digital™

Networks Marketing

OSPF versus IS-IS

A technical discussion of the differences between these link state routing protocols.

Document Number: NIPG-004-01 December, 1993

Abstract

Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) are two routing protocols designed to support routing in a Network Layer Protocol with datagram service. OSPF was designed for IP, the Network Layer protocol in the TCP/IP protocol suite while IS-IS was originally designed for CLNP, the datagram Network Layer protocol in the ISO suite. IS-IS can also support the IP Network Layer as well as the ISO Network Layer.

There has been much debate within IP community regarding the OSPF and IS-IS routing protocols. Much of the debate has focused on which is the superior routing protocol. However, there has also been much debate about "integrated" versus "Ships-in-the-Night" routing. Integrated routing uses a single routing protocol to support multiple Network Layer protocols, while the alternative Ships in the Night uses a separate routing protocol for each Network Layer protocol.

This paper discusses the differences between the OSPF and IS-IS routing protocols. Integrated routing versus Ships-in-the-Night routing is discussed in *Routing Protocols*, NIPG Marketing Document Number NIPG-002.

The information in this document is based on a paper written by Radia Perlman whilst a senior architect within Digital's Networks Engineering organisation.

Introduction

There are two types of routing protocols in use in networks today. One type is known as "distance vector" routing. The other type is known as "link state routing". Both OSPF and IS-IS are link state routing protocols.

In a link state routing protocol, each router is responsible for determining the identity of its neighbours constructing a special packet known as a Link State Packet (LSP) listing the node's neighbours, broadcasting the LSP to all other routers, storing the most recently generated LSP from each other router, and computing routes to all the destinations based on the stored LSP database.

Both OSPF and IS-IS support "hierarchical routing". Routers can only support a limited size of network. If a network grows beyond the size that can practically be supported by the routers, the common technique is to add hierarchy to the network. That means that the network is partitioned into pieces known as "areas". Level 1 routing concerns itself with all the individual nodes and links within an area. Level 2 routing concerns itself with routing between areas, but does not concern itself with the detailed structure inside the areas.

OSPF evolved from an early version of IS-IS, and it adopted many of the ideas, such as:

- The mechanism for distributing LSPs among the routers
- The concept of electing a single router on a LAN, the "Designated Router", which constructs an LSP listing all the nodes on the LAN. This avoids each router having to individually list all the nodes on the LAN.
- The algorithm for calculating and maintaining multiple equivalent paths to a destination, for the purpose of spreading traffic load

The remainder of this section discusses the differences between IS-IS and OSPF.

Packet Encoding Issues

OSPF runs "on top of" IP. That means that an OSPF packet is transmitted with an IP data packet header. In contrast, an IS-IS packet is transmitted directly on top of the Data Link Layer. The two packet formats are shown in the following diagram.

OSPF packet format

IS-IS packet format

Because the Berkeley UNIX implementation does the Data Link differentiation of packet type in the kernel, the claim is that the IS-IS encoding would force modifications of the kernel so as to recognise another type of packet. It is not necessarily a good idea to design protocols around the artefacts of a particular implementation. There are disadvantages to the extra layer of protocol. It increases packet overhead, because of the extra header. Certain routing protocol packets are high priority and must not be dropped. If the kernel is incapable of distinguishing a routing control packet from a data packet, it will be incapable of ensuring priority service for the critical control traffic.

Avoiding Large Control Packets

Both OSPF and IS-IS were designed to avoid the need for control packets to be so large that they would require conventional fragmentation and reassembly. In both protocols, the functionality of a large packet is expressed in several smaller packets, each one of which can be processed independently. With traditional fragmentation and reassembly, all fragments must arrive in order for a packet to be processed. The special mechanisms in both protocols give the following advantages:

- LSPs are propagated more quickly. A fragment can be forwarded as soon as it arrives. If a large LSP were fragmented and reassembled at each hop, the entire LSP would need to be reassembled before the LSP could be forwarded.

- Less bandwidth is required. With the special mechanisms in the two protocols if a single fragment is lost, only that fragment need be retransmitted. With traditional fragmentation and reassembled, if a single fragment is lost, all fragments must be retransmitted.

- It is more robust. If there were a very large packet and an error prone link, the probability of all fragments succeeding would be very low. With traditional fragmentation and reassembly, all fragments from a particular transmission attempt must arrive in order for the transmission to be successful. It is not possible to reassemble fragments from one transmission of a packet with fragments from a later transmission of the same packet.

The mechanisms used by the two protocols for avoiding large control packets are very different. There are two types of packets that can become large enough for fragmentation to be an issue. The first type of packet is an LSP. The second type of packet is called a "Complete Sequence Numbers Packet" in IS-IS, and is roughly analogous to the "Database Description Packet" in OSPF.

Link State Packets (LSPs)

In IS-IS, LSPs list all the neighbours of a node. In the case of the LSP which lists all the nodes on a LAN, the LSP can become very large. Even in the absence of LANs it is possible, but less likely, for a router to have so many links that the neighbour list cannot fit into a single packet.

IS-IS avoids large LSPs by allowing a router *R* to fragment its LSP. The "source" address of an LSP is the router's ID concatenated with an additional octet which indicates the "fragment number". Thus if *R* has an LSP that requires fragmentation into 6 pieces, *R* will generate LSPs from *R.0*, *R.1*, *R.2*, *R.3*, *R.4*, and *R.5*. For the purpose of propagation of the LSP fragments, each fragment is independently propagated and has an independent sequence number. If a few links change, then *R* only needs to reissue the fragments with out of data information.

OSPF avoids large LSPs by having router *R* report different types of information in different LSPs. There are 5 types of LSPs that *R* might report:

1. Type 1 "Router Links". These report router neighbours of *R*, as well as the IP address of each of the LANs to which *R* attaches.

2. Type 2 "Network Links". This is generated by the Designated Router on a LAN and reports all the routers attached to the LAN.

3. Type 3 "IP Network Summary Links". This is generated by a Level 2 router *R*, and reports the cost from *R* to an IP address outside the area, but within the "Autonomous System" (AS). (An Autonomous System is a portion of the network that independently runs an "Interior routing protocol". Routing between Autonomous systems requires another type of routing protocol, known in the IP community as an "Exterior routing protocol". The ISO community refers to an Autonomous System as a "routing domain".) Note that each IP address within the AS is reported in a separate type 3 LSP in OSPF.

4. Type 4 "AS Border Router Summary Link". This is generated by a Level 2 router *R*, and reports the cost from *R* to an AS Border router. (An AS Border

router is one that has information regarding IP addresses in other ASs, acquired either manually or through an Exterior Routing Protocol.) Note that each AS Border Router is reported in a separate type 4 LSP, by each level 2 router in the area.

5. Type 5 "AS External Link" are generated by AS Border Routers, and describe destinations reachable outside the AS. Each destination must be reported in a separate type 5 LSP.

The encoding chosen for OSPF is quite memory and bandwidth intensive, since each LSP requires certain overhead information, such as sequence number, age, ID of router that generated the information, etc. In OSPF type 3, 4, and 5 LSPs can each hold only a single destination. Therefore, in OSPF every destination reachable outside an area requires its own LSP header to be stored and propagated by each router. In contrast, in IS-IS, as many destinations as can fit are all carried in a single LSP, with its own overhead information. In IS-IS an LSP can carry on the order of 100 destinations. In OSPF, each of the destinations requires its own overhead information, so the storage and bandwidth required is about an order of magnitude greater.

OSPF does allow multiple destinations to be reported in type 1 and 2 LSPs. These are less likely to be very large. But if a router had many links, or a LAN had hundreds of routers on it, then it is possible for a type 1 or 2 LSP to required fragmentation in OSPF. In this case, OSPF would utilise the traditional hop by hop fragmentation and reassembly provided by the IP header, with the disadvantages discussed above.

Database Description/CSNPs

Both OSPF and IS-IS require mechanisms to enable neighbour routers to compare their LSP databases. This is done by exchanging the headers of the LSPs. Especially in OSPF where every destination outside an area requires its own LSP header, the information required to summarise the LSP database can be quite large. But in IS-IS it is likely to also be larger than will fit into a single packet.

OSPF deals with the problem by fairly conventional fragmentation. Each fragment has a sequence number, and the neighbour router must acknowledge a fragment before the next fragment can be transmitted. IS-IS deals with the problem by including in a CSNP, the beginning and end LSP source addresses contained in the CSNP. In this way each CSNP can be processed independently in IS-IS.

Variable Length Fields

Another difference between OSPF and IS-IS is that all fields in OSPF are fixed length, and the packet formats specify which fields are present. In contrast, most fields in IS-IS are variable length, and are encoded as:

Fixed format packets (as in OSPF) are both an advantage and disadvantage. It makes it easier for the routers to parse control packets. However, it precludes extendibility of the protocol.

IS-IS allows new fields to be defined in downwardly compatible way. A router that does not support a particular type number merely skips that field and processes the rest of the packet normally. Because of this encoding it was easy to modify IS-IS to support IP addresses as well as ISO addresses. In contrast, OSPF cannot be extended to support new features. There is no way to add new types of fields, and there is no way to change the length of fields. In particular, OSPF assumes addresses require 8 octets to specify, since IP addresses require 8 octets (4 for the address, 4 for the mask).

ISO addresses require 21 octets (up to 20 for the address plus one for the length). So whereas IS-IS was easily enhanced to support IP, OSPF would not be capable of being enhanced to support OSI, or even extended length IP addresses.

Area Partitions

With hierarchical routing, it is possible for an area to partition so that level 1 routing cannot connect the partitions. If both partitions contain level 2 routers, and the level 2 network is connected, then the network as a whole is not physically partitioned. There is a path between the partitions of the area. The path is a level 2 path.

The symptoms of a partitioned area can be difficult to diagnose and annoying for the users. Not only is communication impossible between nodes that should be in the same area but are currently in different partitions of the area (as in nodes B and C in the diagram above), but communication between members of the area and nodes outside the area can be disrupted since the traffic into the area might enter the wrong partition and be undeliverable.

IS-IS has mechanisms in which level 2 routers residing in a partitioned area automatically detect the partition and repair the partition by utilising the level 2 path as a level 1 link.

With ISO addressing the area number is contained in an address. In IP addressing, an address is too small to contain an area number. When IS-IS is used to route IP, although the IP addresses do not contain area numbers, routers do have ISO style addresses.

In both OSPF and IS-IS, the level 2 routers in an area report all the IP addresses within the area as being reachable outside the area. In order to compact the information, level 2 routers are manually configured with "address summaries" for the area. An address summary is an (IP address, subnet mask) pair that describes many IP subnets in the area. In IS-IS the rule is that a level 2 router reports an address summary for the area provided that at least one address in the area matches the address summary. OSPF does not have any explicit area repair mechanisms.

However, the specification claims that the area merely becomes two areas. There are only two methods by which this claim can be accomplished:

1. A network manager might notice that the area has partitioned, and manually reconfigure the level 2 routers in the area, so that level 2 routers in each partition do not contain summary addresses for addresses reachable in other partitions.

2. No summary addresses were used, and each level 2 router reports each IP address individually. If summary addresses are not used, then areas do not become partitioned. They merely break into multiple areas.

The OSPF specification is not clear about which of these alternatives is assumed. Reliance on a human is clearly undesirable. The network remains broken until an appropriate human can be found, the human diagnoses the problem, calculates appropriate summary addresses, and reconfigures the routers. The second alternative is also unappealing, since without summary addresses there really is no savings gained from hierarchical routing. The routing algorithm and the capacity of the routers is unlikely to accommodate a network without hierarchy.

Level 2 Partitions

IS-IS requires a connected level 2 network. This means that there must be a path from every level 2 router to every other level 2 router that traverses only level 2 routers.

OSPF similarly requires a connected level 2 network, but allows a link between a pair of level 2 routers to consist of a manually configured "virtual link" which consists of a path through an area.

Communication over a virtual link between level 2 routers *A* and *B* can be done in two ways:

1. *A* can encapsulate traffic being forwarded to *B* in a Network Layer header giving "*B*" as the destination.

2. *A* can assume that all the level 1 routers on the path towards *B* know enough to forward traffic to the destination towards *B*.

OSPF has chosen the 2nd strategy. Because of this, it is essential that all the level 1 routers within the area know about all destinations outside the area, so that they will be able to forward level 2 traffic.

IS-IS does not mention the possibility of a virtual level 2 link. However it is straightforward to provide the same functionality in a vendor proprietary manner. All that is required is that two level 2 routers both implement the ability to have a virtual level 2 link configured. A virtual level 2 link does not affect any of the other routers in the routing domain. In IS-IS level 2 information is not given to the level 1 routers, so encapsulation would be required for communication over the virtual link.

If it is considered useful to provide the ability for a network manager to notice a level 2 network partition and react by configuring a virtual link, vendors will provide that functionality in their IS-IS level 2 routers. Most likely once a human has diagnosed the problem it will be just as easy to get the physical link which broke repaired, deploy a new physical link, or reconfigure enough level 1 routers to be level 2 routers in order to repair the level 2 network partition.

Injection of Level 2 Information

In IS-IS, level 1 routers know only the information about their own area. If a level 1 router *R* receives a packet with an address that is not reachable within the area, *R* forwards the packet to the level 2 router nearest to *R*.

In OSPF, level 2 information is fed into the area. The form of the level 2 information is interesting. Suppose there is an area *A* in some AS such that:

1. *n* IP destination addresses are reachable within the AS but outside of area *A.*

2. *m* IP destinations reachable outside the AS.

3. *k* level 2 routers in area *A*.

4. *l* AS border routers in the AS.

Each of the *k* level 2 routers reports their own distance to the *n* IP destination addresses and the *l* AS border routers. This information is *O(k * (l+n))*. Also, each of the *l* border routers reports its distance to each of the *m* IP destinations reachable outside the AS. This information is *O(l*m).*

Giving level 2 information to level 1 routers enables them to choose the exit level 2 router which will give the best path to the destination. Thus OSPF yields more optimal interarea routes than IS-IS. The cost of providing more optimal routing, though, is increased bandwidth usage by the routing algorithm (to pass level 2 information around in the area), and increased memory and CPU requirements on level 1 routers (to store information and compute routes for out of area destinations).

Aside from increased bandwidth, CPU, and memory usage, the OSPF requirement for level 1 routers to store level 2 information has an additional problem. In IS-IS, where an area is independent of the rest of the network, database sizes in level 1 routers can be calculated based on the size of the area. If the area never changes, the level 1 routers will be able to function. In contrast, in OSPF, as the entire network grows, the demand on the level 1 routers increases. One small area, with small routers, cannot be sheltered from the growth of the rest of the network.

There is an additional option in OSPF. It is possible, if there is only one level 2 router in an area, to configure the area as a "stub area". In order to do so, the area must meet the topological constraints of a stub area, and every router in the

area must be configured to assume the area is a stub area. If an area is a stub area, then the type 5 LSPs are not flooded into the area. This saves the O(l*m) information in the type 5 LSPs. But surprisingly, type 3 and 4 LSPs are still flooded within an area, even if the area is configured to be a stub area.

Authentication

Both OSPF and IS-IS have the capability of using "authentication". Theoretically, authentication gives some sort of security against hackers programming their nodes to masquerade as routers and confuse the network by sending bad information. But neither protocol has specified an authentication scheme that gives real security. Mostly the authentication schemes specified offer two things:

1. Hooks in the packet formats so that in the future if someone invents a wonderful authentication scheme it can be compatibly slipped into the protocol

2. Marginal insurance against someone accidentally plugging routers together that should not be plugged together.

Both protocols have a field indicating the "type" of authentication, and both only define one type, which is a simple password. There are differences in the two protocols however.

In IS-IS the data associated with the authentication is variable length. In OSPF it is fixed at 64 bits. 64 bits is sufficient for a password scheme but a scheme such as a public key signature scheme which might offer real security would need a field several hundreds of bits long. Thus it would be possible using IS-IS to add a secure authentication scheme in the future, whereas the OSPF scheme would not be able to accommodate a public key cryptographic scheme.

In OSPF, there is a single password per link. A router is configured with a password for each link to which it is attached. It transmits that password when it transmits OSPF messages on that link. It expects all OSPF messages it receives on that link to have that password. In IS-IS a router is configured with a transmit password on a link, which is the password it uses when it transmits IS-IS messages, as well as a set of acceptable receive passwords.

On a point to point link a password scheme in which the receive password and transmit password are different offers some security. If the passwords are the same then the intruder need only wait for the other router to transmit first, and then the intruder will find out the password.

The reason IS-IS configures routers with a set of acceptable receive passwords is so that a link such as a LAN can be migrated from one password to another without disrupting the network. If the routers use password P1, and the network manager would like to change the password to P2, then the network manager adds the password P2 to the configured receive password set of each of the routers, one router at a time. After all the routers have been configured to accept P2 as well as P1, the routers can be configured to transmit P2 instead of P1. After all the routers transmit P2, P1 can be deleted from their configured receive password set, one router at a time.

Since OSPF has only a single password per link it is not possible to change the password in an operational network. The routers would have to all be brought down and locally reconfigured.

Another interesting difference between the authentication schemes in OSPF and IS-IS is the way LSPs are authenticated. In IS-IS the router that originally generated an LSP puts the authentication field into the LSP. This is very dangerous. If a single router in the network has a different configured set of acceptable receive passwords for LSPs, that router will reject an LSP that other routers accept. If routers disagree about the LSP database, routing can be severely disrupted.

In contrast, OSPF does not have the original router put an authentication field into an LSP. Instead, a router that forwards an LSP puts its own authentication field into the packet. This scheme does not suffer from the global disruption possibility that IS-IS suffers from.

Designated Router Election

In IS-IS the Designated Router on a LAN is the router with the highest configured priority, with ID being used as a tie breaker when more than one router has the same highest priority. Routers listen to Hello messages from other routers, and the router with the highest priority/ID is the Designated Router.

In contrast, because of certain mechanisms in OSPF (see "LSP Propagation on LANs", below), it is disruptive in OSPF for a LAN to have a change of Designated Routers. One mechanism used in OSPF to minimise disruption is to elect a Backup Designated Router in addition to a Designated Router. The other mechanism is to have an election protocol for DR and Backup DR that keeps an elected router as DR or Backup DR unless that router fails. A new router with higher priority will not take over from a functioning DR or Backup DR.

The OSPF algorithm consists of requiring a router to wait for some time before "running for" DR or Backup DR. The OSPF Hello includes a field for indicating which router is assumed to be DR, and which router is assumed to be backup DR. If all Hello messages received during the time-out period have those fields empty, the new router will attempt to become Designated Router by putting its own ID into the field in its Hello Messages. If multiple routers happen to run simultaneously, the DR and Backup DR are chosen based on priority and ID.

The OSPF Designated Router election protocol is more complex, and it is not deterministic, which means that it is impossible to predict, based on topology, which router will become DR. This often makes conditions harder to diagnose, because they are not reproducible. And it makes it impossible for the network manager to impose preferences on which routers become DR. If an under powered router which should only take over as a last resort when all other routers have failed ever happens to become DR it will remain DR until it dies.

LSP Propagation on LANs

It is important that when a router generates a new LSP all routers receive the LSP. Unless all routers have identical LSP databases, routing can be severely disrupted. Both OSPF has IS-IS have similar flooding mechanisms for propagating LSPs throughout the network. The main difference in the mechanisms is how LSPs are propagated on a LAN.

In both schemes, it is important for each router to make sure each LSP is properly propagated to each of its neighbours. On point to point links, this is accomplished by transmitting an LSP to a neighbour until an acknowledgement is received from that neighbour for that LSP. The naive extension for LANs is that each router would transmit a copy of each LSP to each neighbour, and collect an acknowledgement for each LSP from each neighbour. This would result in n**2 packet transmissions on the LAN for each LSP. Luckily, neither OSPF nor IS-IS take this approach.

OSPF handles LAN propagation by having the Designated Router on the LAN transmit each LSP, and collect explicit acknowledgements from all the other routers. The protocol is as follows:

1. Some router on the LAN receives an LSP that needs to be forwarded onto the LAN.

2. That router transmits the packet to the DR.

3. The DR multicasts the packet to a Data Link Layer address that is listened to by all routers on the LAN.

4. Each router that receives the LSP transmits an acknowledgement to a Data Link Layer address that is listened to by the DR and the Backup DR.

5. If the DR fails to receive an acknowledgement within some time from some of the routers, it retransmits an individual copy of the LSP to each router from whom it failed to receive an acknowledgement.

IS-IS does not use explicit acknowledgements on the LAN, and any router is allowed to forward an LSP on the LAN to the other routers. It does not have to first get transmitted to the DR. The IS-IS mechanism is simple. A router that receives a new LSP for forwarding onto the LAN marks the LSP in memory as needing to be forwarded onto the LAN. When the router gets a chance, it forwards the packet onto the LAN, to a special Data Link Layer address listened to by all the IS-IS routers. If a router receives an LSP on a LAN before it gets a chance to transmit that LSP on the LAN, the router clears the flag reminding itself to transmit the LSP on the LAN.

To ensure lost LSPs get retransmitted, the IS-IS DR periodically transmits a special packet known as a "Complete Sequence Numbers Packet", which lists all the LSPs and sequence numbers in its LSP database. If a router notices that the DR has missed an LSP, the router will retransmit the lost LSP. If a router R notices, based on the received CSNP, that R has missed an LSP, R explicitly requires the LSP from the DR.

The IS-IS mechanism requires far fewer packet transmissions per LSP. In the absence of lost packets, OSPF requires n packet transmissions per LSP: one for the DR to broadcast the LSP to the other routers, and n-1 for the other routers to transmit acks to the DR. Additionally, if it was a router other than the DR that originally heard the LSP from some other link, there is the packet transmission required to transmit the LSP to the DR. In the absence of lost packets, IS-IS requires only a single packet transmission per LSP. The first router that hears the LSP transmits it on the LAN.

However, OSPF does not have the overhead of the periodic transmission of CSNPs by the DR. If there is no LSP traffic, then OSPF has zero overhead, whereas IS-IS will always have the constant periodic overhead of the CSNPs. The IS-IS mechanism tends to spread out the traffic to a small constant background amount, whereas the OSPF mechanism requires a lot of bandwidth all at once. It is better to have a small constant background overhead than a

burst of traffic followed by zero overhead. Even if only the total number of packets required is used to compare the two schemes, IS-IS is likely to require a lot less bandwidth, because OSPF requires n times as many packets per LSP. Assuming 10 seconds between CSNP transmissions, and 6 routers per LAN, the OSPF mechanism will require more bandwidth if the average amount of time between generation of a new LSP by any router in the network is less than once per minute.

OSPF propagates LSPs throughout the network slightly more slowly than IS-IS, because an LSP cannot be immediately propagated on the LAN. Instead there is an additional packet hop required for initially transmitting the LSP to the DR.
It might seem that OSPF recovers more quickly from lost LSPs than IS-IS, since each LSP is explicitly acknowledged. However, the DR waits some amount of time before assuming an LSP was lost. The periodic time in IS-IS for transmission of CSNPs is equivalent to the DR's time-out interval, so both schemes recover from lost LSPs in the same amount of time.

Synchronised Parameter Setting

In OSPF, there are several parameters that must be configured identically in routers, or else the routers will refuse to communicate with each other. This creates a problem because it is virtually impossible to change the parameter setting via network management. Once a router's parameter setting is changed, it is cut off from the rest of the network since no other routers will be able to communicate with it.

In contrast, a lot of thought has been given in the design of IS-IS to how to modify parameter settings in a running network. In IS-IS there is always a way to migrate from one setting to another by configuring routers one at a time while the network is running.

The parameters in OSPF that must be set identically in neighbouring routers are:

1. HelloTime and DeadTime. The HelloTime indicates how many seconds will elapse between generation of Hello messages by a router. The DeadTime is how many seconds should elapse before the neighbours of the router should consider the router down. OSPF routers report their settings of these parameters inHello messages. If a neighbour router does not have the exact same values as the router that receives the Hello, the router receiving the Hello will refuse to form a link with the transmitting router.

IS-IS reports HoldingTimer in its Hello messages. HoldingTimer is identical with OSPF's DeadTime. There is no reason to transmit the HelloTimer. IS-IS uses the information solely to determine how long to wait between receipt of Hellos from a particular neighbour before declaring the link to that neighbour down. There is no necessity for neighbouring nodes to have the same value.

The reason the decision was made in OSPF to have routers refuse to be neighbours if their HelloTime and DeadTime disagree is that the OSPF designers wanted to make sure that once a Designated Router was elected on a LAN it would remain elected. As will be discussed later, it is very disruptive in OSPF for a new Designated Router to take over on a LAN. If a new router has the DeadTime parameter set to be shorter than the other routers' HelloTime parameter, it might take over as Designated Router, because it would not realise there was already a functioning Designated Router on the LAN.

2. "Stub area" flag. OSPF requires every router in an area to be configured with a flag indicating whether the area is a stub area. If a level 2 router has the stub area flag set, it will not flood type 5 LSPs into the area. The "stub area" flag is reported in OSPF Hello messages. If a router disagrees with a neighbour as to the setting of the "stub area" flag, it will bring down the link to the neighbour.

   IS-IS has no such parameter.

3. Authentication Password. Both OSPF and IS-IS have the optional feature of providing authentication. In OSPF there is a single password per link. The password that a router transmits is the same as the password it will accept on the link.

Neighbour Initialisation Protocol

It is important for routers to have identical LSP databases. This is accomplished in both protocols by ensuring that neighbour routers agree on LSP databases.

In IS-IS, when router *R1* acquires a new router neighbour *R2*, *R1* assumes that *R2* has no LSPs, and *R1* starts transmitting all its LSPs to *R2*. Additionally, *R1* transmits a CSNP to *R2*. Likewise, *R2* will transmit a CSNP to *R1*, but also transmit all its LSPs to *R1*.

If *R2* receives R1's CSNP, *R2* will know which LSPs do not actually need to be transmitted to *R1*. Likewise for *R1*. The CSNP exchange is thus an optimisation to prevent the necessity for transmitting all LSPs on the link.

On a LAN, there is no special protocol for new neighbours. The DR transmits CSNPs periodically. If a new router has missing LSPs, it explicitly requests them from the DR. If the DR is missing LSPs, the other routers transmit their LSPs.

OSPF has a slightly different mechanism. When a link between two routers first comes up, each transmits the equivalent of an IS-IS CSNP. In OSPF it is known as a "Database Description Packet". The difference is that OSPF requires all fragments of a DD packet to be received, or else there might be long term database disagreements. The DD packet is required for correctness, not as an optimisation.

IS-IS additionally has routers on point to point links periodically exchanging CSNPs. This guards against low probability events that might have occurred to corrupt an LSP database without a router's knowledge. OSPF does not have a similar mechanism. The protocol in OSPF for exchanging DD packets is very complex and it would not be practical to go through the procedure periodically.

In OSPF, the DD exchange protocol involves making a decision between the two neighbour routers as to which should be "master" and which "slave". The master transmits each fragment of its DD packet, one at a time. The slave transmits the same numbered fragment of its DD packet in response. It is slow because there is no possibility for pipe lining.

In IS-IS, if a CSNP needs to be fragmented there is no problem. Each fragment of a CSNP states the range of source LSP addresses contained in the CSNP. Fragments can arrive out of order, or fragments can be lost, with no problems. If a fragment of a CSNP is lost in the initial neighbour exchange, the only consequence is that the LSPs that were listed in that fragment will be unnecessarily retransmitted. On a LAN, if a CSNP fragment is lost, assuming LSPs were really lost, the consequence is that it will take an extra CSNP interval for the problem to be noticed and the LSP to be retransmitted.

Exterior Routing Protocol Information

IS-IS has a field in Level 2 LSPs known as "Inter Domain Routing Protocol Information", or IDRPI. The field is ignored by IS-IS routers, but it is there to help the inter domain routers. There are two possible ways in which inclusion of the IDRPI information can help inter domain routers:

1. It helps inter domain routers find other inter domain routers in the domain, so that they can then exchange inter domain routing control traffic.

2. It enables inter domain routers to pass messages to each other without requiring them to explicitly send the information to each other inter domain router in the domain. If there are many inter domain routers, it might be more efficient for them to pass this information in their LSPs, which automatically get flooded to all the other level 2 routers, than to have n**2 conversations going on, between the n inter domain routers in the domain.

This use of the IDRPI field does place a burden on the intra domain routers, because it requires them to store (in the LSP database) the inter domain information.

As with almost all the fields in IS-IS, the IDRPI field is variable length. The IDRPI field in IS-IS was put in because OSPF had a similar field. In OSPF there is a 32 bit "tag" field associated with each destination address reachable outside the routing domain. It is not clearly defined what the purpose of the OSPF tag field is. The reasons given are:

1. The tag might contain the IP address of the inter domain router that advertised the destination. The theory is that this information enables an inter domain router R1 that wishes to forward a packet through the domain to inter domain router R2 to know whether the intra domain routing protocol (OSPF) has stabilised, and will route the packet towards R2. Routing algorithms always take some time to converge. There is no reason for making a special check in this situation.

2. The tag might contain the Autonomous System number from which the information about that destination address was learned. This information might be of use to the obsolescent EGP protocol, since it requires that information learned from AS1 not be transmitted back to another EGP router from AS1. When EGP is replaced by BGP, this field will not be useful for that purpose. The tag field cannot be used like the IDRPI field in IS-IS because it is only 32 bits long. With BGP, there is nothing useful any BGP router would say to another BGP router that would fit in 32 bits.

Database Overload Procedures

There are two reasons why the LSP database might become larger than a router was configured for:

1. Temporary situation. If the DR changes on a LAN, the new DR might issue the large LSP listing all the other nodes on the LAN, before the old DR's LSP times out. Thus for a time other routers will be required to store twice as much information regarding that LAN.

2. Permanent situation. It is possible for the network to grow larger than the routers have been configured for. This may be even more likely in OSPF because the level 1 routers have to keep information about the entire routing domain in OSPF.

There are several strategies an implementation can use when its database overflows. It can:

1. Crash and wait to be manually revived.

2. Crash and immediately reboot.

3. Continue operating based on the subset of information it can hold.

4. Convert itself into an end node so that it can continue to be reachable via network management.

5. Convert itself into an end node, but keep trying to store the LSP database. If it ever succeeds, it reverts back to being a functional router.

Crashing and waiting to be manually revived can be a real annoyance, especially if the problem is a temporary overload situation. A temporary condition might cause all routers in the domain to crash and wait to be manually restarted. Crashing and rebooting will solve the temporary situation, but it does not allow remote management of an under configured router in the case when the overload situation is a permanent situation. Continuing to operate on a subset of the LSP information is very dangerous, since routing can be severely disrupted if routers operate based on different information. Converting into an end node allows the router to be remotely managed. But it requires manual intervention to recover from a temporary situation.

The last strategy is what is specified in IS-IS. A router that must discard LSP information declares itself "overloaded" and sets a bit in its LSP warning other routers that it cannot forward packets. It remains reachable for network management.

It also continues to try to perform the router functions. If a period elapses without the router needing to drop any LSPs, the router declares itself healthy again. In this way a temporary overload situation will resolve without any manual intervention.

In contrast OSPF does not specify any mechanisms to deal with database overload, nor does it have the special flags in the packets to copy the IS-IS scheme. Instead, implementors are free to choose any of the other strategies, such as crashing, or continuing to operate on a subset of the information.

Digital believes the information in this publication is accurate as of its publication date. Such information is subject to change without notice. Digital is not responsible for any errors.

The following are trademarks of Digital Equipment Corporation: Digital, DECnet, and DECNIS.