SPLIT HORIZON DNS SERVER

This section was introduced at the suggestion of Maren Leizaola - many thanks for both taking the time to write and for providing interesting usage examples.

The term **Split Horizon** is normally used to describe a DNS server that will give different responses (IP addresses) based on the source address, or some other characteristic, of the query. While it has similar configuration properties to the Stealth Server it can also be used in a varity of unique situations such as:

1. **Geographic Mapping:** Assume that, for example, a web service is replicated in a number of locations (for either performance or access latency reasons) then a specific IP address may be returned based on the source address of the query to ensure the shortest possible path from the user to the service. For those familiar with **anycast** you could consider this as a poor man's anycast service.

2. **Naming Consistency:** Assume that you have, say, a corporate in-house LDAP service and that you want to keep certain highly secure data on one server only accessible to certain individuals or organizational sections,

which have unique or identifiable IP addresses or address ranges, but for reasons of consistency (scripts, configuration files etc) you want both the secure and insecure LDAP services to be named, say, ldap.example.com.

3. **Load Balancing:** Assume that an analysis of incoming service users shows that their source-ip addresses can be separated into contiguous ranges: 50% from a to b, 50% from b to c. In this case rather than simply provide multiple A/AAAA RRs (where load balancing is essentially random) it may be more effective to use a split-horizon strategy.

Other possibilities may strike imaginative readers. The unifying element is that some characteristic of the incoming query will cause the DNS to generate a query-dependent result.

BIND's view clause provides a method that can be used to build such configurations and example files are provided .