

Network Working Group  
Berkowitz  
Internet Draft  
Communications  
Expiration Date: January 2000  
June 1999

## **Techniques in OSPF-based Network**

Deployment

[draft-ietf-ospf-deploy-00.txt](#)

Status  
of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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''.

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

1.

Abstract

OSPF is the preferred interior routing protocol of the Internet. It is a complex protocol intended to deal with complex networks. While it is a powerful mechanism, it does not handle all situations, and its appropriate use may not be obvious to beginners. Standards track documents deal with protocol design, but deployment of OSPF in many enterprise networks has been limited by lack of information on best current practice information for interior routing. Best Current Practices documents have focused on general exterior connectivity. This memorandum is intended to complement the protocol specification by describing the experience-based, vendor-independent techniques of OSPF and complementary technologies in representative networks. Better understanding of the use of OSPF features to help exterior connectivity will help reduce the demand for complex user BGP configuration.

## **2. Introduction**

[RFC1812](#), Requirements for IPv4 Routers, says "a router that implements any routing protocol (other than static routes) MUST IMPLEMENT OSPF... A router MAY implement additional IGPs." This is a well-crafted statement, as it recognizes that static may be a useful complement to OSPF. All too often, network operators work from the limiting belief that if they use OSPF, it must run everywhere in their networks. They see the two-level hierarchy suggested by an Area 0 and a subordinate set of nonzero areas, and assume that networks using OSPF must be strictly hierarchical with only two hierarchical levels.

### **3. Barriers to Understanding OSPF Deployment**

Some terminology used in OSPF documents may not match more recent usage, or may be confusing to a reader encountering them the first time. OSPF, for example, uses the term "autonomous system" differently than the most recent BGP-oriented usage. OSPF's usage is discussed below.

In addition, the first-time reader should be careful to understand OSPF's use of the concept of external routes. Another point is not to confuse the hierarchy in a specific OSPF domain with the overall hierarchy of the enterprise in which it is used; you are not limited to the two apparent levels of OSPF, backbone versus non-backbone.

#### **[3.1](#) OSPF and Autonomous**

Systems

OSPF specifications assume an enterprise's set of networks equate to an autonomous system. These specifications use the term autonomous system in a manner inconsistent with the current usage, which speaks of an AS in terms of exterior routing. [RFC1930](#) defines an AS as a set of routers, under one or more administrations, that presents a common routing policy to the global Internet.

It is useful to define an OSPF domain as an Area 0 and some number of nonzero areas. In the broader sense, a routing domain is a set of routers, with a common set of routing mechanisms and routing policies, under a single administration.

A general observation: OSPF has a terminology problem here. "Inter-zone" communication is the role of what OSPF calls a "Autonomous System Border

Router," but the "zones" don't need, in reality, to be distinct AS. A term from the routing literature that probably is more recognized than "AS" or "zone" is "routing domain."

OSPF has a built-in mechanism for connecting different domains, the ASBR. ASBRs interconnect routing domains. Any OSPF router that advertises external routes is an ASBR. Examples include:

1. An OSPF process that accepts routes from another separate

OSPF process:

2. An OSPF process that accepts routes from another dynamic interior routing process:

Another option is to have static routes going to an enterprise core network from each OSPF Area 0, such that the Area 0's advertise the default route into their subordinate areas, but the ASBR containing the static route advertises the default into Area 0.

3. An OSPF process that accepts static routes and advertises the default.
4. An OSPF process that accepts routes from a BGP exterior routing process and advertises them into OSPF. Details of this are not shown, because configurations involving redistribution to and from OSPF usually involve fairly complex filtering and other mechanisms to avoid loops and injections of huge numbers of routes.

### **3.2 Thinking about Externals**

The path determination workload of an OSPF area is influenced

most  
strongly by the number of network, summary, and external LSAs  
involved  
in the routing computation. The various stub area schemes, stub,  
not-  
so-stubby, and the vendor-specific totally stubby, all help when a  
large  
number of externals would be injected into an area.

In medicine,  
there is a classic piece of advice from a physician to a  
patient who  
complains that some odd movement of his arm hurts: if that  
hurts, don't do  
it. If the nature of a routing environment is such that  
large numbers of  
externals do not exist, then there is no benefit to  
exploring stub area  
techniques to control the propagation of externals.

Most beginners tend to  
think of floods of externals into OSPF due  
primarily to Internet  
connectivity. They imagine nearly 50,000 routes of  
a current Internet  
default-free routing table overwhelming their  
routers. In practice, it  
would be extremely unlikely to find any  
significant benefit would come from  
injecting such a table into the OSPF  
routing system. In most cases, no  
more than a default route needs to  
be injected to achieve connectivity. Even  
if preferential exits for  
certain exterior destinations are desired, that  
is more likely to be an  
Interior BGP problem than an OSPF one.

In  
practice, a major source of external routes can come from one's  
own  
enterprise, as part of a migration from RIP or IGRP, or  
from  
static/default routes used to connect edge routers to the main  
routing  
system. Another way to think of this is that large numbers of  
external  
routes can come from other routing domains in your own  
organization.

If the intra-organizational external routes are injected

into OSPF  
through ASBRs connected to Area 0, stub techniques may be effective in  
controlling the number of externals that impose workload on nonzero  
areas. Even in this case, do summarize the externals as much as possible.

More challenging situations arise when the intra-organizational routes come in at the edge of a nonzero area. Assume, for example, that an organization has a large number of small offices, each with a single frame relay PVC that carries its traffic to the next level of the corporate routing hierarchy. Each of the small edge routers needs a default route toward the next level.

The next level, however, is composed of OSPF-speaking routers. Even though they are in a nonzero area, they are still ASBRs. Since these routers will not speak OSPF to the edge routers, these second-level routers will need to have static routes that point to these edge subnets, and then the static routing information needs to be advertised as external routes into OSPF. Remember that these externals can be summarized on the ASBR.

### **3.3 Hierarchy and OSPF**

It should never be forgotten there are two ways to achieve hierarchy and aggregation: address summarization, and physical topology. Many newcomers to OSPF think only of the first, since OSPF summarizes on ABRs and ASBRs. This leads to an incorrect assumption that OSPF has at best three levels of hierarchy:

1. Intra-area
2. Inter-area
- 3.

External

It is true that reduction in routing information and consequent reduction of the routing computation workload only can come at these levels. But there are other benefits of hierarchy, such as minimizing hop count, simplifying debugging as opposed to complex meshes, etc.

Inside a nonzero area, it is perfectly reasonable to have a hierarchy of internal routers, leading up to area border routers at the apex of the area hierarchy. Several small OSPF speakers at "branch offices" might single- or dual home to "district level" interior routers. Groups of district routers might in turn single or dual home to regional routers, and groups of regional routers might single or dual home to area border routers.

Hierarchy can begin anew inside area 0. ABRs can be concentrated topologically by homing to backbone routers, and the backbone routers may in turn home on ASBRs.

A higher-level, non-OSPF core network may usefully link the top hierarchical ASBRs of multiple OSPF routing domains. Such a core network could be statically routed or use BGP.

### 3.4

#### Virtual Links

A virtual link is a tunnel that has at least one end in Area

#### **0. Virtual**

Links (VL) are a mechanism that can be used within OSPF to handle certain connectivity patterns. The standard is a bit "soft" on their applications, and support engineers have seen a variety of VL applications. For some of the problems being raised, it appears better OSPF solutions may exist.

A matter of particular interest is the

potential advantages of NSSAs over some of the VL solutions to bringing in a new "community of interest." See the discussion below of VL applicability in other than backbone robustness.

The perception of virtual links' utility seems to be as a means of accomodating specical connectivity requiremnts "inside" a single "OSPF domain," the latter defined as a set of an Area 0 and some number of nonzero areas.

In some of these requirements, virtual links may be completely appropriate, one of several potential solutions, or definitiely not an appropriate solution.

Some designers may use virtual links to avoid other mechanisms that they do not like, such as defining the enterprise's network with multiple interconnected OSPF domains.

Virtual links are not necessarily the appropriate solution, but cases have been seen where they are used for:

- Protecting against Area 0 partitioning
- Making one area 0 following the merger of two enterprises, both of which ran independent OSPF
- Providing connectivity to a newly acquired enterprise whose best connectivity is to a router in a nonzero area of the acquiring enterprise

## **4. Area Sizing and Numbering Strategies**

### **4.1 Communities of Interest**

Select a set of clients and servers that primarily speak to one another.

Is the number of routers required, after growth projections have been applied, less than the vendor-recommended limit?

Let's review basic OSPF, emphasizing points that can help establish hierarchies of more than two levels. Let's also review some sizing considerations for areas, bearing in mind that this can be as much art and science. Experienced routing designers know they won't always hit the correct area structure, and expect to have to monitor and tune the design of a large OSPF system. Especially experienced routing designers know this tuning will primarily involve topology and bandwidth, rather than "knob twisting" for timers and buffers.

The guideline about limiting the number of nonzero areas in an ABR, in practice, means that large numbers of areas within a single OSPF system tend to require large and expensive numbers of ABRs. This is especially true attempting to avoid a single point of failure -- a single ABR -- per area.

Many factors go into choosing the number of OSPF speakers inside an area. A few conservative guidelines for these and other sizing considerations:

1. Begin setting up area definitions based on communities of interest. It is highly desirable if the majority of traffic can stay intra-area.

2. Do not exceed your vendor guidelines on routers per nonzero area. Counts from 50 to 200 routes are often cited, although much larger numbers have worked in specific environments. Use the lower numbers when media tend to "bounce" up and down. As we will see below, this doesn't preclude

a large number of routers in what can be considered an area, because the "stub" routers need not speak OSPF.

This is conservative, and many vendor implementations have wored well with larger numbers of OSPF speakers per area. The specific limit will vary with the OSPF software implementation, the processing power of the router platform, and the size and stability of the topological database.

If areas defined on community of interest contain too many routers under rule #2, consider splitting the area. Look for geographic/provider natural boundaries for splitting.

3. Think carefully about the relationships you want between the backbone and each non-backbone area. Some small OSPF environments put everything into Area **0. This can be reasonable if the only reasons for** using OSPF are fast convergence and flexible addressing among a small set of routers. The power of OSPF is only fully realized with multiple areas.

4. Area 0 is intended as a transit area. Capacity planning and troubleshooting tend to be easiest when there are no application servers in area 0. Given the connectivity of area 0, it may be reasonable to place network management or DNS servers there.

If the application topology is hierarchical, as might be found where a mainframes or server farms provides most application services, still use caution in putting this server in area 0. Mainframes or server farms often have local inter-server communications, such as backup, that should be kept out of area 0. It may be wise to put the central servers in a small area of their own.

5. Several techniques exist for setting up areas and controlling

advertisements in and out of them. You don't need to use the same techniques in every nonzero area. Remember that you don't use these simply to reduce the amount of routing traffic, but also for stability. Stability increases as the number of routing table recomputations decreases.

Summarization and the various kinds of stub areas reduce routing table recomputation by hiding specific routes. In other words, an interior router in Area 1, doesn't need to recompute the Area 1 routing table if a link goes up or down in Area 0. Media, especially, have a habit of bouncing up and down; hiding links outside my area reduces "route flapping" leading to recomputation.

An Area Border Router between Area 1 and Area 0 probably needs to know if a link in Area 0 changes state, but it doesn't need to know if a link in Area 51 changes state. The exception to this case is where finding the absolutely optimal path from Area 1 to Area 0 to Area 51 is more important than stability. This might be the case in some international networks with low-bandwidth links.

The reason to use the less general types of areas is they allow reducing the amount of routing information advertised from the backbone into nonzero areas.

Inside a given area, be it backbone or non-backbone, there can be a hierarchy among OSPF speakers. For example, it is perfectly reasonable to reduce the peering in the backbone by connecting the Area 0 side of ABRs to "collapsed backbone" router(s). In large networks, it can avoid problems due to excessive numbers of neighbors to

any given Area 0 router.

## **4.2 Numbering**

There is a widespread misperception that areas will only work with a single contiguous address range. Even a small amount of summarization, however, will considerably increase the stability of OSPF systems because it hides route flap.

A common recommendation, assuming that a single contiguous block, such as an /18 prefix is given to the enterprise, is to "bit split" the prefix such that the first high-order bits below the global prefix identifies the area. For a routing domain with four nonzero areas, this would allocate a /20 prefix to each area:

```
    /20 starting 00xxxx...   area 1
starting 01xxxx...   area 2
    starting 10xxxx...   area 3
starting 11xxxx...   area 4
```

This approach seems straightforward, but suffers from several limitations.

First, what about area 0? No space has been left for it. In any case, it is likely that a /20 assigned to area 0 would waste a great deal of space, since area 0 should have a small number of router interfaces in it.

One technique when using the bit-split method, and assuming registered addresses are used in the nonzero areas, is to use [RFC1918](#) private address space for area 0. This can be quite reasonable, because there are few legitimate reasons why an arbitrary external Internet host would need to access a backbone interface internal to an enterprise network. One possible criticism of this approach is that traceroutes that

traversed the backbone might show the private address space, but it is usually apparent when this is happening. Another reason why area 0 interfaces might need registered addresses is that the management of the network is outsourced. In outsourcing situations, the service provider commonly can assign some of its allocated address space for interfaces it will manage,

Another and more general approach is to bit-split to a level deeper than the number of areas. In the example above, there were 4 nonzero areas, and **4 /20 blocks. The /20 comes from it being the first power of two that can contain 4 areas.**

Consider, however, going 2 or 3 powers of two deeper. Divide the available address space into 8, 16, or 32 blocks. These, respectively, would be /21, /22, or /23 address prefixes.

One of these blocks can be assigned to Area 0. It is a reality that the number of users in individual areas will vary, so area 1 might, for example, need three /21 blocks, area 2 might need two such blocks, area 1 might need one. Several blocks can be reserved for growth as needed.

These blocks can still be summarized to avoid route flap.

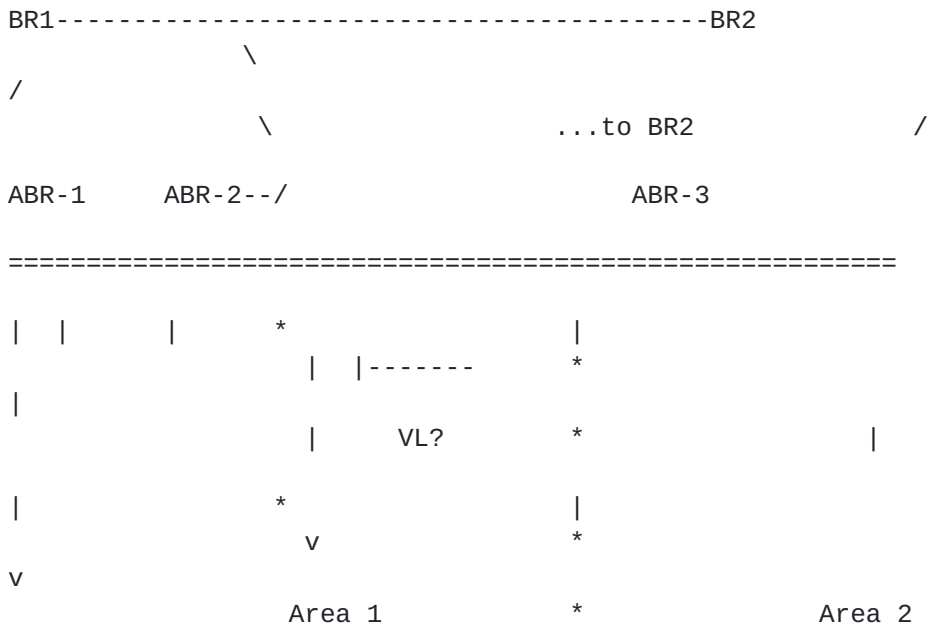
## **5. Increasing Backbone Reliability**

A failure in Area 0 is critical.

A single Area 0 has a single point of failure. Hopefully, the ASCII graphic shows this. Area 0 has two interconnected Border Routers, BR1 and BR2. To avoid single router points of failure, there are two ABRs, each with three interfaces: Area 0, Area 1, and Area 2.

## 5.1 VL Solution

If the BR1-BR2 link fails, a VL can be defined between the Area 1 interface of ABR-1 and between the Area 1 interface of ABR-2. This reconstitutes backbone connectivity with a tunnel through Area 1.



## 5.2

Alternative using Adding Circuit(s)

The preferred way to solve this problem is adding a parallel link between BR1 and BR2. Especially if these links can be per-packet load balanced, convergence would be extremely fast.

This solution, however, incurs additional cost for the additional circuit. Balanced against this cost is the performance impact the VL would have on the routers and links in the nonzero area through which the VL is tunneled. Those routers and links may have been engineered for the traffic estimates and performance goals under the workload of that single area, not of that area with backbone traffic added to it.

## 5.3

Alternative using non-OSPF network

An alternative approach, especially useful if demand OSPF is not available, is to split Area 0 and create two OSPF domains. Each domain has a static route that points to the address range in the other domain. This route is advertised into the local domain as a Type 1 external.

## **6. Backbones of Backbones**

The method described in [5.3](#) above can be generalized to give more effective levels of hierarchy in an overall network that uses OSPF for its dynamic routing.

True, OSPF's area structure has two levels, Area 0 and everything else. A routing architecture for an enterprise that uses OSPF, however, can be more than a simple hierarchy of backbone and non-backbone areas. We can extend its notion of hierarchy both above Area 0 and below the nonzero areas. Even inside an area, there are some forms of hierarchy.

Such an architecture also lends itself to evolution to a high-performance layer 2 or Multiprotocol Label Switching (MPLS) service in the core. It can also be reasonable to have interior routers that "concentrate" traffic, or act as collapsed backbones within an area.

It is also perfectly reasonable to have a broader OSPF routing environment, in which some routers do not speak OSPF but cooperate with those that do. At the "high" end, this involves redistribution of external routes into OSPF. At the "low" end, this involves static and default routes from stub routers to the lowest level of OSPF router inside a nonzero area. Let's talk about the lowest level first.

The lowest level is a "stub" or "edge" router with a single outgoing path, such as a branch office router with several LAN links, perhaps a STUN serial interface for IBM support, and a single WAN link. There are no routers on any of the LANs.

Such a router really doesn't need to run any dynamic routing protocol. It should default to a higher-level router.

Assuming the higher-level router has multiple links going toward the backbone, that router does need to run OSPF. It doesn't need to connect directly to the backbone.

## **7. Transition and Network Consolidation**

A management imperative that often occurs after the merger of two enterprises is eliminating the "expense of duplicate backbones." At a slightly more technical level, this means an OSPF design that has a single Area 0.

Among many OSPF users, it is a matter of faith and morals that there entire enterprise appear as a single OSPF domain. As a consequence, I've seen designs with a single area 0 spanning serveral continents, with each continent having significantly different line quality and speed, needs for demand backup circuits between parts of the area, etc.

Two companies, A, and B, merge. Each has an existing OSPF system, with its own area 0. These Area 0s are widely separated geographically, but they each have an Area 1 whose boundaries are in fairly close proximity. Each ABR has a single interface to Area 0.

In designing solutions, engineers should consider not only pure OSPF mechanisms, but the use of non-OSPF tunneling mechanisms. OSPF virtual

links are, of course, appropriate for some topologies.

Also, consider having separate OSPF domains linked into a "backbone of backbones," using static or BGP routing among the domains.

### **7.1 Non-OSPF tunneling**

Generic Route Encapsulation (GRE) can solve numerous OSPF topology problems, both in backbone area 0.0.0.0 and in non-backbone areas. GRE is appropriate

when the underlying transmission infrastructure is considered adequately secure. If, however, the traffic needs to cross the untrusted global Internet,

IPsec tunnel mode [[RFC2401](#)] may be more appropriate.

Let's say that the

two companies merge, and want to put their Western Region into the same area because the geographic divisions will share a good deal of information that does not need to cross the backbone.

Both areas, of course, could be connected separately to area 0.0.0.0. There might not be sufficient backbone bandwidth to carry the combined Western Region traffic.

If there were a management directive to merge these areas immediately, and dedicated WAN circuits between them were not immediately available, an alternative could be to tunnel between them. The next design issue would be deciding if the tunnels need to be secure. Tunneling, of course, will add overhead bytes to every routed packet, and will also impose additional processing workload on the routers at the endpoints of the tunnel.

In spite of the overhead it produces, tunneling may still be a good interim step to achieving desired connectivity during network redesign. OSPF has its own tunneling mechanism for routing information, the virtual link, which is appropriate for other topological problems.

### **7.2 Solution using VLS**

The two companies perceive they want "a single backbone". They do not wish to renumber every network and router, but are willing to merge two nonzero areas if that were helpful. They are also unwilling, at the present time, to merge their two backbones, principally due to the geographic distance between them.

They merge their Area 1's, renumbering appropriately adding physical connectivity between the company A area 1 and the company B area 1. They

now have a common Area 1, and then define a VL between the two ABRs.

There is now a "single backbone." Potentially, a large amount of inter-area traffic now flows through the combined Area 1. Neither Area **1** presumably was designed to support that flow, but that of the traffic of the community of routers for which it originally was built.

### **7.3 An Alternate Solution using External Links** between the Area 0's

Do not attempt to create a single Area 0. Rather than trying to merge the backbones or the area 1's, define an ASBR function in each Area 0 that has one or more static (or even BGP) routes to the ASBR of the other Area 0, and vice versa. Each ASBR advertises the default, and may or may not advertise the external routes to the other Area 0. Doing this requires adding one or more links between A-Area 0 and B-Area 0.

This requires no changes to any of the nonzero areas, which still can default to their existing Area 0. If tuning or topology changes are needed to handle traffic flows, most of this can be done at the core tier. Such tuning would involve a smaller set of routers not directly connected, in most cases, to end systems.

The core tier now consists of two Area 0's and a set of links between them.

### **7.4 An Alternate** Solution using Two Area 0's

Both companies have OSPF, but topological or other factors make it difficult to make a direct connection between Area 0s. In this case, establish a Company B ASBR that will advertise Company B routes to

Company A. This ASBR might be on the Company B Area 0, or in another area. In either case, it advertises routes to the other ASBR, using multiple static routes or BGP.

The company A ASBR can generate a default and advertise Company A addresses to Company B, using appropriate filtering. If the Company A ASBR is in a non-zero area, NSSA may be appropriate.

Company B has OSPF in place. Add an ASBR to its Area 0, with a static route to the Area 0 of Company A.

## **8. Transition of Legacy**

Routing Protocol Domains to OSPF

### **8.1 Problems of Integrating a Newly**

Acquired Enterprise; OSPF not used by new acquisition

Company A acquires a smaller Company B. Company B now runs RIP, or possibly has a limited OSPF implementation. Company B is not geographically close to the Area 0 of Company A.

There are several ways to deal with this situation while imposing minimal impact on Company B users.

### **8.2 An OSPF and VL**

Solution.

Replace RIP with OSPF on the Company B routers, assigning the Company B address space to a new nonzero area. Define a VL from one of the new area's routers, tunneled through an existing Company A area, that attaches to the backbone at the edge of the Company A area.

Eventually, provide direct connectivity to Area 0 from the Company B area, and delete the VL.

Disadvantages here include an immediate

conversion to OSPF, the cost of connectivity between the new area and the existing area that carries the VL, and the additional traffic load imposed on the nonzero area with the Area 0 connection.

### **8.3 Run RIP in Company B**

area, with link to Area 0 ASBR that learns RIP

Leave Company B running RIP. Run a link to a Company A, Area 0 router. Either redistribute the RIP routes into Area 0, or simply run RIP on the ASBR and have it originate the default into Area 0. If additional bandwidth or connectivity changes are needed to handle the Company B traffic, this primarily can be restricted to Area 0 where it is not visible to end users.

The disadvantage here is the cost of additional circuits from the Company B area to Area 0. Also, ASBR default injection into a nonzero area can create loops if conflicting defaults are being advertised elsewhere, as by an ABR.

### **8.4 Run RIP in**

Company B area, with an ASBR in the transit nonzero area of Company A.

Again Company B runs RIP. It connects to the nearest potential ASBR in an established Company A area, and is advertised as an external route into that area. This nonzero area then injects the external, possibly summarizing it, into Area 0. This method is less flexible in terms of advertising defaults than is 5.2, and also has the disadvantage of injecting external traffic into an area not originally designed to handle it. Since the area with the ASBR can no longer be stubby, it may be flooded with significant numbers of Type 5 LSAs from the backbone and

other areas.

## **8.5 ASBR Solution with NSSAs**

Again Company B runs RIP.

It connects to the nearest potential ASBR in an established Company A area, and is advertised as an external route into that area. This ASBR supports NSSA, and translates the RIP routes into Type 7 LSAs, which are sent to ABRs of this area and translated to externals in Area 0.

The established nonzero area maintains its generally stubby quality, and will not be flooded by unnecessary Type 5 LSAs.

## **9. Traffic Management**

In the real world, situations may arise where a client in one area has a significant traffic exchange with a server in another area. The volume of this traffic is such that a special virtual circuit or dedicated line would be warranted, but establishing such a link would violate the OSPF hierarchy. If routers on both ends were in different areas, the hello protocol handshake would fail and there would be no routing between them.

Again, this is a matter of if it hurts to do something, don't do it. Most routers have a mechanism for preferring one source of routing information over another. In this case, establish a static route between the client and server prefixes.

In the routers along this path between client and server, give the static route a preference factor that makes it more preferable than OSPF. While the details will be specific to individual router implementations, it is usually quite practical to establish OSPF routing, through area 0, as a

backup to the preferred static path.

## **10. Multiple Exit**

### Points/Multihoming

Use the power of OSPF's two types of externals when defining your policy for Internet connectivity. Many enterprise designers assume, incorrectly, they must run BGP to deal properly with connection to multiple Points of Presence (POP) of the same ISP, or to multiple ISPs when a primary/secondary policy is desired. See [[Multihome](#)] for a discussion of multihoming.

## **11. Security Considerations**

Another potentially confusing area is OSPF's use of authentication. Do not confuse OSPF authentication mechanisms with security mechanisms for user traffic. The purpose of the OSPF authentication mechanism is to protect the OSPF system itself from denial of service attacks.

Cleartext authentication is useful when combining different OSPF domains. By putting a password on the new or merged domain, you get some protection against configuration errors. An old router, not configured with the password, will not have its routing advertisements accepted by the new. Cleartext passwords are a simple method to confirm all routers in the routing domain have been examined by an administrator with the new design in mind.

Cryptographic authentication protects an OSPF domain against malicious attacks, typically launched from inside the OSPF domain. Such protection might, for example, be useful in a university network.

## **12. Acknowledgments**

## **13. References**

[RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

[RFC2328] Moy, J., "Open Shortest Path First version 2,"

[Multihome] Work in progress, H. Berkowitz, "To Be Multihomed: Requirements & Definitions," [draft-berkowitz-multirgmt-01.txt](#).

[Berkowitz 99a] H. Berkowitz. \_Designing Routing and Switching Architectures for Routing and Switching\_. Indianapolis: Macmillan Technical Publishing, (August 1999 publishing).

[Moy 1998] J. Moy. \_OSPF: Anatomy of an Internet Routing Protocol\_. Reading, MA: Addison-Wesley, 1998

[Thomas] T. Thomas. \_OSPF Network Design Solutions\_. Indianapolis: Cisco Press, 1998

#### **14. Author's Address**

C. Berkowitz  
Gett  
Communications  
PO Box 6897  
Arlington VA 22206  
Phone: +1 703 998 5819  
EMail:  
hcb@clark.net