowed values are 0 or 1.

If an RD wants to utilize partial order as a mechanism to control its transitivity, BISs within that RD should append the Hierarchical Recording path attribute (if not present) to the routeing information they send, according to the following rules:

1. Originating Information

The default value of the HIERARCHICAL RECORDING attribute shall be set to 1 by the originator of the routeing information.

2. Propagating UPDATE PDUs

If a given BIS receives an inbound UPDATE PDU whose HIERARCHICAL RECORDING attribute is equal to 1, it may propagate the routeing information to BISs in any adjacent RD:

-
- a. If it is necessary to enter a confederation in order to reach the destination BIS, then the advertising BIS shall set the HIERARCHICAL RECORDING attribute of the outbound UPDATE PDU to 0.
 - b. If the destination BIS can be reached without entering a confederation, the value of the HIERARCHICAL RECORDING attribute shall remain equal to the value contained in the inbound UPDATE PDU.
3. If a given BIS receives an inbound UPDATE PDU whose HIERARCHICAL RECORDING attribute is equal to 0, then it may propagate the routing information only to those BISs that can be reached without exiting any confederation to which the advertising BIS belongs. The value of the HIERARCHICAL RECORDING attribute in the outbound UPDATE PDU shall be set to 0.

Transitivity Constraints

The rules listed above will support several general policies, for example:

- A BIS may freely propagate an UPDATE PDU to any adjacent RD that can be reached without exiting any confederation to which the advertising BIS belongs
- Systems located in a region where confederations overlap cannot be used to propagate UPDATE PDUs between systems located in the non-overlapped regions
- RDs that form a confederation can not be used as a transit for RDs outside the confederation.

Figure Figure 1 on page 3 will be used to illustrate the effect of the attribute usage rules. This figure shows 12 routing domains which have been organized into 7 confederations, 6 of which are nested inside RDC-1. It also shows the links between the routing domains.

Consider the following examples:

- Consider an UPDATE PDU that is originated by RD-A and travels to RD-C. The rules permit RD-A to advertise the update to RD-C, and since RD-C can be reached without entering a new RDC, the HIERARCHICAL RECORDING attribute remain equal to 1. Routing information that corresponds to traversing a set of nested confederations from the inside out is permitted.
- Consider an UPDATE PDU that is originated by RD-G and travels along a path <RD-G, RD-F>. RD-G is permitted to advertise to RD-F, and it sets the HIERARCHICAL RECORDING attribute to a value of 1, because RDC-3 will be entered. Routing information that corresponds to traversing a set of nested confederations from the outside in is permitted.
- Consider an UPDATE PDU originated by RD-B and travelling to RD-I along path <RD-B, RD-J, RD-I>. When RD-B advertises to RD-J, the HIERARCHICAL attribute will have a value of 1, because it was not necessary to enter any new RDC in going from B to J. When J subsequently advertises to I, it will set the HIERARCHICAL RECORDING attribute to 0, because RDC-3 and RDC-7 are entered. Routing information flow that corresponds to traversing a set of nested confederations from the inside out, and then traversing a different set of nested confederations from the outside in is permitted.
- Consider a potential UPDATE PDU originated by RD-J. It cannot travel along the path <RD-J, RD-I, RD-H>. The UPDATE PDU that arrives at I will have a HIERARCHICAL RECORDING attribute with value 0; therefore this update can not be advertised to H since it would be necessary to exit RDC-7 in order to reach RD-H. Thus, the proposed method provides a simple mechanism to control the transitivity of the overlap: that is, it allows the RDCs that form the overlap to constrain

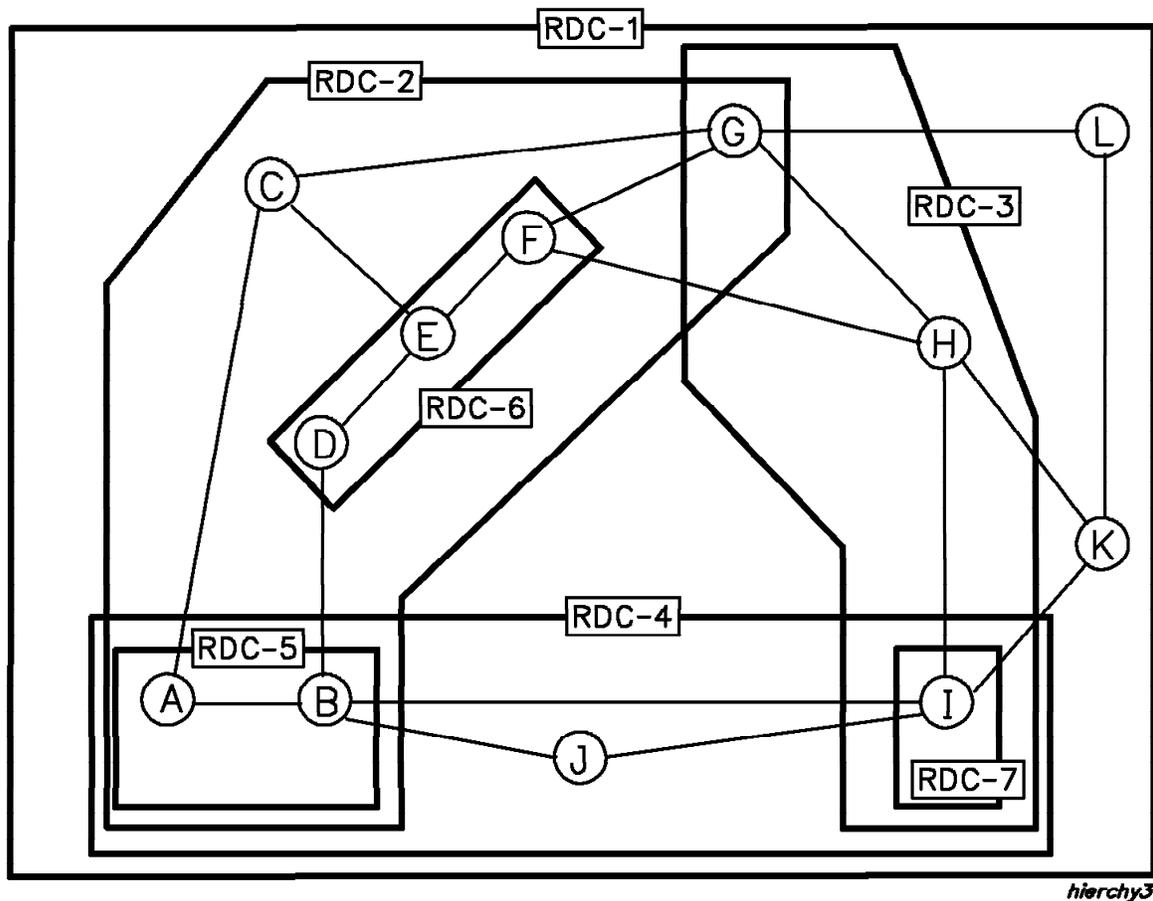


Figure 1. Transitivity and Nesting Relationships

their transitivity in such a way that traffic between the non-overlapping parts of the hierarchy could not traverse the region of the overlap.

- RD-D can advertise any UPDATE PDU to RD-E and RD-F regardless of who originated it or what the value of the HIERARCHICAL RECORDING attribute was. This is so because both of these destinations can be reached without exiting RDC-6. Note also that the value of the attribute will remain unchanged because it is not necessary to enter a confederation anywhere along this path.
- It would not be permitted to advertise an UPDATE PDU along a path <RD-C, RD-A, RD-B, RD-D>. The path segment from C to A is allowed, and A will receive an UPDATE PDU with a HIERARCHICAL RECORDING attribute of 0, because it was necessary to enter confederation RDC-5. Flow from A to B is permitted also. However, flow from B to D is not allowed because the incoming value of the HIERARCHICAL RECORDING attribute was 0, and therefore the UPDATE PDU can not exit RDC-5.

Complementary Operation with Distribution Lists

Hierarchical-based rules can be complemented with rules that govern propagation of the routing information in presence of the DIST_LIST_INCL and DIST_LIST_EXCL path attributes. In certain cases, the constraints on propagation of UPDATE PDUs which are imposed by the HIERARCHICAL RECORDING attribute may not be sufficiently broad to support all of the policies that a given RD wishes to implement. Therefore, an RD is free to include the DIST_LIST_INCL(EXCL) attribute in its UPDATE PDUs as well. In such cases, the contents of the DIST_LIST_EXCL (or DIST_LIST_EXCL) shall take precedence over the propagation rules based on the HIERARCHICAL RECORDING attribute.

When aggregating routes (see Section 7.12.9), if at least one of the routes to be aggregated contains the Hierarchical Recording path attribute, then the aggregated route contains the Hierarchical Recording path attribute as well (with the value of the aggregated route). If several routes to be aggregated contain the Hierarchical Recording path attribute with different values, the value of the Hierarchical Recording attribute of the aggregated route is set to 0.

When announcing unfeasible routes, the Hierarchical Recording path attribute may be omitted. However, propagation of the unfeasible routes is governed by the rules associated with the corresponding feasible routes. Presence of the Hierarchical Recording path attribute in the unfeasible route is ignored.

Attribute Encoding:

Since the HIERARCHICAL RECORDING attribute is now a single 1-octet field, there is no need to maintain the semantics of either a "level" or an "Up/Down Flag".