

Network Working Group  
Internet Draft  
Expiration Date: July 1997  
File name: [draft-ietf-ospf-nssa-02.txt](#)

R. Coltun  
FORE Systems  
V. Fuller  
BBN Planet  
P. Murphy  
US Geological Survey  
April 1997

## **The OSPF NSSA Option**

### Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress".

To learn the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

## Table Of Contents

<a href="#">1.0</a>	Abstract .....	<a href="#">1</a>
<a href="#">2.0</a>	Overview .....	<a href="#">2</a>
<a href="#">2.1</a>	Motivation - transit networks .....	<a href="#">2</a>
<a href="#">2.2</a>	Motivation - corporate networks .....	<a href="#">3</a>
<a href="#">2.3</a>	Proposed Solution .....	<a href="#">4</a>
<a href="#">3.0</a>	Implementation Details .....	<a href="#">6</a>
<a href="#">3.1</a>	The N-bit .....	<a href="#">6</a>
<a href="#">3.2</a>	Type-7 Address Ranges .....	<a href="#">7</a>
<a href="#">3.3</a>	Type-7 LSAs .....	<a href="#">7</a>
<a href="#">3.4</a>	Originating Type-7 LSAs .....	<a href="#">8</a>
<a href="#">3.5</a>	Calculating Type-7 AS External Routes .....	<a href="#">9</a>
<a href="#">3.6</a>	Incremental Updates .....	<a href="#">12</a>
<a href="#">4.0</a>	Originating Type-5 LSAs .....	<a href="#">13</a>
<a href="#">4.1</a>	Translating Type-7 LSAs .....	<a href="#">13</a>
<a href="#">4.2</a>	Flushing Translated Type-7 LSAs .....	<a href="#">16</a>
<a href="#">5.0</a>	Acknowledgments .....	<a href="#">17</a>
<a href="#">6.0</a>	References .....	<a href="#">17</a>
<a href="#">7.0</a>	Security Considerations .....	<a href="#">17</a>
<a href="#">8.0</a>	Authors' Addresses .....	<a href="#">18</a>
<a href="#">Appendix A</a>	Type-7 LSA Packet Format .....	<a href="#">19</a>
<a href="#">Appendix B</a>	The Options Field .....	<a href="#">20</a>
<a href="#">Appendix C</a>	Router-LSAs .....	<a href="#">21</a>
<a href="#">Appendix D</a>	Configuration Parameters .....	<a href="#">23</a>
<a href="#">Appendix E</a>	Differences from <a href="#">RFC 1587</a> .....	<a href="#">24</a>

**[1.0](#) Abstract**

This memo documents of an optional type of OSPF area which is somewhat humorously referred to as a "not-so-stubby" area (or NSSA). NSSAs are similar to the existing OSPF stub area configuration option but have the additional capability of importing AS external routes in a limited fashion.

The OSPF NSSA Option was originally defined in [RFC 1587](#). The functional differences between this memo and [RFC 1587](#) are explained in [Appendix D](#). All differences, while expanding capability, are backward-compatible in nature. Implementations of this memo and of [RFC 1587](#) will interoperate.

Please send comments to [ospf@gated.cornell.edu](mailto:ospf@gated.cornell.edu).

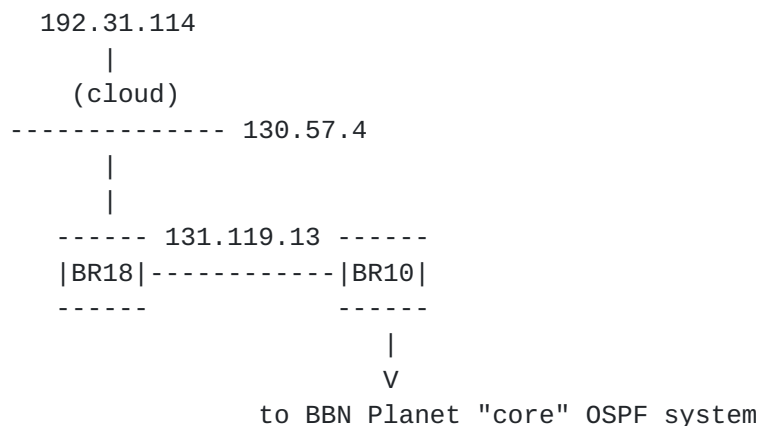


## [2.0](#) Overview

### [2.1](#) Motivation - transit networks

Wide-area transit networks (such as the NSFNET regionals) often have connections to moderately-complex "leaf" sites. A leaf site may have multiple IP network numbers assigned to it.

Typically, one of the leaf site's networks is directly connected to a router provided and administered by the transit network while the others are distributed throughout and administered by the site. From the transit network's perspective, all of the network numbers associated with the site make up a single "stub" entity. For example, BBN Planet has one site composed of a class-B network, 130.57.0.0, and a class-C network, 192.31.114.0. From BBN Planet's perspective, this configuration looks something like this:



where the "cloud" consists of the subnets of 130.57 and network 192.31.114, all of which are learned by RIP on router BR18. Topologically, this cloud looks very much like an OSPF stub area. The advantages of running the cloud as an OSPF stub area are:

1. Type-5 routes (OSPF external link-state advertisements (LSAs)) are not advertised beyond the router labeled "BR10". This is advantageous because the link between BR10 and BR18 may be a low-speed link or the router BR18 may have limited resources.
2. The transit network is abstracted to the "leaf" router BR18 by advertising only a default route across the link between BR10 and BR18.
3. The cloud becomes a single, manageable "leaf" with respect to the transit network.



4. The cloud can become, logically, a part of the transit network's OSPF routing system.
5. Translated type-5 LSAs that are sent into the backbone from the cloud (which is a separate stub area) may be considered "leaf" nodes when performing the Dijkstra calculation.

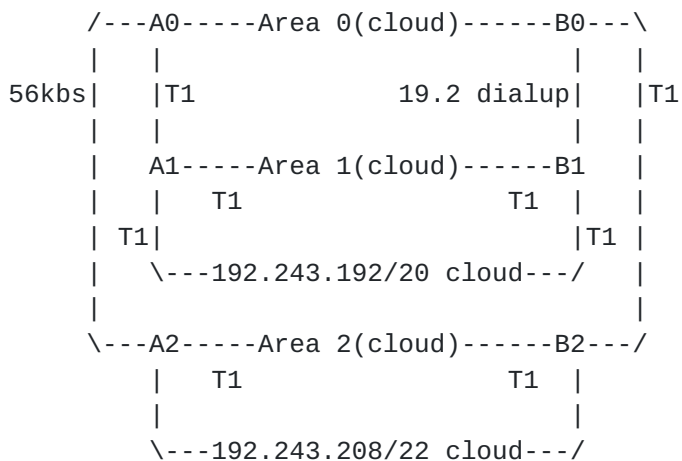
However, the current definition of the OSPF protocol [[1](#)] imposes topological limitations which restrict simple cloud topologies from becoming OSPF stub areas. In particular, it is illegal for a stub area to import routes external to OSPF; it is not possible for routers BR18 and BR10 to both be members of the stub area and to import the routes learned from RIP or other IP routing protocols as type-5 (OSPF external LSAs) into the OSPF system. In order to run OSPF out to BR18, BR18 must be a member of a non-stub area or the OSPF backbone before it can import routes other than its directly-connected network(s). Since it is not acceptable for BR18 to maintain all of BBN Planet's external (type-5) routes, BBN Planet is forced by OSPF's topological limitations to only run OSPF out to BR10 and to run RIP between BR18 and BR10.

## **2.2 Motivation - corporate networks**

In a corporate network which supports a large corporate infrastructure it is not uncommon for OSPF area 0 to have a large area infrastructure injecting large routing tables into area 0. Organizations within the corporate infrastructure may routinely multi-home their non-0 OSPF areas to strategically located backbone area 0 routers, either to provide backbone redundancy or increase backbone connectivity or both. Because of these large routing tables, OSPF aggregation via summarization is routinely used and recommended. Stub areas are also recommended to keep the size of the non-zero routing tables small. Organizations within the corporation are administratively autonomous and compete for corporate backbone resources. They also want isolation from each other in order protect their own network resources within the organization.

Consider a typical backbone connection, as shown on the next page, where routers An, Bn are connected to their respective area n's and routers A0 and B0 are border routers to both area 1 and area 2. Serial lines are displayed, but several ethernets, not displayed, may also connect from the connected areas at each router depicted internal to Areas 1 and 2. Assume the 192.243.192/20 and 192.243.208/22 clouds are subnetted with a protocol foreign to the corporate OSPF process. These processes could be RIP, IGRP, or second and third OSPF processes separate from the corporate OSPF backbone process.





As a matter of policy, the corporate network administrators want 192.243.192/20 and 192.243.208/22 aggregated to the backbone. This policy conflicts with both Area 1 and Area 2's desire to see the aggregate's subnetted infrastructures learned from Area 1's internal routers A1 and B1 and Area 2's internal routers A2 and B2 from which it can make efficient routing decisions. The current standard OSPF stub area has no mechanism to support the redistribution of routes for 192.243.192/20 and 192.243.208/22 subnets within their respective areas. If one assumes that both the Area 1 and Area 2 clouds are extremely large with internally very large routing tables originating from a complex OSPF link state topology and subnetting scheme, neither Area 1 or Area 2 wants to be at the mercy of either the other area's summary links advertisements or external links advertisements. Thus standard OSPF areas are not an option either.

Any solution to this dilemma must honor Area 1's path of choice through A0 with redundancy through B0 while at the same time honoring Area 2's path of choice through B0 with redundancy through A0. Furthermore, such a solution must support the aggregation of the externally learned subnetted routing subdomain.

### 2.3 Proposed Solution

This document describes a new optional type of OSPF area, somewhat humorously referred to as a "not-so-stubby" area (or NSSA), which has the capability of importing external routes in a limited fashion.

The OSPF specification defines two general classes of area configuration. The first allows type-5 LSAs to be flooded throughout the area. In this configuration, type-5 LSAs may be originated by routers internal to the area or flooded into the area by area border routers. These areas, referred to herein as type-5 capable areas (or just plain areas in the OSPF specification), are distinguished by the fact that they can carry transit traffic. The backbone is always a type-5 capable area. The

second type of area configuration, called stub, allows no type-5 LSAs to

be propagated into/throughout the area and instead depends on default routing to external destinations.

NSSAs are defined in much the same manner as existing stub areas. To support NSSAs, a new option bit (the "N" bit) and a new type of LSA (type-7) are defined. The "N" bit ensures that routers belonging to a NSSA agree on its configuration. Similar to the stub area's use of the "E" bit, both NSSA neighbors must agree on the setting of the "N" bit or the OSPF neighbor adjacency will not form.

Type-7 LSAs provide for carrying external route information within a NSSA. Type-7 AS External LSAs have virtually the same syntax as the Type-5 AS External LSAs with the obvious exception of the link-state type (see [section 3.2](#) for more details). There are two major semantic differences between type-5 and type-7 LSAs.

- o Type-7 LSAs may be originated by and advertised throughout a NSSA; as with stub areas, type-5 LSAs are not flooded into NSSAs and do not originate there, except on NSSA border routers.
- o Type-7 LSAs are advertised only within a single NSSA; they are not flooded into the backbone area or any other area by border routers, though the information which they contain can be propagated into the backbone area (see [section 3.6](#)).

In order to allow limited exchange of external information across a NSSA border, NSSA border routers will translate selected type-7 LSAs received from the NSSA into type-5 LSAs. These type-5 LSAs will be flooded to all type-5 capable areas. NSSA border routers may be configured with address ranges so that several type-7 LSAs may be represented by a single type-5 LSA. The NSSA border routers which perform translation are configurable thus creating efficient forwarding to type-5 LSA originating from aggregated type-7s.

In addition, a NSSA border router may originate a default type-7 LSA (IP address of 0.0.0.0) into the NSSA, and must if no NSSA internal type-7 default route exists. Default routes are necessary because NSSAs do not receive full routing information and must have a default route to route to AS-external destinations. Like stub areas, NSSAs may be connected to the backbone at more than one area border router, but may not be used as a transit area. Note that the default route originated by a NSSA border router is never translated into a type-5 LSA, however, a default route originated by a NSSA internal AS boundary router (one that is not also an area border router) may be translated into a type-5 LSA.

Like stub areas, the importing of OSPF summary routes (type-3 LSAs) into NSSAs is a configuration option. However particular care should be taken

to ensure that OSPF internal routes are always chosen over OSPF external (type-7) routes. This may happen when other IGP's, like RIP and ISIS,

leak routing information between an OSPF NSSA and another OSPF area. In these cases, all OSPF summary routes should be imported into the effected NSSAs. The recommended default behavior is to import OSPF summary routes into NSSAs. Note that if a type-7 default originates from an internal NSSA router, all summary routes must automatically be imported into the NSSA. This insures that OSPF internal routes are preferred over an internal type-7 LSA default path which may cause inter-AS traffic to exit the AS.

In our transit example topologies the subnets of 130.57 and network 192.31.114 will still be learned by RIP on router BR18 but now both BR10 and BR18 can be in a NSSA and all of BBN Planet's external routes are hidden from BR18; BR10 becomes a NSSA border router and BR18 becomes an AS boundary router internal to the NSSA. BR18 will import the subnets of 130.57 and network 192.31.114 as type-7 LSAs into the NSSA. BR10 then translates these routes into type-5 LSAs and floods them into BBN Planet's backbone.

In our corporate example, the subnets of 192.243.192/20 and 192.243.208/22 are learned via their respective routing process, redistributed throughout NSSAreas 1 and 2, and then aggregated during the translation process into a single Type-5 LSA which is flooded into Area 0. Area 1 may configure A0 to perform translation with B0 standing by as a backup translator, while Area 2 configures B0 as its translator with A0 its backup.

### **3.0 Implementation Details**

#### **3.1 The N-bit**

The N-bit ensures that all members of a NSSA agree on the area's configuration. Together, the N-bit and E-bit reflect an interface's (and consequently the interface's associated area) external LSA flooding capability. As explained in [section 10.5](#) of the OSPF specification, if type-5 LSAs are not flooded into/throughout the area, the E-bit must be clear in the option field of the received Hello packets. Interfaces associated with a NSSA will not send or receive type-5 LSAs on that interface but may send and receive type-7 LSAs. Therefore, if the N-bit is set in the options field, the E-bit must be cleared.

To support the NSSA option an additional check must be made in the function that handles the receiving of the Hello packet to verify that both the N-bit and the E-bit found in the Hello packet's option field match the value of the options that have been configured in the receiving interface. A mismatch in the options causes processing of the received Hello packet to stop and the packet to be dropped.



### **3.2 Type-7 Address Ranges**

NSSA border routers may be configured with type-7 address ranges.

Each address range is defined as an [address,mask] pair. Many separate type-7 networks may then be represented by a single address range, just as a subnetted network is composed of many separate subnets. NSSA border routers may then summarize type-7 routes by advertising a single type-5 route for each type-7 address range. The type-5 route, resulting from a type-7 address range match will be distributed to all type-5 capable areas. [Section 4.1](#) gives the details of generating type-5 routes from type-7 address ranges.

A type-7 address range includes the following configurable items.

- o An [address,mask] pair.
- o A status indication of either Advertise or DoNotAdvertise.
- o An external route tag.

### **3.3 Type-7 LSAs: NSSA External Link-State Advertisements**

External routes are imported into NSSAs as type-7 LSAs by NSSA AS boundary routers. A NSSA AS boundary routers (ASBR) is a router which has an interface associated with the NSSA and is exchanging routing information with routers belonging to another AS. Like ASBRs, a NSSA router indicates it is a NSSA ASBR by setting the E-bit in its router links advertisement. As with type-5 LSAs a separate type-7 LSA is originated for each destination network. To support NSSAs, the link-state database must therefore be expanded to contain a type-7 LSA.

Type 7-LSAs are identical to type-5 LSAs except for the following (see [section 12.4.4](#) "AS external links" in the OSPF specification).

1. The type field in the LSA header is 7.
2. Type-7 LSAs are only flooded within the originating NSSA. The flooding of type-7 LSAs follows the same rules as the flooding of type 1-2 LSAs.
3. Type-7 LSAs, which are kept within the NSSA's LSDB, are area specific. Type-5 LSAs, which are flooded to all type-5 capable areas, have global scope and are kept in the router's LSDB.
4. At the NSSA border router, selected type-7 LSAs are translated into type 5-LSAs and flooded into the backbone.



5. Type 7 LSAs have a propagate (P) bit which is used to flag the NSSA border router to translate the type-7 LSA into a type-5 LSA. Examples of how the P-bit is used for loop avoidance are in the following sections.
6. Those type-7 LSAs that are to be translated into type-5 LSAs must have their forwarding address set. Type-5 LSAs that have been translated from type-7 LSAs for the most part must contain a forwarding address. The exception to this is if the translation to a type-5 LSA is the result of an address range match, in which case the type-5 LSA will not contain a forwarding address (see [section 4.1](#) for details). The forwarding address contained in type-5 LSAs will result in more efficient routing to the AS external networks when there are multiple NSSA border routers. Having the forwarding address in the type-7 LSAs will ease the translation of type-7 into type-5 LSAs as the NSSA border router will not be required to compute the forwarding address.

If the network between the NSSA AS boundary router and the adjacent AS is advertised into OSPF as an internal OSPF route, the forwarding address should be the next hop address as is currently done in type-5 LSAs, but unlike type-5 LSAs if the intervening network is not advertised into OSPF as an internal OSPF route, the forwarding address should be any one of the router's active OSPF interface addresses.

Type-5 and type-7 metrics and path types are directly comparable.

### **[3.4](#) Originating Type-7 LSAs**

NSSA AS boundary routers may originate type-7 LSAs. All NSSA border routers must also be AS boundary routers since they all must have the capability of translating type-7 LSAs into type-5 LSAs (see [section 4.1 for the translation algorithm](#)). NSSA border routers must set the E-bit (external bit) as well as the B-bit (border bit) in their router (type-1) LSAs (both in the backbone and in the NSSA).

When a NSSA internal AS boundary router originates a type-7 LSA that it wants to be translated into a type-5 LSA by NSSA border routers (and subsequently flooded into the backbone), it must set the P-bit in the LSA header's option field and add a valid forwarding address in the type-7 LSA.

If a router is attached to another AS and is also a NSSA border router, it may originate both a type-5 and a type-7 LSA for the same network. The type-5 LSA will be flooded to the backbone (and all attached type-5 capable areas). The type-7 LSA will be flooded into

the NSSA. If this is the case, the P-bit must be reset in the type-7

NSSA so the type-7 LSA isn't again translated into a type-5 LSA by another NSSA border router. If the border router only originates a type-7 LSA, it may set the P-bit, thus allowing the network to be aggregated/propagated during the type-7 translation.

A type-7 default route (network 0.0.0.0) may be originated into the NSSA by a NSSA border router or by a NSSA ASBR which is internal to the NSSA. The type-7 default route originated by the NSSA border router must have the P-bit reset so that the default route originated by the NSSA border router will not find its way out of the NSSA into the rest of the AS system via another NSSA border router.

The type-7 default route originated by a NSSA ASBR which is not a NSSA border router may have the P-bit set. Type-7 routes which are originated by a NSSA border router will not get added to the routing tables of other NSSA border routers. If no NSSA internal router originates a type-7 default route in the NSSA, then, like stub areas, a type-7 LSA with default destination must be originated by all the NSSA's border routers in order to support inter-area AS routing and inter-AS routing.

Note that a NSSA area border router which originates a type-7 default route would have to originate a type-5 default route before other NSSA area border routers would see that default route. An internal type-7 default route whose P-bit is not set may only be installed on a NSSA border router when it is the only non-backbone OSPF area connected to it. This restriction protects the default routing of other areas attached to the NSSA border router as well any ISP agreements of the NSSArea.

In order for backbone summary internal routes to be preferred over external Type 7 routes, all implementations must support the optional import of summary LSAs from the backbone into a NSSA, with the exception of a (type-3) summary LSA for the default route. The import of summary LSAs is automatically activated when an type-7 default route is detected as originating from an internal NSSA ASBR, regardless of the no-summary setting. This protects the NSSA from routing intra-AS traffic out the AS via a type-7 default route.

Unlike the stub area case, a default route must not be injected into the NSSA as a summary (type-3) LSA. The reason for this is that the summary default route would be chosen over all more preferred type-7 default routes.

### **3.5 Calculating Type-7 AS External Routes**

This calculation must be run when type-7 LSAs are processed during the AS external route calculation. This calculation will also process type-5 LSAs and may replace [section 16.4](#) when processing type-5 LSAs. If

[section 16.4](#) is still used to process type-5 LSAs, NSSA ASBR routing table entries are not to be used for the ASBR address calculation of type-5 LSAs.

A NSSA border router should examine both type-5 LSAs and type-7 LSAs if either type-5 or type-7 routes need to be updated or recalculated. This is done as part of the AS external route calculation. A NSSA internal router should examine type-7 LSAs when type-7 routes need to be recalculated. When OSPF [Section 16.4.1](#) path preference is applied in step (6.c), NSSA and non-NSSA intra-AS paths have equal preference.

What follows is only a modest modification of the OSPF Version 2 Specification [Section 16.4](#). Original text is suffixed with <Moy>. NSSA specific text is suffixed with <NSSA>.

AS external routes are calculated by examining AS-external-LSAs, be they type-5 or type-7. Each of the AS-external-LSAs is considered in turn. Most AS-external-LSAs describe routes to specific IP destinations. An AS-external-LSA can also describe a default route for the Autonomous System (Destination ID = DefaultDestination, network/subnet mask = 0x00000000). For each AS-external-LSA <Moy with "or type-7">:

- (1) If the metric specified by the LSA is LSInfinity, or if the age of the LSA equals MaxAge, then examine the next LSA. <Moy>
- (2) If the LSA was originated by the calculating router itself, examine the next LSA. <Moy>
- (3) Call the destination described by the LSA N. N's address is obtained by masking the LSA's Link State ID with the network/subnet mask contained in the body of the LSA. Look up the routing table entries (potentially one per attached area) for the AS boundary router (ASBR) that originated the LSA. If no entries exist for router ASBR (i.e., ASBR is unreachable), do nothing with this LSA and consider the next in the list. <Moy>

Else if the destination is a type-7 default route (destination = DefaultDestination) and one of the following is true, then do nothing with this LSA and consider the next in the list:

- o The originator of the type-7 LSA and the calculating router are both NSSA border routers.
- o The calculating router is a border router, the LSA has its P-bit clear, and at least two non-backbone OSPF areas connect to the calculating router. <NSSA>

Else, this LSA describes an AS external path to destination N. Examine the forwarding address specified in the AS-external-LSA.

This indicates the IP address to which packets for the destination should be forwarded. <Moy>



If the forwarding address is set to 0.0.0.0, packets should be sent to the ASBR itself. <Moy>

If the current LSA is type-5, among the multiple non-NSSA ASBR routing table entries for the ASBR (both NSSA and non-NSSA ASBR entries might exist on an NSSA border router), select the preferred entry as follows <NSSA>:

If RFC1583Compatibility is set to "disabled", prune the set of routing table entries for the ASBR as described in OSPF [Section 16.4.1](#). In any case, among the remaining routing table entries, select the routing table entry with the least cost; when there are multiple least cost routing table entries the entry whose associated area has the largest OSPF Area ID (when considered as an unsigned 32-bit integer) is chosen. <Moy>

If the forwarding address is non-zero, look up the forwarding address in the routing table. The matching routing table entry must specify an intra-area or inter-area path; if no such path exists, do nothing with the LSA and consider the next in the list. <Moy>

- (4) Let X be the cost specified by the preferred routing table entry for the ASBR/forwarding address, and Y the cost specified in the LSA. X is in terms of the link state metric, and Y is a type 1 or 2 external metric. <Moy>
- (5) Now, look up the routing table entry for the destination N. If no entry exists for N, install the AS external path to N, with the next hop equal to the list of next hops to the ASBR/forwarding address, and advertising router equal to ASBR. If the external metric type is 1, then the path-type is set to Type-1 external and the cost is equal to X + Y. If the external metric type is 2, the path-type is set to Type-2 external, the link-state component of the route's cost is X, and the Type-2 cost is Y. <Moy>
- (6) Otherwise compare the AS external path described by the LSA with the existing paths in N's routing table entry, as follows. If the new path is preferred, it replaces the present paths in N's routing table entry. If the new path is of equal preference, it is added to N's routing table entry's list of paths. <Moy>

Preference is defined as follows:

- (a) Intra-area and inter-area paths are always preferred over AS external paths. <Moy>

(b) Type 1 external paths are always preferred over type 2

external paths. When all paths are type 2 external paths, the paths with the smallest advertised type 2 metric are always preferred. <Moy>

- (c) If the new AS external path is still indistinguishable from the current paths in N's routing table entry, and RFC1583Compatibility is set to "disabled", select the preferred paths based on the intra-AS paths to the ASBR/forwarding addresses, as specified in [Section 16.4.1](#). <Moy>
- (d) If the new AS external path is still indistinguishable from the current paths in N's routing table entry, select the preferred path based on a least cost comparison. Type 1 external paths are compared by looking at the sum of the distance to the forwarding address and the advertised type 1 metric (X+Y). Type 2 external paths advertising equal type 2 metrics are compared by looking at the distance to the forwarding addresses. <Moy>
- (e) If the new paths are still indistinguishable the following priorities apply (listed from highest to lowest) for breaking the tie.
  - a. Any type-5 LSA.
  - b. A type-7 LSA with the P-bit set and the forwarding address non-zero.
  - c. Any other type-7 LSA. <NSSA>

### **[3.6](#) Incremental Updates**

Incremental updates for type-7 LSAs should be treated the same as incremental updates for type-5 LSAs (see [section 16.6](#) of the OSPF specification). That is, if a new instance of a type-7 LSA is received it is not necessary to recalculate the entire routing table. If there is already an OSPF internal route to the destination represented by the type-7 LSA, no recalculation is necessary. Otherwise, the procedure in the proceeding section will have to be performed but only for the external routes (type-5 and type-7) whose networks describe the same networks as the newly received LSA.



## **4.0 Originating Type-5 LSAs**

### **4.1 Translating Type-7 LSAs Into Type-5 LSAs**

This step is performed as part of the NSSA's Dijkstra calculation after type-5 and type-7 routes have been calculated. If the calculating router is not a NSSA border router this translation algorithm should be skipped. It is not recommended that multiple NSSA border routers perform the translation unless the efficient routing of packets through area 0 to a NSSA partitioned by aggregation requires it. It is normally sufficient to have only one NSSA border router perform the translation. Excessive numbers of type-7 translators unnecessarily increase the size of the OSPF link state data base.

A new bit called bit Nt is added to the router links advertisement. All NSSA area border routers which are performing the translation set bit Nt in their router links advertisement into the NSSA.

A new parameter called the NSSATranslateState is added to the OSPF area data structure. If a NSSA border router has

NSSATranslateState = enabled

then this router always translates Type-7 LSAs into Type-5 LSAs for the NSSA.

If a NSSA border router has

NSSATranslateState = disabled

and no NSSA border router for the NSSA has bit Nt set in its router links advertisement and this router has the highest router ID among the NSSA border router set, then this router performs the translation of type-7 LSAs into type-5 LSAs for the NSSA and NSSATranslateState should be set to elected.

Otherwise the translation algorithm should not be performed and NSSATranslateState should remain set to disabled.

Note that during the translation of type-7 LSAs into aggregated type-5 LSAs, the highest router ID is not necessarily the best choice for an advertising router as all packets which are forwarded by a type-5 LSA routing table entry originating from an aggregated type-7 LSA translation are sent to the translator (As described below, the forwarding address is not set in type-5 LSAs which originate from aggregated type-7 LSAs). The NSSATranslateState allows the network designer to configure the most cost effective route to the NSSA's type-5 LSAs which originate from the aggregation of type-7 LSAs. In cases of aggregate partitioning of the NSSA, multiple translators may be required to effect efficient routing.



Indeed, all NSSA border routers could be set to perform translation, if required.

All installed type-7 LSA's should be examined including those type-7s originated by the router itself (a set which may differ between NSSA border routers of the same NSSAarea). This allows NSSA border routers to propagate and/or aggregate locally originated type-7s during the translation. Locally originated type-7s are skipped during the external route calculation. Their installation status is external to OSPF.

If the type-7 LSA (associated with the route being examined) has the P-bit set and a non-zero forwarding address, the following steps should be taken.

The translation procedure must first check for a configured type-7 address range. Recall that a type-7 address range consists of an [address,mask] pair and a status indication of either Advertise or DoNotAdvertise. At most a single type-5 LSA is made for each range. If the route being examined falls within the type-7 address range, (i.e., the [address,mask] pair of the route is equal to or a more specific instance of the [address,mask] pair of the type-7 address range), one of following three actions may take place.

1. When the range's status indicates Advertise and the route's address and mask are equal to the address and mask of the type-7 range, a type-5 LSA should be originated if
  - o there currently is no type-5 LSA originated from this router corresponding to the type-7 LSA, or there is and
  - o the path type or the metric in the corresponding type-5 LSA is different from the type-7 LSA or
  - o the forwarding address in the corresponding type-5 LSA is different from the type-7 LSA.

The newly originated type-5 LSA will describe the same network and have the same network mask, metrics, forwarding address, external route tag and path type as the type-7 LSA, however, the advertising router field will be the router ID of this area border router.

2. When the range's status indicates Advertise and the route's address or mask indicates a more specific route (i.e., the route's address is subsumed by the range or the route has a longer mask), a type-5 LSA is generated with link-state ID



equal to the range's address (if necessary, the link-state ID can also have one or more of the range's "host" bits set; see [Appendix F](#) of the OSPF specification for details), the network mask, external route tag and path type will be set to the configured type-7 range values. The advertising router field will be the router ID of this area border router. The forwarding address will not be set. The path type should always be set to the highest path type that is subsumed by the net range. The metric for the type-5 LSA will be set as follows:

- o if the path type is external type 2, the type-5 metric should be set to the largest type-7 metric subsumed by this net range + 1.
- o if the path type is external type 1, the type-5 metric should be set to the largest metric.

For example, given a net range of [10.0.0.0, 255.0.0.0] for an area that has type-7 routes of:

```
10.1.0.0 path type 1, metric 10
10.2.0.0 path type 1, metric 11
10.3.0.0 path type 2, metric 5
```

a type-5 LSA will be generated with a path type of 2 and a metric of 6.

As another example, given a net range of [10.0.0.0, 255.0.0.0] for an area that has type-7 routes of:

```
10.1.0.0 path type 1, metric 10
10.2.0.0 path type 1, metric 11
10.3.0.0 path type 1, metric 5
```

a type-5 LSA will be generated with a path type of 1 and a metric of 11.

These metric and path type rules will avoid routing loops in the event that path type 1 and 2 are both used within the area.

3. When the range's status indicates DoNotAdvertise, the type-5 LSA is suppressed and the component networks remain hidden from the rest of the AS.

By default (given that the P-bit is set and the LSA has a non-zero forwarding address) if a network is not contained in any



explicitly configured address range, a type-7 to type-5 LSA translation will occur.

A new instance of a type-5 LSA should be originated and flooded to all attached type-5 capable areas if

- o there currently is no type-5 LSA originated from this router corresponding to the type-7 LSA, or there is and
- o the path type or the metric in the corresponding type-5 LSA is different from the type-7 LSA or
- o the forwarding address in the corresponding type-5 LSA is different from the type-7 LSA.

The newly originated type-5 LSAs will describe the same network and have the same network mask, metrics, forwarding address, external route tag and path type as the type-7 LSA. The advertising router field will be the router ID of this area border router.

As with all newly originated type-5 LSAs, a type-5 LSA that is the result of a type-7 to type-5 translation (type-7 range or default case) is flooded to all attached type-5 capable areas.

#### **4.2 Flushing Translated Type-7 LSAs**

If a NSSA border router has translated a type-7 LSA to a type-5 LSA that should no longer be translated, the type-5 LSA should be flushed (set to MaxAge and flooded). The translated type-5 LSA should be flushed whenever the routing table entry that caused the translation changes so that either the routing table entry is unreachable or the entry's associated LSA is not a type-7 with the P-bit set and a non-zero forwarding address.

If a NSSA border router is translating type-7 LSA's into type-5 LSA's and

NSSATranslateState = elected

and a new router links advertisement arrives in the NSSA with bit Nt set, then it should flush any of its type-5 LSAs originating from translated type-7 LSAs and readvertise its router links advertisement into the NSSA with bit Nt clear.

This flushing keeps the number of type-7 translators at the minimum number configured, either explicitly or by default. Since the default translator election process only occurs in the absence of a translator, should a router with a higher router ID join the NSSA border router set it will not necessarily assume translator duties. Indeed, a previously default elected translator should continue to perform translation duties until supplanted by an NSSA border router whose Nt bit is set to true.

Such an event might happen due to by setting the NSSATranslateState to

enabled or a topological rejoining of a partitioned NSSA. This behavior reduces the flushing of translated type-7 LSA's in the AS.

Any change in the membership of the NSSA border router set or the setting of their Nt bits resulting from updated router links advertisements should force a NSSA border router whose NSSATranslateState is not set to recheck and possibly elect or disable its type-7 translation status.

## **5.0 Acknowledgments**

This document was produced by the OSPF Working Group, chaired by John Moy.

In addition, the comments of the following individuals are also acknowledged:

Phani Jajjarvarpu	cisco
Dino Farinacci	cisco
Jeff Honig	Cornell University
John Moy	Proteon, Inc.
Doug Williams	IBM

## **6.0 References**

- [1] Moy, J., "OSPF Version 2", [RFC 1583](#), Proteon, Inc., March 1994.
- [2] Moy, J., "Multicast Extensions to OSPF", [RFC 1584](#), Proteon, Inc., Proteon, Inc., March 1994.

## **7.0 Security Considerations**

Security issues are not discussed in this memo.



## **8.0 Authors' Addresses**

This update uses much of the original text from [RFC 1587](#) authored by

Rob Coltun  
Fore Systems  
6905 Rockledge Drive  
Suite 800  
Bethesda, Maryland 20817

Phone: (301) 571-2521  
EMail: rcoltun@fore.com

Vince Fuller  
BBN Planet  
3801 East Bayshore Road  
Palo Alto, California 94303

Phone: (415) 528-7227  
EMail: vaf@wr.BBNPlanet.com

New Sections, edits and revisions are authored by

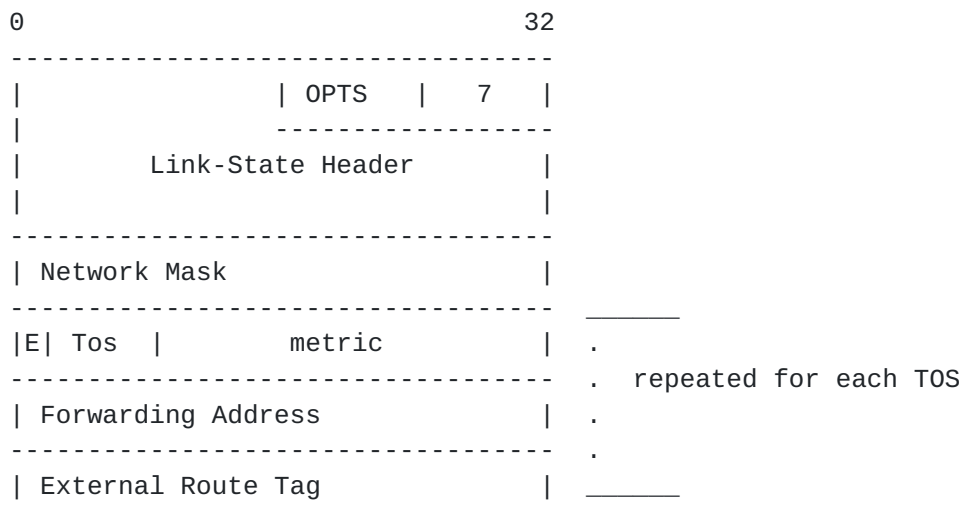
Pat Murphy  
US Geological Survey  
345 Middlefield Road  
Menlo Park, California 94560

Phone: (415) 329-4044  
EMail: pmurphy@usgs.gov

in support of the new features suggested by Pat Murphy.



## Appendix A: Type-7 Packet Format



The definitions of the link-state ID, network mask, metrics and external route tag are the same as the definitions for the type-5 LSAs (see A.4.5 in the OSPF specification) except for:

#### The Forwarding Address

If the network between the NSSA AS boundary router and the adjacent AS is advertised into OSPF as an internal OSPF route, the forwarding address should be the next hop address but if the intervening network is not advertised into OSPF as an internal OSPF route, the forwarding address should be any one of the router's active OSPF interface addresses.



## Appendix B: The Options Field

The OSPF options field is present in OSPF Hello packets, Database Description packets and all link-state advertisements. See [appendix A.2](#) in the OSPF specification for a description of option field. Six bits are assigned but only two (the E-bit and the N/P bit) are described completely in this section.

```
-----
| * | * | DC | EA | N/P | MC | E | T |
-----
```

### The Type-7 LSA options field

- E-bit:** Type-5 AS external link advertisements are not flooded into/through OSPF stub areas and NSSAs. The E-bit ensures that all members of a stub area agree on that area configuration. The E-bit is meaningful only in OSPF Hello packets. When the E-bit is reset in the Hello packet sent out a particular interface, it means that the router will neither send nor receive type-5 AS external link state advertisements on that interface (in other words, the interface connects to a stub area or NSSArea). Two routers will not become neighbors unless they agree on the state of the E-bit.
- N-bit:** The N-bit describes the router's NSSArea capability. The N-bit is used only in Hello packets and ensures that all members of a NSSArea agree on that area's configuration. When the N-bit is set in the Hello packet and sent out a particular interface, it means that the router will send and receive type-7 LSAs on that interface. Two routers will not form an adjacency unless they agree on the state of the N-bit. If the N-bit is set in the options field, the E-bit must be reset.
- P-bit:** The P-bit is used only in the type-7 LSA header. It flags the NSSA border router to translate the type-7 LSA into a type-5 LSA.



## Appendix C: Router-LSAs

Router-LSAs are the Type 1 LSAs. Each router in an area originates a router-LSA. The LSA describes the state and cost of the router's links (i.e., interfaces) to the area. All of the router's links to the area must be described in a single router-LSA. For details concerning the construction of router-LSAs, see the OSPF Specification, [Section 12.4.1](#).

0										1										2										3																			
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
LS age										Options										1																													
Link State ID																																																	
Advertising Router																																																	
LS sequence number																																																	
LS checksum															length																																		
0 Nt W V E B										0										# links																													
Link ID																																																	
Link Data																																																	
Type										# TOS										TOS 0 metric																													
TOS										0										metric																													
...																																																	
TOS										0										metric																													
Link ID																																																	
Link Data																																																	
...																																																	

In router-LSAs, the Link State ID field is set to the router's OSPF Router ID. The T-bit is set in the LSA's Option field if and only if the router is able to calculate a separate set of routes for each IP TOS. Router-LSAs are flooded throughout a single area only.



**bit V**

When set, the router is an endpoint of one or more fully adjacent virtual links having the described area as Transit area (V is for virtual link endpoint).

**bit E**

When set, the router is an AS boundary router (E is for external).

**bit B**

When set, the router is an area border router (B is for border).

**bit W**

When set, the router is a wild-card multicast receiver (W is for for wild).

**bit Nt**

When set, the router is a NSSA border router which translates type-7 LSAs into type-5 LSAs (Nt is for NSSA translation).

The remainder of the router links specification is as defined in the OSPF Specification, Section A.4.2

## Appendix D: Configuration Parameters

[Appendix C.2](#) in the OSPF specification lists the area parameters. The area ID, list of address ranges for type-3 summary routes and authentication type remain unchanged. [Section 3.2](#) of this document lists the configuration parameters for type-7 address ranges. The following parameter is added to the NSSArea data structure:

## NSSATranslateState

This parameter indicates whether or not an NSSA Border router is performing NSSA translation of type-7 LSAs into type-5 LSAs and flooding the translated type-5 LSAs into the AS. If the parameter is set to "enabled", translation is always performed. If the parameter is set to "elected", it means translation is being performed because the router was chosen during the default election process. If the parameter is set to "disabled" the NSSA border router is not currently performing type-7 translation. "disabled" is the default setting. (See [Section 4.1](#).)

For NSSAs the external capabilities of the area must be set to accept type-7 external routes. Additionally there must be a way of configuring the NSSA border router to send a default route into the NSSArea using a specific metric (type 1 or type 2 and the actual cost).



## Appendix E: Differences from [RFC 1587](#)

This section documents the differences between this memo and [RFC 1587](#). All differences are backward-compatible. Implementations of this memo and of [RFC 1587](#) will interoperate.

### E.1 Enhancements to OSPF summary LSAs.

The flooding of backbone summary LSAs (type-3 LSAs) into the NSSA is now optional. In [RFC 1587](#) the flooding of backbone summary LSAs was mandated in order to guarantee inter-area routes are preferred over external routes. The current recommended default behavior is to flood summary LSAs.

See sections [2.2](#) and [3.4](#) for details.

### E.2 Changes to the type-7 AS external routing calculation.

The type-7 external route calculation has been revised. Most notably:

- o The path preference defined in OSPF [Section 16.4.1](#) has been included.
- o A type-7 default route with the P-bit clear will not be installed on a NSSA border router which connects multiple non-backbone areas. This protects the default routing of other OSPF Areas as well as any ISP agreements of the originating NSSArea.
- o The type-7 AS external route calculation may now compute type-5 external paths.

See [Section 3.5](#) for details.

### E.3 Changes to translating type-7 LSAs into type-5 LSAs

NSSA border routers which perform the translation are now optionally configurable, with more than one allowed. This allows the network designer to choose the most cost effective intra-AS route for NSSA type-7 aggregates propagated into the AS. Furthermore, since different NSSA border routers may install different sets of type-7 LSA routes, the cost effectiveness of non-aggregated type-7 propagation may also be maximized.

All installed type-7 LSA's including those originated by the NSSA border router are processed for translation. This allows the NSSA border router to propagate and/or aggregate locally originated type-7s utilizing the type-5 translation process. NSSA [RFC 1587](#) required locally originated type-7 LSAs be paired



with locally originated type-5 LSAs for the external path to be seen by the AS (Note that locally originated type-7s are skipped during the external route calculation).

The default translator election process occurs only in the absence of a translator amongst the NSSA border router set.

See [Section 4.1](#) for details.

#### E.4 Changes to flushing translated type-7 LSAs

A NSSA border router which was elected by the default selection process of [RFC 1587](#), terminates its translation duties and flushes its translated type-7 LSAs from the AS when another translator presents itself in the NSSA. This keeps the number of translators at a minimum.

See [Section 4.2](#) for details.

