Network Working Group
Internet Draft
Expiration Date: June 2001
File name: draft-ietf-ospf-floodgates-00.txt

P. Murphy
US Geological Survey
December 2000

OSPF Floodgates

Status of this Memo

Table Of Contents

1.0 Abstract

This memo describes an option to the OSPF Version 2 specification which allows the partial suppression of flooding between neighboring routers. This suppression acts much like a floodgate, opening and closing automatically when triggered by the proper conditions. Two types of flooding suppression are described, 1-Way and 2-Way. The option applies to standard areas, stub areas, and NSSAs. It works over any OSPF interface. Routers with this option configured are backward compatible with routers running an OSPFv2 compliant implementation as defined in RFC 2328, and can be restricted to a subset of the OSPF routing domain. This option is applied only between neighboring OSPF routers.

Please send comments to ospf@discuss.microsoft.com.

2.0 Overview

2.1 Requirement for OSPF Floodgates

Corporate networks in todays modern Internet require a low cost means for maintaining a fault tolerant network. Many of the smaller sites in a corporate network connect via dedicated service which is backed up via traditional analog dialup service or digital ISDN service. These backup links are normally metered for billing purposes. Excessive charges may be incurred when a flapping OSPF area topology activates every one of these backup links for link state update each time the area's topology changes.

Consider the following typical OSPF example:

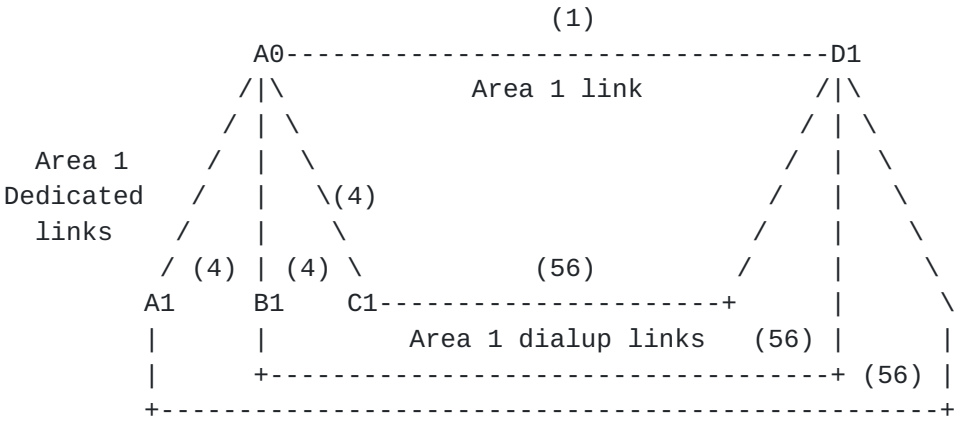

where A0 is a border router attached to Area 1. A1, B1, C1, and D1 are Area 1 internal routers. D1 is also an Area 1 access router which is attached to A0 via a dedicated circuit and which is connected to A1, B1, and C1 via dialup links. A0 connects to A1, B1, and C1 via dedicated circuits. OSPF costs are shown in ()s. If the dedicated link between A0 and C1 is flapping up and down frequently, the current OSPF demand circuit specification, [DC], requires link state update across the adjacencies from A1 and B1 to D1 for each flap, even though D1 learns about the link state change across its adjacency to A0. When there are 50 routers like A1 and B1 with dialup links to D1, activating their dial-up links just for link state update can be very expensive. Indeed it is enough of a reason not to use the OSPF demand circuit specification.

What is needed is a tool which suppresses the link state update of the flapping circuit between A0 and C1 over the links from A1 and B1 to D1 as long as A1 and B1 see a flooding path to D1 other than their dialup links to D1. However C1 still needs to raise its dialup link to D1 when its link to A0 fails, as its only path to the rest of Area

1 would be through D1.

Consider another example:

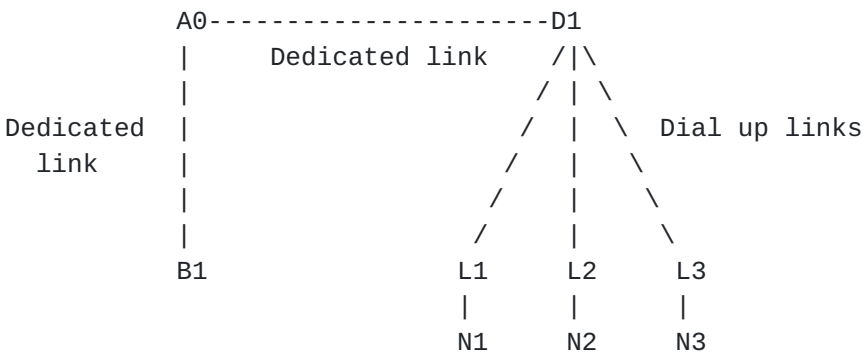

where again A0 is a border router; B1 is an Area 1 internal router; and D1 is an Area 1 internal access router. L1, L2, and L3 are Area 1 leaf internal routers linking stub networks N1, N2, and N3 to Area 1. The OSPF demand circuit specification provides a vehicle whereby the dialup links from D1 to the leaf sites L1/N1, L2/N2, and L3/N3 are used only as required. However should the link between A0 and B1 flap up and down frequently, D1's dialup links to L1, L2, and L3 would activate with each change for no purpose other than needless link state update. In this example a default path though D1 is normally sufficient to meet the basic routing requirements of L1, L2, and L3. In an OSPF area where there are dozens of these leaf sites, such anomalous behavior could be very expensive, making the use of OSPF demand circuit impractical.

What is needed here is a means of filtering a portion of the link state database maintenance between D1 and L1, L2, and L3. L1, L2, and L3 need to advertise their router LSAs to D1, but only need receive LSA updates from D1 when their default path through D1 is impacted. This would provide L1, L2, and L3 with sufficient link state information to compute default through D1.

2.2 Proposed Solution

The OSPF floodgates option provides a vehicle whereby neighboring routers can control the amount of link state information which is exchanged across their adjacencies. In all cases opaque Type 9 LSAs are always synchronized across a floodgated adjacency. Thus all flooding suppression between neighbors is partial. The effect of flooding an opaque Type 9 LSA can be minimized by setting its DoNotAge bit, providing DoNotAging is enable on all routers connected to its link-local network. What is presented here are two floodgate tools. The first one opens its LSA floodgates fully during times of

topological need and then thins the LSA flow to a trickle during times of topological plenty. The second floodgate tool thins LSA flow in only one floodgated direction so that dynamic update in this direction happens only when absolutely necessary or during periods of normal application based traffic flow. Both floodgates are fully open during adjacency formation, and close only after the adjacency reaches Full state.

Partial floodgating is triggered between neighbors via Type 9 opaque LSAs exchanged across their adjacent link. Each opaque LSA requires a unique opaque type. Currently the floodgate option uses experimental opaque type 225 until one is assigned. Opaque type 225 is referred to as the floodgate opaque type in this document (See Appendix A).

Floodgating is enabled via a new interface configuration parameter called FoodgateType (See Appendix A). This parameter defines the type of floodgating which is performed across the interface. When set to "1-Way", 1-Way floodgating is configured. When set to "2-Way", 2-Way floodgating is configured. Both 1-Way and 2-Way floodgates work best with the demand circuit processing of Hello packets, as defined in [DC] Section 3.2, and the DoNotAge processing of LSAs, as defined in [DC] Section 3.3. The requirements of Section 2.1's examples are only met when floodgates, demand circuits, and DoNotAge processing are all used. However, without demand circuit and DoNotAge processing, OSPF floodgates still perform flooding optimization.

The base topology of an area is defined as all its network links plus those router links which have no configured floodgate (1-Way or 2-Way). In order for the area's routers to compute link state paths through it's base topology, any router which has links with floodgating configured must originate a floodgate Type 10 opaque LSA into the area listing those links (See Appendix B). All routers along the area's paths between routers with configured floodgated links should be opaque capable. This guarantees router LSAs and floodgate Type 10 opaque LSAs will be synchronized. As a precaution implementations may choose to disable floodgating in an area when any of the area's routers are not opaque capable.

2.2.1 2-Way Floodgates

2-Way floodgates partially suppress flooding in both directions. Both routers must agree on the floodgate type (either 1-Way or 2-Way) and meet each other's floodgate activation criteria (See Sections 4.1 and 5.1) before floodgating may occur. 2-Way floodgating is enabled via transmitted and received floodgate Type 9 LSAs (See Appendix A). When performing SPF calculations, routers with configured 2-way floodgates must confirm the reachability across the area's base topology of their neighboring routers on the other side of their 2-Way

floodgates.

If an adjacent neighbor on the other side of a 2-Way floodgate becomes unreachable across the area's base topology, then the floodgate is opened for normal link state update. When the floodgate initially opens, the neighbor state machine drops to ExStart state forcing the link state database to resynchronize (See Section 3.5).

2.2.2 1-Way Floodgates

1-Way floodgates allow full updating in one direction and suppress updating in the other floodgated direction. Both routers must agree on the floodgate type before 1-Way floodgating may occur. 1-Way floodgating is activated via a floodgate Type 9 LSA (See Appendix A). This is necessary to provide a vehicle by which each neighbor may open the floodgate.

The implementation of 1-Way Floodgates is facilitated by partitioning an area into sub-areas. By default a router belongs to sub-area 0, but may be configured to belong to a distinct non-zero sub-area. A 1-Way floodgate is always configured between a sub-area 0 router and a non-zero sub-area router. The direction of the partial flow is always from sub-area 0 to the non-zero sub-area. A transit network belongs to each sub-area of any router which connects to it. Area border routers will always belong to sub-area 0. When a 1-Way floodgate is configured on a non-zero sub-area router, in addition to listing its 1-Way and 2-Way floodgated links in its floodgate Type 10 LSA, the router must also define its non-zero sub-area ID (See Appendix B).

A router with a configured floodgate (either 1-Way or 2-Way) is called a gatekeeper. Each floodgate has two gatekeepers, one on each side of the floodgate. A gatekeeper always opens its 1-Way floodgates when it receives a default advertisement from a reachable non-zero sub-area router. Furthermore it keeps its 1-Way floodgates open as long as there exists a default advertisement from a reachable non-zero sub-area router. All of an area's default advertisements must originate either from a sub-area 0 router or from outside the area before any of the area's 1-Way floodgates will close. If there is no installed default through sub-area 0 then a gatekeeper's 1-Way floodgates are opened and remain open until default is restored through sub-area 0.

Since flooding is suppressed across 1-Way floodgates, sub-area routers should have a SPF tree which ensures a stable flooding and forwarding path within each sub-area. During the Shortest Path First (SPF) calculation a gatekeeper should verify that the SPF Tree's branches which leave the local router's sub-area never return as they

grow toward the extremities of the area. This check need only be performed by the gatekeepers. If a gatekeeper fails to meet this criteria it opens all of its floodgates.

2.3 Application

Consider now the first example mentioned in Section 2.1 and assume 2-Way floodgates are configured over the dialup links from D1 to A1, B1, and C1. Furthermore assume DoNotAge processing is in effect within Area 1 and that the dialup links from D1 to A1, B1, and C1 have Demand Circuit enabled. Since a full flooding path exists from A1 and B1 to D1 through A0 the floodgates between D1 and its neighbors, A1 and B1, remain closed and the flapping of the dedicated link between A0 and C1 only effects the floodgate between D1 and C1. The dialup backup links from D1 to A1 and B1 are unaffected and no link state update occurs over these links during this outage.

In the second example we again assume DoNotAge processing is in effect within Area 1 and that the dialup links from D1 to L1, L2, and L3 have Demand Circuit enabled. The L1/N1, L2/N2, and L3/N3 topologies are configured in sub-areas 1, 2 and 3 respectively, while the remaining Area 1 routers are left in sub-area 0. 1-Way floodgates are configured between D1 and L1, L2 and L3. Assume the default path for Area 1 is through A0. When the link between B1 and A0 flaps up and down, the 1-Way floodgates between D1 and L1, L2 and L3 remain closed since the default path of their respective sub-areas has not changed. No link state update occurs over these links during this outage.

3.0 Floodgate Implementation Details

3.1 Interface Data Structure

Floodgating is enabled via a new interface configuration parameter called FoodgateType (See Appendix A). This parameter defines the type of floodgating which is performed across the interface. When set to "None", there is no floodgate configured for the interface. When set to "1-Way", 1-Way floodgating is configured. When set to "2-Way", 2-Way floodgating is configured. The default setting of FloodgateType for any interface is "None". Floodgates are activated on a per neighbor basis. Only neighbors at Full state may have active floodgates.

3.2 Neighbor Data Structure

Floodgate state transition is added to the neighbor state machine. In support of floodgate state transition, three new elements are added to the neighbor data structure: the Floodgate Delay Timer, the

Floodgate State, and the UpdateLinkState flag. The Floodgate State parameter shows the current state of the floodgate. When the neighbor data structure is first created, its Floodgate State parameter is initialized to either "Disabled" or "InActive" depending on whether the interface FloodgateType parameter is respectively "None" or either "1-Way" or "2-Way". As a floodgate transitions from InActive to Active state it assumes intermediate values of Init, 1-Way, and Waiting (See Section 3.3 below).

When a floodgate transitions out of Active state, the link state data base must be resynchronized with that of its adjacent neighbor. This is accomplished by dropping the neighbor to ExchangeStart state. The UpdateLinkState parameter is set to "disabled" so that a new router-LSA and possible network-LSA is not reoriginated. All other events which drop a neighbor out of full state set the UpdateLinkState parameter to "enabled".

A floodgate's gatekeeper resets and starts the single shot Floodgate Delay Timer of one its floodgated neighbors's when all of the following conditions are all met:

- the neighbor meets the gatekeeper's floodgate activation criteria (See Sections 4.1 and 5.1).
- the neighbor's signals that the gatekeeper meets its floodgate activation criteria by listing it as opaque data in its Type 9 LSA (See Sections 4.2 and 5.2).
- the neighbor has reached Full state.

See Sections 3.4 and 3.5 below for details. Normal flooding continues until the timer fires, at which time the floodgate closes and starts suppressing flooding across the link. This ensures a stable flooding path within the area. The current value of the Floodgate Delay Timer is 300 seconds.

When a lower layer protocol activity timer is in effect for the floodgated link, non-OSPF traffic over the link should reset the Floodgate Delay Timer using the ResetFloodgate event (See Section 3.4 and Appendix C). When a floodgate is closed lower layer activity may also open a floodgate by triggering the ResetFloodgate event (See Section 3.4). While the Floodgate Delay Timer is counting down, the floodgate is still open.

3.3 Floodgate States

The various states that OSPF floodgates may attain is documented in this section. These states are listed in order of progressing

functionality. For example, the inoperative state, InActive, is listed before the intermediate states Init, 1-Way, and Waiting. The last state, Active, is the fully operational floodgate state. The specification makes use of this ordering by sometimes making references such as "in state greater than X". During all floodgate states less than Active full link state update with the neighbor is in effect and the floodgate is said to be fully open.

Disabled
This state indicates no floodgate is configured.

InActive
This state indicates the floodgated neighbor does not meet the local router's floodgate activation criteria (See Sections 4.1 and 5.1).

Init
This state indicates the floodgated neighbor meets the local router's floodgate activation criteria but a floodgate Type 9 LSA with matching floodgate type has not been received from the neighbor.

1-Way
This state indicates the neighbor's floodgate Type 9 LSA was received with matching floodgate type but the local router is not listed in the LSA's opaque data, implicating the local router does not yet meet the neighbor's floodgate activation criteria.

Waiting
This state indicates the neighbor's floodgate Type 9 LSA has been received with matching floodgate type and it lists the local router in the LSA's opaque data, implying the local router does meet the neighbor's floodgate activation criteria.

Active
In this state, the floodgate between the neighbor has closed and it is actively suppressing link state update.

3.4 Floodgate Events

The following new events are defined in support of floodgates.

OpenFloodgate
Both 1-Way and 2-Way floodgates have floodgate activation criteria which must be met for these floodgates to remain closed. When a gatekeeper observes that the floodgate activation criteria of its floodgated neighbor is no longer

met, this event is triggered (See Sections 4.1 and 5.1). Following this event the floodgate transitions from closed to open. Floodgates may also be opened by lower layer protocols when dial-up links are active.

ReleaseFloodgate
When a gatekeeper observes the floodgate activation criteria of its floodgated neighbor is now met, whereas before it was not, then this event is triggered, releasing the floodgate for closure process. (See Sections 4.1 and 5.1).

FloodgateDelayTimer
When this timer fires, it signals the fact that the neighbor's floodgated has been in Waiting state for the full duration of the timer during which all necessary criteria for closing the floodgate have been met. Its firing causes the floodgate to close and flooding suppression to begin.

ResetFloodgate
Lower layer protocols have indicated the floodgate should be opened due to non-OSPF traffic flow.

1-Way9Received
A floodgate Type 9 LSA has been received in which the local router is not mentioned. Either the floodgated neighbor has not received a floodgate Type 9 LSA from the local router, or the local router no longer meets the neighbor's floodgate activation criteria.

2-Way9Received
A floodgate Type 9 LSA has been received in which the local router is listed. Thus the neighbor has received the local router's floodgate Type 9 LSA and the local router meets the neighbor's floodgate activation criteria.

3.5 Floodgate State Transitions

Several neighbor state machine events effect floodgate state. If the floodgate's state is greater than InActive, any change in the floodgated neighbor's Full state results in the floodgate's state dropping to Init state. If the floodgate's state is Waiting when the neighbor event LoadingDone is triggered, the Floodgate Delay Timer is reset and started. Otherwise no action is taken. All neighbor state machine events which drop a neighbor out of full state must set the UpdateLinkState parameter to "enabled".

When an OpenFloodgate event occurs and the floodgate's state is greater than InActive then the new state is set to InActive and a new

floodgate Type 9 LSA is originated omitting the neighbor from its opaque data. If the floodgate's old state is Active it sets the UpdateLinkState to disabled and triggers the neighbor's state machine to perform the same action as a SeqNumMismatch event, thus dropping the neighbor's state to ExStart. No new router-LSA or network-LSA is generated during the data base resynchronization unless problems occur.

When a ResetFloodgate event occurs and the floodgate's state is Active then set the UpdateLinkState to "disabled" and trigger the neighbor's state machine to perform the same action as a SeqNumMismatch event, thus dropping the neighbor's state to ExStart.

Following any ResetFloodgate event where the floodgate state is greater than InActive then

(1) If there does not exist an acknowledged floodgate Type 9 LSA for the neighbor set the floodgate state to Init,

(2) Else if an acknowledged floodgate Type 9 LSA exists but it does not list the local router then set the floodgate state to 1-Way,

(3) Else the acknowledged floodgate Type 9 LSA does list the local router, so set the floodgate state to Waiting. If the floodgated neighbor's state machine is in Full state, start the Floodgate Delay Timer.

When a ReleaseFloodgate event occurs and the Floodgate state is InActive then

(1) If there does not exist an acknowledged floodgate Type 9 LSA for the neighbor set the floodgate state to Init,

(2) Else if an acknowledged floodgate Type 9 LSA exists but it does not list the local router, then set the floodgate state to 1-Way,

(3) Else the acknowledged floodgate Type 9 LSA does list the local router, so set the floodgate state to Waiting. If the floodgated neighbor's state machine is in Full state, start the Floodgate Delay Timer.

Once a floodgate is released it may transition between the Init, 1-Way and Waiting states toward activation and closure. The events which effect these state transitions are described below.

When a 1-Way9Received event occurs and the floodgate's state is

Active, set the UpdateLinkState to disabled and trigger the neighbor's state machine to perform the same action as a SeqNumMismatch event, thus dropping the neighbor's state to ExStart.

When a 1-Way9Received event occurs and the floodgate's state is Waiting, stop a running Floodgate Delay Timer.

Following any 1-Way9Received event where the floodgate's state is greater than InActive, set the floodgate's state to 1-Way.

When a 2-Way9Received event occurs, if the floodgate's state is Init or 1-Way then set the floodgate's state to Waiting and, if the floodgated neighbor's state machine is in Full state, reset and start the Floodgate Delay Timer.

When the Floodgate Delay Timer fires it triggers the FloodgateDelayTimer event. If the Floodgate state is Waiting then the Floodgate is closed and flooding suppression begins.

3.6 Receiving Floodgate Type 9 LSAs

When a new functionally different floodgate Type 9 LSA is received or acknowledged across a floodgated interface during link state update or database exchange, its floodgate type is compared with that of the receiving interface. If they do not match then any existing neighbor state machine for the neighbor is executed with the OpenFloodgate event, and processing of this received floodgate Type 9 LSA terminates.

Otherwise the opaque data of the floodgate Type 9 LSA is checked to see if the local router is listed. If it is not, then the neighbor's state machine is executed with the event 1-Way9Received. If it is listed, then the neighbor state machine is executed with the event 2-Way9Received.

3.7 Originating Floodgate Type 10 LSAs

In order to compute link state paths through an area's base topology, all of the area's routers with configured floodgates (1-Way and 2-Way) must originate a floodgate Type 10 LSA which lists those links (See Appendix B). In addition a router advertises its sub-area ID in its floodgate Type 10 LSA. By default a router's sub-area ID is set to 0. only routers with a non-zero sub-area ID in their area data structures must originate a Type 10 LSA specifying its value in the Sub-area ID field and must do so even when they do not have any configured floodgates. All routers in an area with configured floodgates should be opaque capable, otherwise the flooding of floodgate Type 10 LSAs may be incomplete and some floodgates may

close inappropriately. When a new router LSA is originated by a router for a area in which the router has configured floodgates, the router should also originate a new floodgate Type 10 LSA into the area. This keeps router LSAs and floodgate Type 10 LSAs closely synchronized. As a precaution implementations may choose to disable floodgating in an area when any of the area's routers are not opaque capable.

3.8 Link State Acknowledgment and Retransmission

It is entirely possible that two gatekeepers on each side of a floodgate may close their side of the floodgate asynchronously. Any received link state update packets must still be properly received and acknowledged. Any link state requests or non-empty link state retransmission list must still be processed. Thus residual link state update may continue for a while following the closure of a floodgate. Floodgate closure simply means any newly received or originated LSAs, except for Type 9 LSAs, are not flooded out an interface with a closed floodgate unless requested. Note that 1-Way floodgates are only closed in one direction. Full flooding occurs in the other direction.

This fact has a direct impact on interfaces of broadcast type. Floodgates do work over these types of interfaces, but functionality is impacted by the multicasting of link state update. For the most part floodgates over broadcast networks are only effective when there are closed foodgates between the DR and Backup DR, plus all of their adjacent neighbors. These floodgates will still eventually close as the closure of floodgates is not effected by current link state update once two neighbors have reached Full state.

4.0 2-Way Floodgate Implementation Details

While floodgate state transition is basically the same for both 2-Way and 1-Way floodgates, floodgate activation criteria are notably different. When a 2-Way Floodgate is active, it is assumed that an interface's floodgated neighbors are reachable via the area's base topology. Every router with configured floodgates (1-Way and 2-Way) must originate a Type 10 LSA specifying their floodgated links. These links are removed from the area's link state data base to form the base topology. The base topology provides a flooding path between the two neighbors which avoids any floodgated links. When two floodgated neighbors are reachable across the base topology, the flooding of all LSAs, except for Type 9 LSAs, may be suppressed across an active 2-Way floodgate. The effects of Type 9 LSA flooding and Hello packet exchange can be further mitigated with demand circuit and DoNotAge processing.

4.1 2-Way Floodgate State Transitions

When the link state data base of an area changes, each of its routers with configured 2-Way floodgates must recheck the reachability of it's floodgated neighbors across the area's base topology. If a 2-Way floodgated neighbor becomes newly unreachable across the area's base topology, then the OpenFloodgate event should be triggered. This causes the neighbor's 2-Way floodgate to fully open and its floodgate state to transition to InActive. A new floodgate Type 9 LSA is originated at the floodgated interface with the neighbor omitted from its opaque data. If a 2-Way floodgated neighbor becomes newly reachable across the area's base topology, then the ReleaseFloodgate event should be triggered. This causes the 2-Way floodgate to transition out of InActive state and to begin the process of swinging shut.

Lower layer protocols may open a 2-Way floodgate with the ResetFloodgate event. Non-OSPF traffic in either direction may trigger this event and will keep the floodgate in Waiting state until a lower layer protocol inactivity timer fires. Lower layer protocol inactivity timers need to be modified to ignore OSPF protocol 89 in both directions across 2-Way floodgated links. Note that floodgates do not stop the normal exchange of Hello packets. Only OSPF DDR suppresses Hello packets.

4.2 Originating 2-Way Floodgate Type 9 LSAs

Floodgate Type 9 LSAs are originated across every 2-Way floodgated interface with their Floodgate Type set to 2-Way. In addition those neighbors which are reachable across the area's base topology and from whom acknowledged floodgate Type 9 LSAs have been received (or acknowledged during database exchange) with matching floodgate type, are listed in the LSAs opaque data (See Appendix A).

5.0 1-Way Floodgate Implementation Details

For a 1-Way Floodgate to be active in an area, the area must be partitioned into sub-areas. A router may belong to only one sub-area. Each router of a non-zero sub-area must declare its Sub-area ID even when it has no configured floodgates. This is done via a floodgate Type 10 LSA. By default a router belongs to sub-area 0. Thus routers in sub-area 0 only need originate a floodgate Type 10 LSA if they have configured floodgates (See Section 3.7). A new area configuration parameter called Sub-area ID is used to set the sub-area ID of an area's internal routers. Border routers are always in sub-area 0.

Like a stub area, a non-zero sub-area uses OSPF's classless default

to forward traffic destined outside its borders. Default forwarding is not always optimum when there are multiple floodgates bordering a non-zero sub-area, as the part of the areas's LSDB which is external to the non-zero sub-area may age out or become obsolete. The non-zero sub-area side of a 1-Way floodgate always performs dynamic link state update with the sub-area 0 side of the floodgate. Thus sub-area 0 has a full link state database of the area's entire topology. When a 1-Way floodgate is closed, the sub-area 0 side of a floodgate suppresses all LSA updates, except for Type 9 LSAs, to the non-zero sub-area side of the floodgate.

5.1 1-Way Floodgate State Transitions

Maintaining a default path and border router reachability are very important floodgate activation criteria to the health of a 1-Way floodgate. A sub-area 0 router will keep its 1-Way floodgates closed only when all of its default next-hops are over non-floodgated sub-area 0 links. If a sub-area 0 router's new default installation has a next-hop in a non-zero sub-area, whereas its old default next hop was in sub-area 0, the OpenFloodgate event is triggered for all of the router's 1-Way floodgated neighbors. This causes these neighbors' floodgates to fully open and their floodgate states to transition to InActive. If a sub-area 0 router looses its default through sub-area 0 it triggers the OpenFloodgate event for all of its 1-Way floodgated neighbors.

A non-zero sub-area router must install a default next-hop from an advertisement originating outside the area's non-zero sub-areas. Unfortunately a non-zero sub-area router with a configured 1-Way floodgate has no easy way of knowing what default path(s) the other routers in its sub-area install. To circumvent this problem, a non-zero sub-area router should trigger an OpenFloodgate event for any 1-Way floodgated neighbor when it receives a default advertisement from any reachable non-zero sub-area router, even ones from a different non-zero sub-area. This rule ensures that a default path will eventually trail out of the area through sub-area 0. It even allows for the unwise case when two non-zero sub-areas neighbor each other. In configuration terms, a non-zero sub-area router should never originate a default advertisement, as that may disable 1-Way floodgating throughout the entire area. Since a non-zero sub-area router performs full flooding over all of its interfaces, even those with configured 1-Way floodgates, it has no other floodgate activation criteria. Note that non-zero sub-areas may even partition with no effect, as they simply become separate non-zero sub-areas.

Even though the link state database of a non-zero sub-area is synchronized within the sub-area, because of its 1-Way floodgates it is highly unlikely a non-zero sub-area will have sufficient link

state information to transit the rest of the area's traffic through its topology unless its floodgates are open. Hence any sub-area 0 router with configured 1-Way floodgates should ensure during the Shortest Path First calculation that no sub-area 0 router is directly below a non-zero sub-area router on its Shortest Path First tree. When this is the case sub-area 0 is said to be topologically concave. When a sub-area 0 router's Shortest Path First calculation indicates sub-area 0 is no longer topologically concave, it should trigger an OpenFloodgate event for all of its 1-Way floodgates.

If, as a result of a sub-area 0 router's Shortest Path First calculation, the router determines that all of its installed defaults have next hop in sub-area 0 and that sub-area 0 is topologically concave, then the ReleaseFoodgate event is triggered for all of its 1-Way floodgated neighbors. If on a non-zero sub-area router the last advertisement for default originating from a non-zero sub-area router ages out or is flushed, or the router which originated this default advertisement becomes unreachable and there still exists an installed SPF default, then the ReleaseFloodgate event should be triggered for all of its 1-Way floodgated neighbors. These conditions cause a 1-Way floodgate to transition out of InActive state and to begin the process of swinging shut.

5.2 Originating 1-Way Floodgate Type 9 LSAs

Floodgate Type 9 LSAs are originated across every 1-Way floodgated interface with the Floodgate Type set to 1-Way. Furthermore a 1-Way floodgated neighbor is listed in the LSA's opaque data when

- a floodgate Type 9 LSA has been received from the neighbor with matching 1-Way Floodgate Type,
- a floodgate Type 10 LSA has been received from the neighbor with the appropriate Sub-area ID, and
- the neighbor meets the local router's floodgate activation criteria.

Two non-zero sub-area routers may never form a 1-Way floodgate, even though they can still be neighbors. In practice two distinct non-zero sub-areas which have mutually adjacent neighbors simply form a single non-zero sub-area even though they have different non-zero Sub-area IDs. While the Sub-area ID is a good mental tool for keeping tract of sub-areas, the boundaries of a non-zero sub-area are determined by where it meets sub-area 0, and not by the sub-area ID sameness or difference of its router members.

6.0 Floodgate Extensions

The floodgate machinery defined in Section 3.0 through Section 3.7 provides the necessary tools for adding additional floodgate extensions. The Floodgate Type field in the floodgate Type 9 LSA reserves values 0 through 127 for future floodgate extensions. Values 128 through 255 may be used for experimentation or implementation specific value added features. Currently only three values are assigned, 0 for no floodgating, 1 for 1-Way floodgating and, 2 for 2-Way floodgating.

A notable floodgate extension omission is a floodgate which suppresses the flooding of summary-LSAs and AS-external-LSAs from a sub-area 0 router to a non-zero sub-area router. It would function much like a 1-Way floodgate but its floodgate activation criteria is much simpler.

7.0 Security Considerations

Security issues are not discussed in this memo.

8.0 Acknowledgments

This document was produced by the OSPF Working Group, chaired by John Moy. In addition, comments from the following individuals are also acknowledged:

Acee Lindem IBM

9.0 References

[DC] Moy, J., "Extending OSPF to Support Demand Circuits", RFC 1793, Cascade Communication Corp., April 1995.

[MIB] McCloghrie, K., and M. Rose, "Management Information Base for network management of TCP/IP-based internets: MIB-II", STD 17, RFC 1213, Hughes LAN Systems, Performance Systems International, March 1991.

[OPAQ] Coltun, Rob, "The OSPF Opaque LSA Option", FORE Systems, August 1998.

[OSPF] Moy, J., "OSPF Version 2", RFC 2328, Ascend Communications, Inc., April 1998.

[1583] Moy, J., "OSPF Version 2", RFC 1583, Proteon, Inc., March 1994.

10.0 Author's Address

Pat Murphy
US Geological Survey
345 Middlefield Road, Mail Stop 870
Menlo Park, California 94560

Phone: (650)329-4044
EMail: pmurphy@noc.doi.net

Appendix A. Floodgate Type 9 LSAs

Floodgate Type 9 LSAs are used directly by OSPF to synchronize floodgate state between neighboring routers over links. These LSAs are flooded across the floodgated link. The flooding scope of floodgate type 9 LSAs is link-local, which means floodgate Type 9 LSAs are never forwarded beyond the local (sub)network which contains their floodgates.

Section 3.5 of this document describes the purpose of this floodgate opaque LSA in more detail.

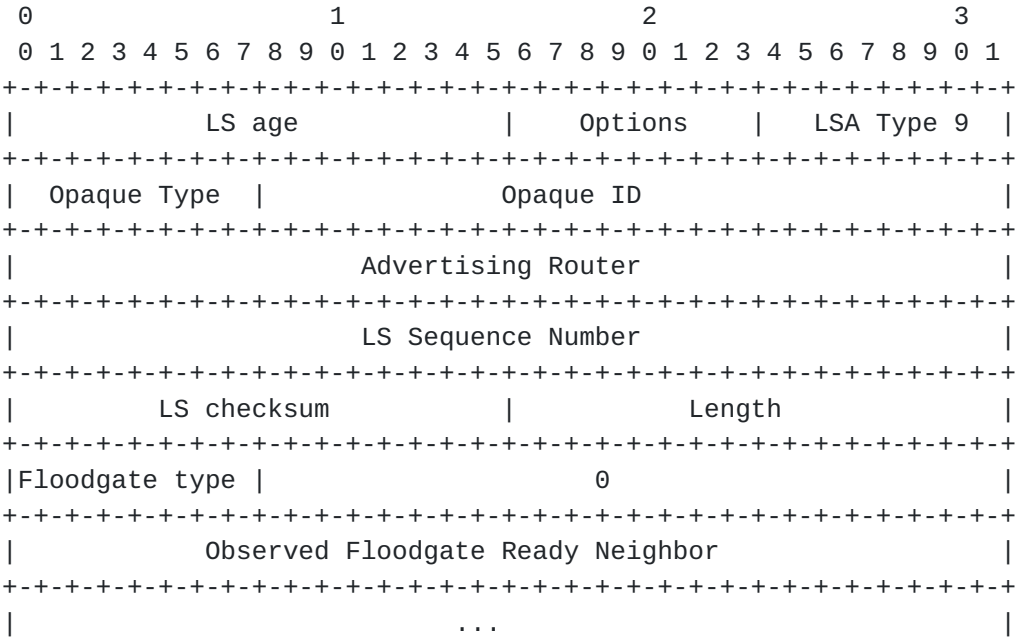

Syntax of the floodgate Type 9 LSA's Link State ID

The link-state ID of the floodgate Type 9 LSA is divided into an Opaque Type field (the first 8 bits), which has the floodgate opaque type (225) value, and the Opaque ID (the remaining 24 bits). The floodgate Type 9 LSA's Opaque ID is set to the last 24 bits of the interface's IP address, if it has one. In the case of unnumbered point-to-point connections, the floodgate Type 9 LSA's Opaque ID is set to the last 24 bits of the interface's MIB-II [MIB] ifIndex value. The most significant bits may be changed to ensure uniqueness.

Floodgate Type

This parameter is set to 0 when no floodgates are configured for this link. It is set to 1 when 1-Way floodgating is configured for the link and it is set to 2 when 2-way floodgating is configured for the link.

Observed Floodgate Ready Neighbor

All neighbors from whom Floodgate Type 9 opaque LSAs have been received and who meet the local router's floodgate activation criteria are listed in at the end of Floodgate Type 9 LSA (See Section 4.2. and 5.2).

Appendix B. Floodgate Type 10 LSAs

Floodgate Type 10 LSAs are used directly by OSPF to advertise inside an area a router's list of floodgated links. The flooding scope of floodgate Type 10 LSAs is areawide, which means floodgate Type 10 LSAs are never forwarded beyond the area border. A router which has floodgated links must originate a floodgate Type 10 LSA listing those links. If a router belongs to a non-zero sub-area, it must also originate a floodgate Type 10 LSA which specifies its sub-area in the Sub-area ID field. If a router has no floodgated links and belongs to sub-area 0, it need not originate a floodgate Type 10 LSA.

Section 3.7 of this document describes the purpose of the floodgate Type 10 LSA in more detail.

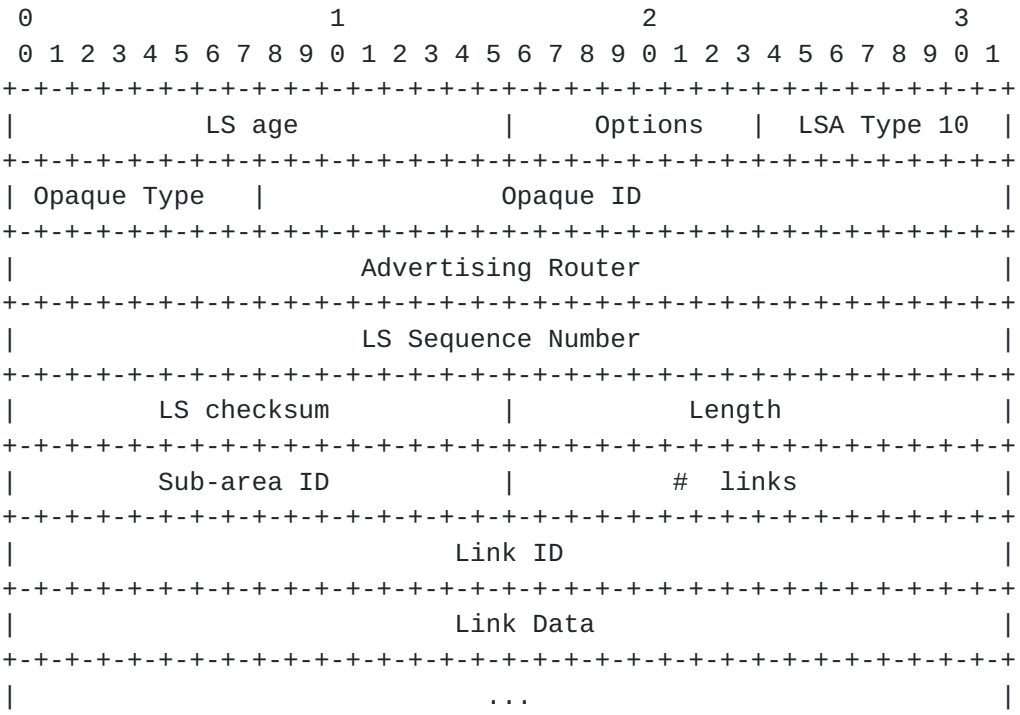

Syntax Of The Opaque LSA's Link State ID

The link-state ID of the floodgate Type 10 LSA is divided into an Opaque Type field (the first 8 bits), which has value floodgate (currently 224), and the Opaque ID (the remaining 24 bits). The floodgate Type 10 Opaque ID is set to the last 24 bits of the OSPF interface's IP address, if it has one. In the case of unnumbered point-to-point connections, the floodgate Type 10 opaque ID is set to the last 24 bits of the interface's MIB-II [MIB] ifIndex value. The most significant bits may be changed to ensure uniqueness.

links
The number of floodgated links listed in this LSA. This must be the local router's total collection of floodgated links with flooding suppression configured for the area.

Sub-area ID
An OSPF area with configured 1-Way floodgates must be broken into OSPF sub-areas. 1-Way floodgates are configured between routers in sub-area 0 and routers in non-zero sub-areas. By default the Sub-area ID field is set to 0. A router specifies its non-zero sub-area in the the Sub-area ID field.

The following fields are used to describe each of the router's links. They must match their counterparts in the router's Type 1 (router) LSA for this area.

Link ID
The Router ID of the neighboring router that this router connects to via a link. This provides the key for looking up the neighboring LSA in the link state database during the routing table calculation. See [OSPF] Section 12.2 for more details.

Link Data
For unnumbered point-to-point connections, it specifies the interface's MIB-II [MIB] ifIndex value. For the other link types it specifies the router interface's IP address. This latter piece of information is needed during the routing table build process, when calculating the IP address of the next hop. See [OSPF] Section 16.1.1 for more details.

Appendix C. Configuration Parameters

Floodgating is enabled via a new interface configuration parameter called FoodgateType. When set to "None", there is no floodgate configured for the interface. When set to "1-Way", 1-Way floodgating is configured. When set to "2-Way", 2-Way floodgating is configured. The default setting of FloodgateType for any interface is "None". A maximum of 255 different floodgate types may be defined. Currently only these three are defined. The first 128 floodgate types are reserved for future additions to this specification. The remaining 128 floodgate types are available for experimentation and implementation specific value added features (See Section 3.1).

A new area configuration parameter called Sub-area ID is used to set the sub-area of an area's internal routers. It should not be possible to set this parameter on a border router, as the border routers of an area are always in sub-area 0.

A new interface configuration parameter called the ActivityTrigger has been added. Its function is implemented by lower level protocols. When set to enabled, non-OSPF IP traffic between two floodgated neighbors may open a floodgate. When set to disabled, non-OSPF IP traffic has no effect on the state of a floodgate.