

## OSPF Database Overflow

### Status of this Memo

This memo defines an Experimental Protocol for the Internet community. This memo does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

### Abstract

Proper operation of the OSPF protocol requires that all OSPF routers maintain an identical copy of the OSPF link-state database. However, when the size of the link-state database becomes very large, some routers may be unable to keep the entire database due to resource shortages; we term this "database overflow". When database overflow is anticipated, the routers with limited resources can be accommodated by configuring OSPF stub areas and NSSAs. This memo details a way of gracefully handling unanticipated database overflows.

This memo is a product of the OSPF Working Group. Please send comments to [ospf@gated.cornell.edu](mailto:ospf@gated.cornell.edu).

### Table of Contents

1.	Overview .....	2
2.	Implementation details .....	3
2.1	Configuration .....	3
2.2	Entering OverflowState .....	4
2.3	Operation while in OverflowState .....	5
2.3.1	Modifications to Flooding .....	5
2.3.2	Originating AS-external-LSAs .....	6
2.3.3	Receiving self-originated LSAs .....	6
2.4	Leaving OverflowState .....	6
3.	An example .....	6
4.	Administrative response to database overflow .....	7
5.	Operational experience .....	8
6.	Possible enhancements .....	8
A.	Related MIB parameters .....	8
	References .....	9
	Security Considerations .....	9
	Author's Address .....	9

## 1. Overview

OSPF requires that all OSPF routers within a single area maintain an identical copy of the OSPF link-state database. However, when the size of the link-state database becomes very large, some routers may be unable to keep the entire database due to resource shortages; we term this "database overflow". For example, a regional network may have a very large OSPF database because it is importing a large number of external routes into OSPF. Unless database overflow is handled correctly, routers will end up with inconsistent views of the network, possibly leading to incorrect routing.

One way of handling database overflow is to encase routers having limited resources within OSPF stub areas (see Section 3.6 of [1]) or NSSAs ([2]). AS-external-LSAs are omitted from these areas' link-state databases, thereby controlling database size.

However, unexpected database overflows cannot be handled in the above manner. This memo describes a way of dynamically limiting database size under overflow conditions. The basic mechanism is as follows:

- (1) A parameter, `ospfExtLsdbLimit`, is configured in each router indicating the maximum number of AS-external-LSAs (excluding those describing the default route) that are allowed in the link-state database. This parameter must be the same in all routers in the routing domain (see Section 2.1); synchronization of the parameter is achieved via network management.
- (2) In any router's database, the number of AS-external-LSAs (excluding default) is never allowed to exceed `ospfExtLsdbLimit`. If a router receives a non-default AS-external-LSA that would cause the limit of `ospfExtLsdbLimit` to be exceeded, it drops the LSA and does NOT acknowledge it.
- (3) If the number of non-default AS-external-LSAs in a router's database hits `ospfExtLsdbLimit`, the router a) flushes all non-default AS-external-LSAs that it has itself originated (see Section 2.2) and b) goes into "OverflowState".
- (4) While in OverflowState, the router refuses to originate any non-default AS-external-LSAs (see Section 2.3.2).
- (5) Optionally, the router can attempt to leave OverflowState after the configurable parameter `ospfExitOverflowInterval` has elapsed since entering OverflowState (see Section 2.4). Only at this point can the router resume originating non-default AS-external-LSAs.

The reason for limiting non-default AS-external-LSAs, but not other LSA types, is twofold. First of all, the non-default AS-external LSAs are the most likely to dominate database size in those networks with huge databases (e.g., regional networks; see [5]). Second, the non-default AS-external-LSAs can be viewed as "optional" in the following sense: the router can probably be monitored/reconfigured without them. (However, using similar strategies, other LSA types can also be limited; see Section 5.)

The method of dealing with database overflow described herein has the following desirable properties:

- o After a short period of convergence, all routers will have identical link-state databases. This database will contain less than `ospfExtLsdbLimit` non-default AS-external-LSAs.
- o At all times, routing WITHIN the OSPF Autonomous System will remain intact. Among other things, this means that the routers will continue to be manageable.
- o Default routing to external destinations will also remain intact. This hopefully will mean that a large amount of external connectivity will be preserved, although possibly taking less efficient routes.
- o If parameter `ospfExitOverflowInterval` is configured, the OSPF system will recover fully and automatically (i.e., without network management intervention) from transient database overflow conditions (see Section 2.4).

## 2. Implementation details

This section describes the mechanism for dealing with database overflow in more detail. The section is organized around the concept `OverflowState`, describing how routers enter the `OverflowState`, the operation of the router while in `OverflowState`, and when the router leaves `OverflowState`.

### 2.1. Configuration

The following configuration parameters are added to support the database overflow functionality. These parameters are set by network management.

#### `ospfExtLsdbLimit`

When the number of non-default AS-external-LSAs in a router's link-state database reaches `ospfExtLsdbLimit`, the router enters `OverflowState`. The router never holds more

than `ospfExtLsdbLimit` non-default AS-external-LSAs in its database.

`ospfExtLsdbLimit` MUST be set identically in all routers attached to the OSPF backbone and/or any "regular" OSPF area. (This memo does not pertain to routers contained within OSPF stub areas nor NSSAs, since such routers do not receive AS-external-LSAs.) If `ospfExtLsdbLimit` is not set identically in all routers, then when the database overflows: 1) the routers will NOT converge on a common link-state database, 2) incorrect routing, possibly including routing loops, will result and 3) constant retransmission of AS-external-LSAs will occur. Identical setting of `ospfExtLsdbLimit` is achieved/ensured by network management.

When `ospfExtLsdbLimit` is set in a router, the router must have some way to guarantee that it can hold that many non-default AS-external-LSAs in its link-state database. One way of doing this is to preallocate resources (e.g., memory) for the configured number of LSAs.

#### `ospfExitOverflowInterval`

The number of seconds that, after entering `OverflowState`, a router will attempt to leave `OverflowState`. This allows the router to again originate non-default AS-external-LSAs. When set to 0, the router will not leave `OverflowState` until restarted. The default setting for `ospfExitOverflowInterval` is 0.

It is not necessary for `ospfExitOverflowInterval` to be configured the same in all routers. A smaller value may be configured in those routers that originate the "more important" AS-external-LSAs. In fact, setting `ospfExitOverflowInterval` the same may cause problems, as multiple routers attempt to leave `OverflowState` simultaneously. For this reason, the value of `ospfExitOverflowInterval` must be "jittered" by randomly varying its value within the range of plus or minus 10 percent before using.

## 2.2. Entering `OverflowState`

The router enters `OverflowState` when the number of non-default AS-external-LSAs in the database hits `ospfExtLsdbLimit`. There are two cases when this can occur. First, when receiving an LSA during flooding. In this case, an LSA which does not already have a database instance is added in Step 5 of Section 13 of [1]. The

second case is when the router originates a non-default AS-external-LSA itself.

Whenever the router enters OverflowState it flushes all non-default AS-external-LSAs that it itself had originated. Flushing is accomplished through the premature aging scheme described in Section 14.1 of [1]. Only self-originated LSAs are flushed; those originated by other routers are kept in the link-state database.

### 2.3. Operation while in OverflowState

While in OverflowState, the flooding and origination of non-default AS-external-LSAs are modified in the following fashion.

#### 2.3.1. Modifications to Flooding

Flooding while in OverflowState is modified as follows. If in Step 5 of Section 13 of [1], a non-default AS-external-LSA has been received that a) has no current database instance and b) would cause the count of non-default AS-external-LSAs to exceed `ospfExtLsdbLimit`, then that LSA is discarded. Such an LSA is not installed in the link-state database, nor is it acknowledged.

When all routers have identical values for `ospfExtLsdbLimit` (as required), the above flooding modification will only be invoked during a short period of convergence. During convergence, there will be retransmissions of LSAs. However, after convergence the retransmissions will cease, as the routers settle on a database having less than `ospfExtLsdbLimit` non-default AS-external-LSAs.

In OverflowState, non-default AS-external-LSAs ARE still accepted in the following conditions:

- (1) If the LSA updates an LSA that currently exists in the router's link-state database.
- (2) LSAs having LS age of `MaxAge` are always accepted. The processing of these LSAs follows the procedures described in Sections 13 and 14 of [1].
- (3) If adding the LSA to the router's database would keep the number of non-default AS-external-LSAs less than or equal to `ospfExtLsdbLimit`, the LSA is accepted.

### 2.3.2. Originating AS-external-LSAs

Originating AS-external-LSAs is described in Section 12.4.5 of [1]. When a router is in OverflowState, it does not originate non-default AS-external-LSAs. In other words, the only AS-external-LSAs originated by a router in OverflowState have Link State ID 0.0.0.0.

### 2.3.3. Receiving self-originated LSAs

Receiving self-originated LSAs is described in Section 13.4 of [1]. When in OverflowState, a router receiving a self-originated non-default AS-external-LSA responds by flushing it from the routing domain using the premature aging scheme described in Section 14.1 of [1].

## 2.4. Leaving OverflowState

If `ospfExitOverflowInterval` is non-zero, then as soon as a router enters OverflowState, it sets a timer equal to the value of `ospfExitOverflowInterval` (plus or minus a random value in the range of 10 percent). When this timer fires, the router leaves OverflowState and begins originating non-default AS-external-LSAs again.

This allows a router to automatically recover from transient overflow conditions. For example, an AS boundary router that imports a great many AS-external-LSAs may crash. Other routers may then start importing the routes, but until the crashed AS boundary router is either a) restarted or b) its AS-external-LSAs age out, there will be a much larger database than usual. Since such an overflow is guaranteed to go away in `MaxAge` seconds (1 hour), automatic recovery may be appropriate (and fast enough) if the overflow happens off-hours.

As soon as the router leaves OverflowState, it is again eligible to reenter OverflowState according to the text of Section 2.2.

## 3. An example

As an example, suppose that a router implements the database overflow logic, and that its `ospfExtLsdbLimit` is 10,000 and its `ospfExitOverflowInterval` is set to 600 seconds. Suppose further that the router itself is originating 400 non-default AS-external-LSAs, and that the current number of non-default AS-external-LSAs in the router's database is equal to 9,997.

Next, it receives a Link State Update packet from a neighbor, containing 6 non-default AS-external-LSAs, none of which have current database copies. The first two LSAs are then installed in the database. The third LSA is also installed in the database, but causes the router to go into OverflowState. Going into OverflowState causes the router to flush (via premature aging) its 400 self-originated non-default LSAs. However, these 400 LSAs are still considered to be part of the link-state database until their re-flooding (with age set to MaxAge) is acknowledged (see Section 14 of [1]); for this reason, the last three LSAs in the received update are discarded without being acknowledged.

After some small period of time all routers will converge on a common database, having less than 10,000 non-default AS-external-LSAs. During this convergence period there may be some link-state retransmissions; for example, the sender of the above Link State Update packet may retransmit the three LSAs that were discarded. If this retransmission happens after the flushing of the 400 self-originated LSAs is acknowledged, the 3 LSAs will then be accepted.

Going into OverflowState also causes the router to set a timer that will fire some time between 540 and 660 seconds later. When this timer fires, the router will leave OverflowState and re-originate its 400 non-default AS-external-LSAs, provided that the current database has less than 9600 (10,000 - 400) non-default AS-external-LSAs. If there are more than 9600, the timer is simply restarted.

#### 4. Administrative response to database overflow

Once the link-state database has overflowed, it may take intervention by network management before all routing is restored. (If the overflow condition is transient, routing may be restored automatically; see Section 2.4 for details.) An overflow condition is indicated by SNMP traps (see Appendix B). Possible responses by a network manager may include:

- o Increasing the value of `ospfExtLsdbLimit`. Perhaps it had been set too conservatively, and the routers are able to support larger databases than they are currently configured for.
- o Isolating routers having limited resources within OSPF stub areas or NSSAs. This would allow increasing the value of `ospfExtLsdbLimit` in the remaining routers.
- o Reevaluating the need to import certain external routes. If `ospfExtLsdbLimit` cannot be increased, the network manager will want to make sure that the more important routes continue to be imported; this is accomplished by turning off the importing of

less important routes.

#### 5. Operational experience

The database overflow scheme described in this memo has been implemented in the Proteon router for a number of years, with the following differences. First, the router did not leave OverflowState until it was restarted (i.e., ospfExitOverflowInterval was always 0). Second, default AS-external-LSAs were not separated from non-default AS-external-LSAs. Operationally the scheme performed as expected: during overflow conditions, the routers converged on a common database having less than a configured number of AS-external-LSAs.

#### 6. Possible enhancements

Possible enhancements to the overflow scheme include the following:

- o Other LSA types, with the exception of the transit LSAs (router-LSAs and network-LSAs), could be limited in a similar fashion. For example, one could limit the number of summary-LSAs, or group-membership-LSAs (see [6]).
- o Rather than flushing all of its non-default AS-external-LSAs when entering OverflowState, a router could flush a fixed number whenever the database size hits ospfExtLsdbLimit. This would allow the router to prioritize its AS-external-LSAs, flushing the least important ones first.

#### A. Related MIB parameters

The following OSPF MIB variables have been defined to support the database overflow procedure described in this memo (see [4] for more information):

##### ospfExtLsdbLimit

As in Section 2.1 of this memo, the maximum number of non-default AS-external-LSAs that can be stored within the database. If set to -1, there is no limit.

##### ospfExitOverflowInterval

As in Section 2.1 of this memo, the number of seconds that, after entering OverflowState, a router will attempt to leave OverflowState. This allows the router to again originate non-default AS-external-LSAs. When set to 0, the router will not leave OverflowState until restarted.



`ospfLsdbOverflow`

A trap indicating that the number of non-default AS-external-LSAs has exceeded or equaled `ospfExtLsdbLimit`. In other words, this trap indicates that the router is entering `OverflowState`.

`ospfLsdbApproachingOverflow`

A trap indicating that the number of non-default AS-external-LSAs has exceeded ninety percent of "`ospfExtLsdbLimit`".

## References

- [1] Moy, J., "OSPF Version 2", RFC 1583, Proteon, Inc., March 1994.
- [2] Coltun, R., and V. Fuller, "The OSPF NSSA Option", RFC 1587, RainbowBridge Communications, Stanford University, March 1994.
- [3] Moy, J., Editor, "OSPF Protocol Analysis", RFC 1245, Proteon, Inc., July 1991.
- [4] Baker F., and R. Coltun, "OSPF Version 2 Management Information Base", Work in Progress.
- [5] Moy, J., Editor, "Experience with the OSPF Protocol", RFC 1246, Proteon, Inc., July 1991.
- [6] Moy, J., "Multicast Extensions to OSPF", RFC 1584, Proteon, Inc., March 1994.

## Security Considerations

Security issues are not discussed in this memo.

## Author's Address

John Moy  
Cascade Communications Corp.  
5 Carlisle Road  
Westford, MA 01886

Phone: 508-692-2600 Ext. 394  
Fax: 508-692-9214  
EMail: [jmoy@casc.com](mailto:jmoy@casc.com)