Locator ID Separation Protocol (LISP)

dino@cisco.com
oran@cisco.com
vaf@cisco.com
dmm@cisco.com

Saturday, March 17, 2007

Background

- Stimulated from problem statement effort at the Amsterdam IAB Routing Workshop on October 18/19 2006
- The solution started at dinner between Dino Farinacci, Dave Oran, and Jason Schiller on day-1 of workshop
- Discussions continued with various people on day-2 of workshop, primarily with Lixia Zhang and Vince Fuller

Agenda

- Problem Statement for LISP
- What is LISP
- Packet Flow Example
- Deployment Scenarios
- Prototype Schedule
- Sneak Preview of -01 Draft

Problem Statement

- We need a mechanism to:
  - Associate an ID with a set of Locator addresses
  - Forward packets using Locator addresses
  - Maintain the reachability status of Locator addresses
- The mechanism needs to be:
  - Simple so it can be easily and incrementally deployed
  - Does not depend on a lot of Internet infrastructure
  - Does not require transit non-TE routers to carry state
  - No specialized ID/Locator binding service
  - Flexible so both Sites and Providers can benefit
  - Pragmatic so it can be deployed in <= 12 months

Site-Based Requirements/Goals

- Sites need to be multihomed
  - Connected to more than one provider
- Sites need flexibility to change providers
  - While maintaining session survivability
- Site-supported devices need to be mobile & roam
  - While maintaining session survivability
- Sites need to easily renumber their devices
  - While maintaining session survivability

Provider-Based Requirements/Goals

- Providers need their routers to scale or they can't deliver any service
- Providers need to maximize their resources to deliver cost effective connectivity
  - Providers want the ability to do Traffic Engineering
- Provider-supported devices need to be mobile & roam
  - While maintaining session survivability
  - While achieving scalability

What is LISP

- A Map-n-Encap Scheme
  - Formal definition for separating an ID and a Locator
- Procedures for tunneling where:
  - EIDs are in inner headers
  - Locators are in outer headers
- Procedures for obtaining EID-to-Locator mappings
- Procedures for determining Locator Reachability
- Formal definition for tunnel router placement and TE usage

Where to get mappings

- 4 variants depending on your tradeoff:
  - LISP 1
    - Routable IDs over existing topology to probe for mapping reply
  - LISP 1.5
    - Routable IDs over another topology to probe for mapping reply
  - LISP 2
    - EIDs are not routable and mappings are in DNS
  - LISP 3
    - EIDs are not routable, mappings obtained using new mechanisms (DHTs perhaps)
- `draft-farinacci-lisp-00.txt` documents LISP 1 and 1.5 using ICMP

LISP Terminology

- EIDs
  - Endpoint IDs, obtained from DNS as we do today
  - Third-party referral behavior same as today
- RLOCs
  - Routing Locators, IP addresses of routers in a destination site
- Ingress Tunnel Router (ITR)
  - Encapsulates first set of Locators in source site
  - Second set of Locators optionally by TE ITRs
- Egress Tunnel Router (ETR)
  - Decapsulates in destination site
  - Decapsulates optionally by TE ETRs

LISP Terminology

- The LISP "Cache" is:
  - The EID->RLOC(s) mappings
  - The cache is built on demand - learned from ICMP (in LISP 1.x), or ??? in the case of LISP ≥ 2
  - Caches have the information to get you somewhere
- The LISP "Database" is:
  - The configured IP addresses of routers which are used as Locator addresses for hosts that have IDs assigned from subnets attach to the routers
  - Advertised in ICMP messages
  - Databases have the information for others to get to you

LISP Terminology

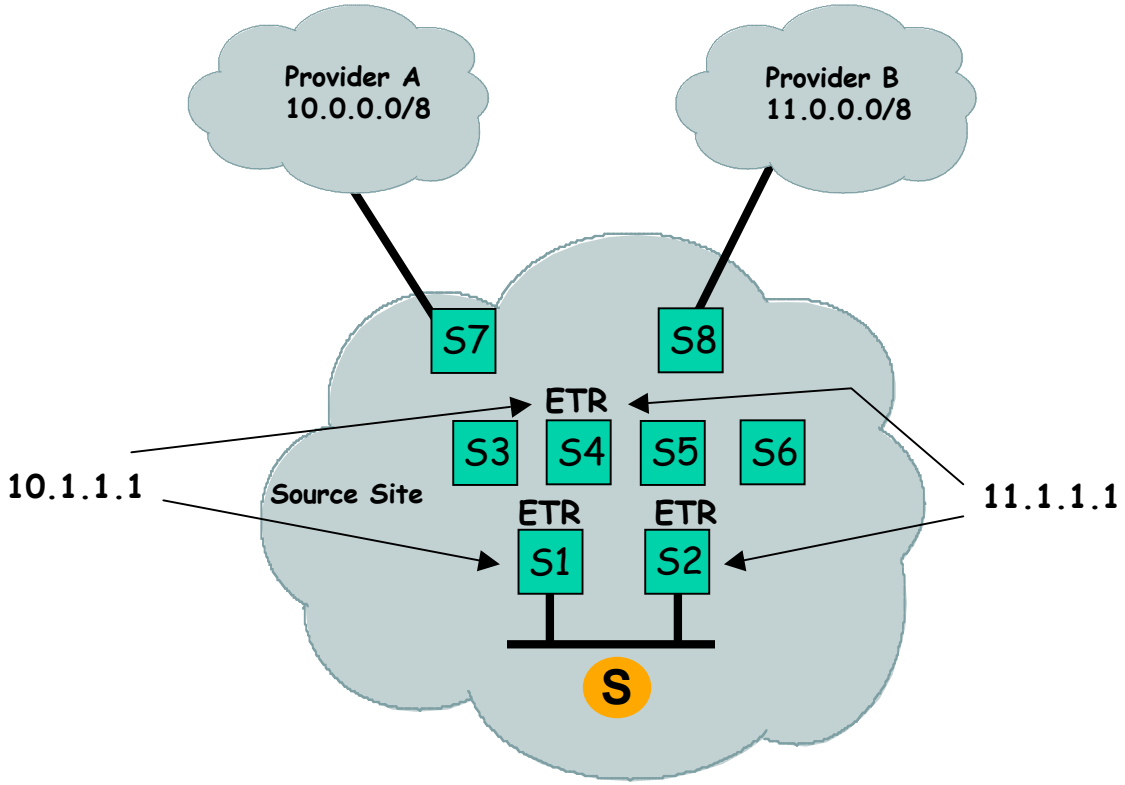

The "LISP Database" is already configured in the site network

Packet Flow Example

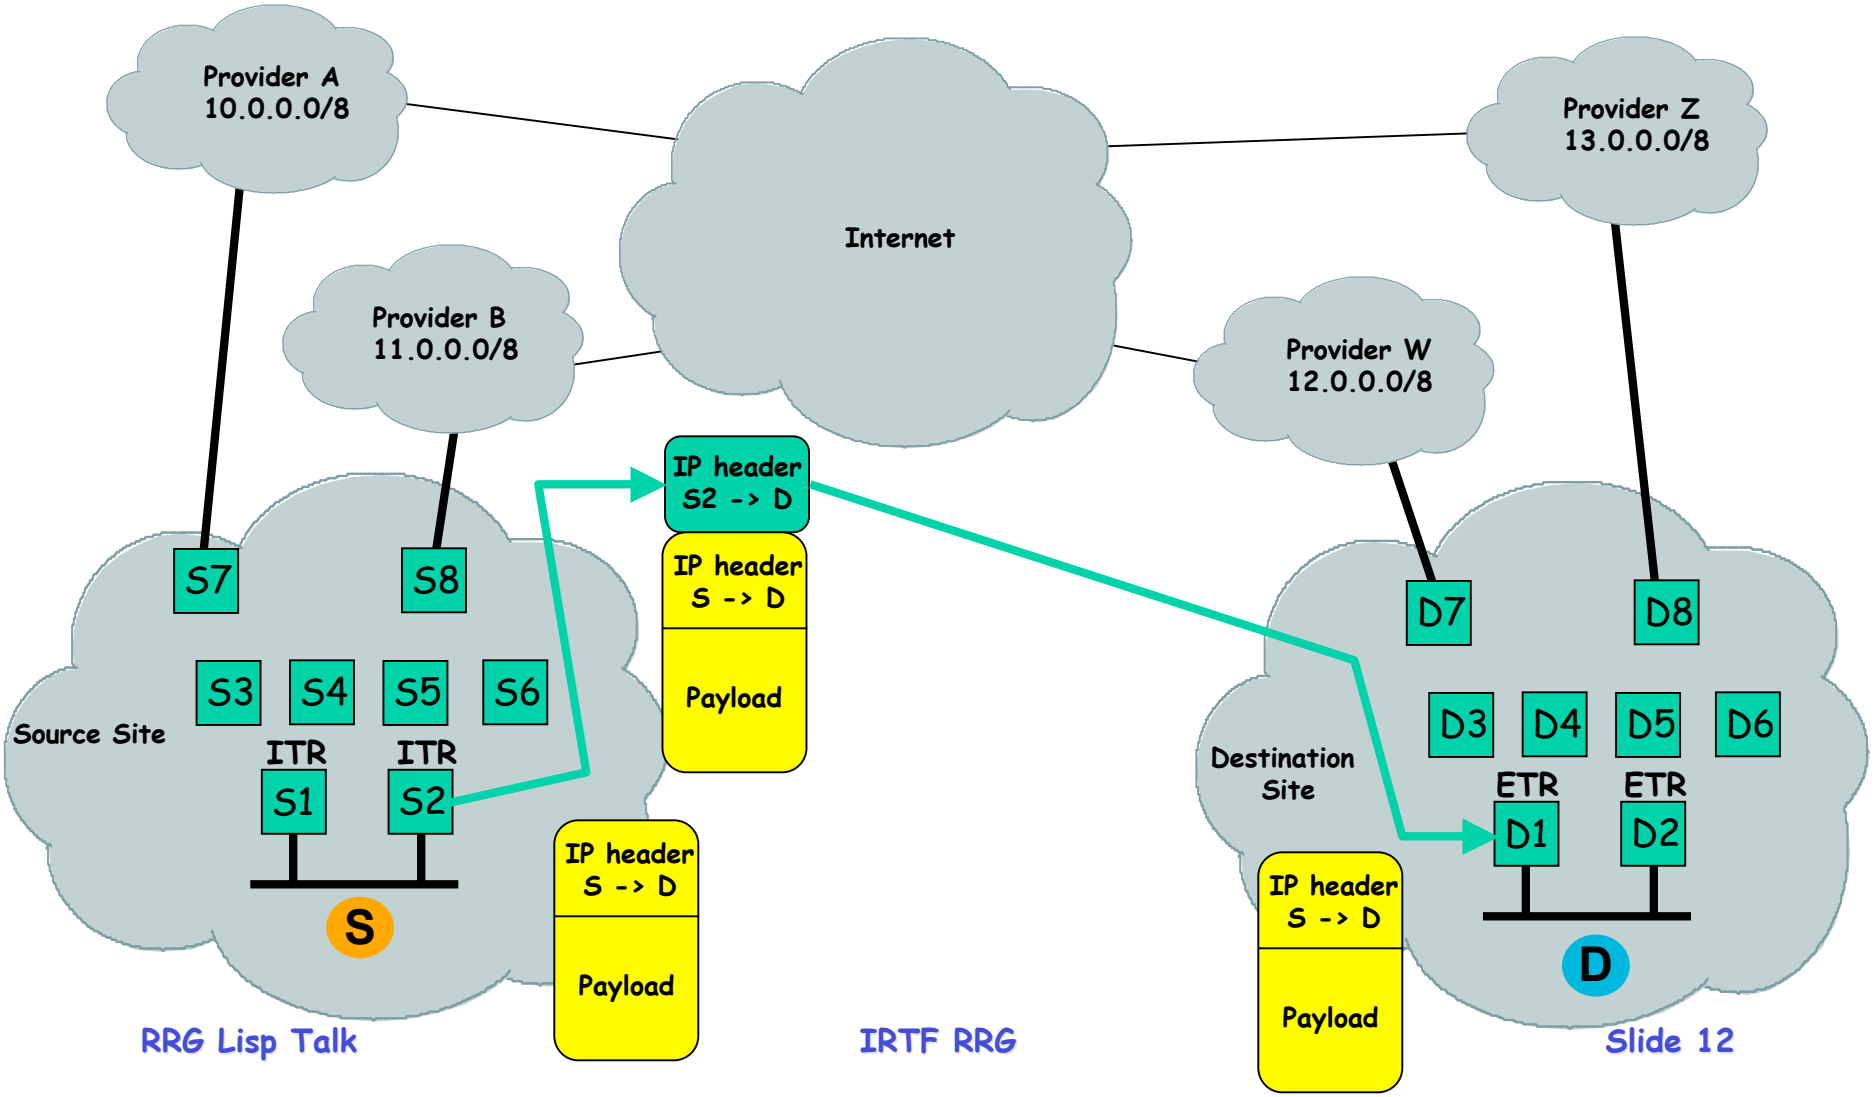

Packet Flow Example

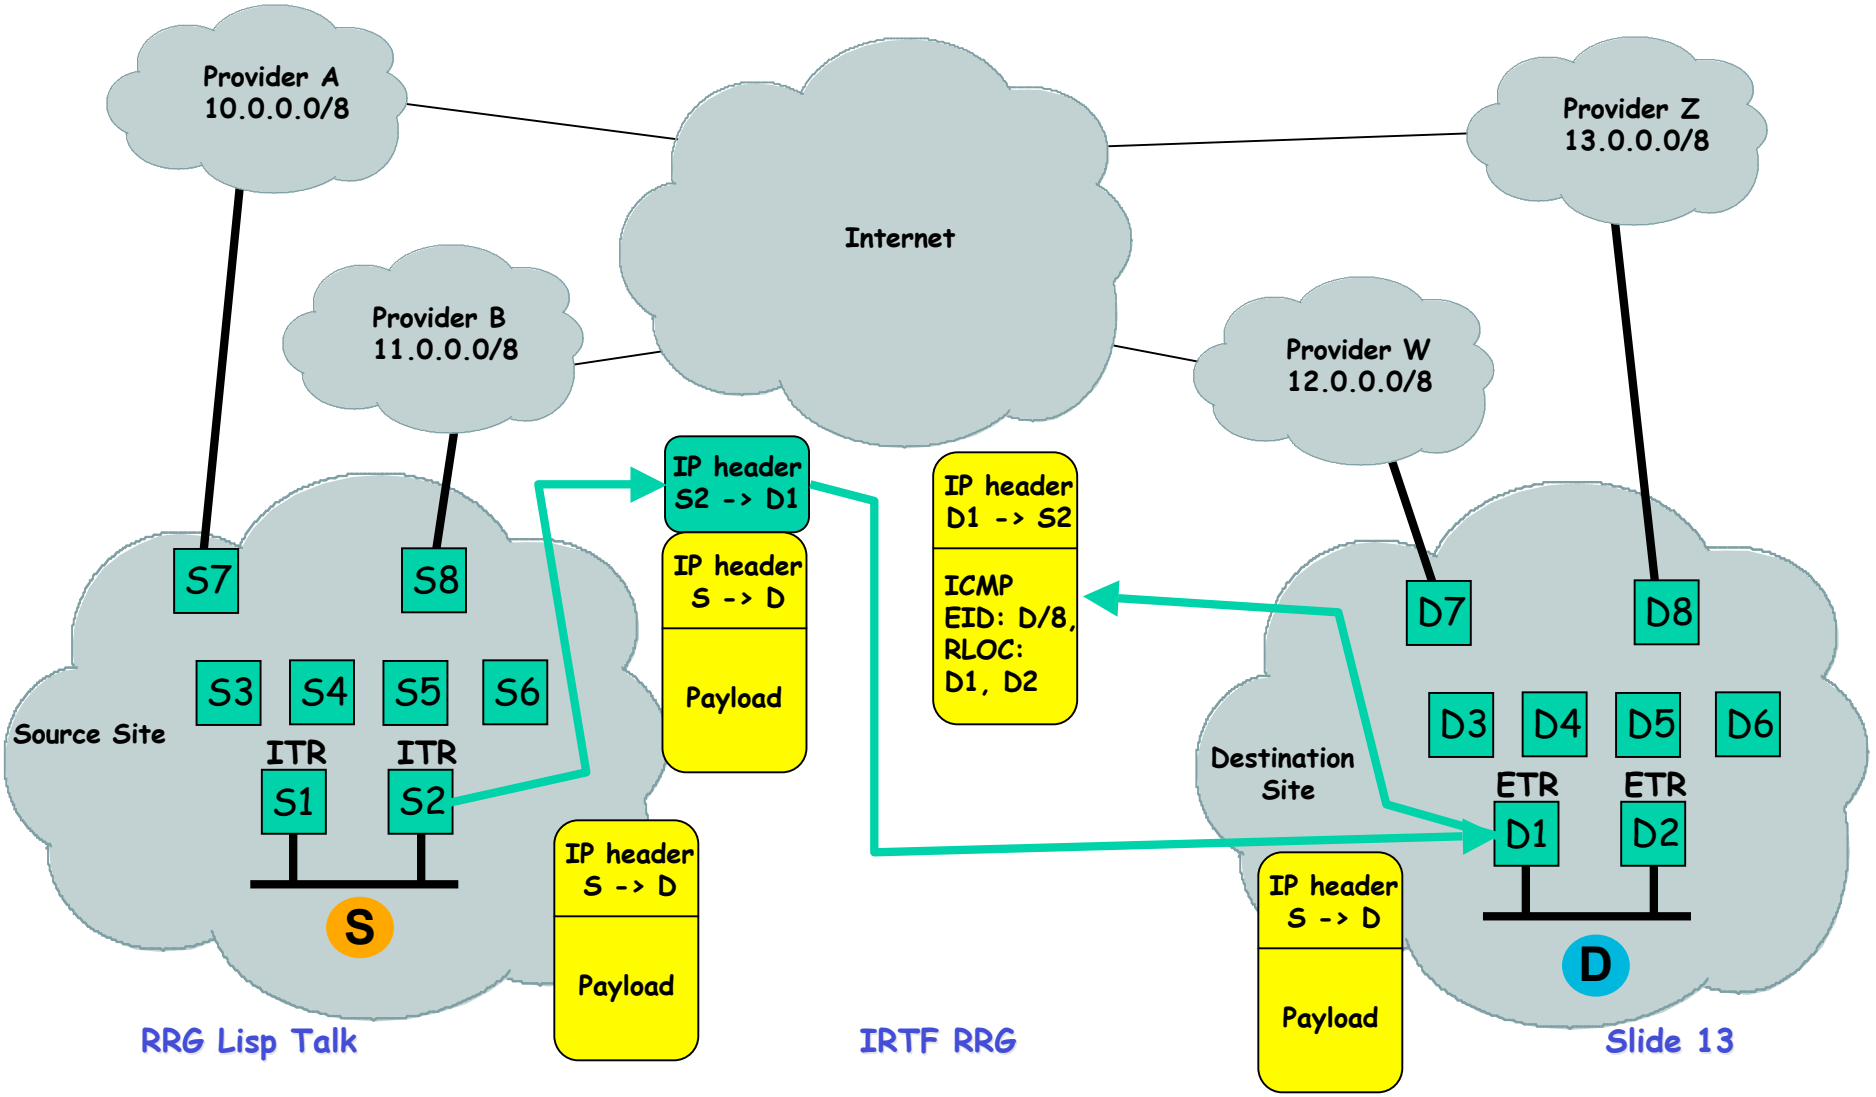

Deployment Scenarios

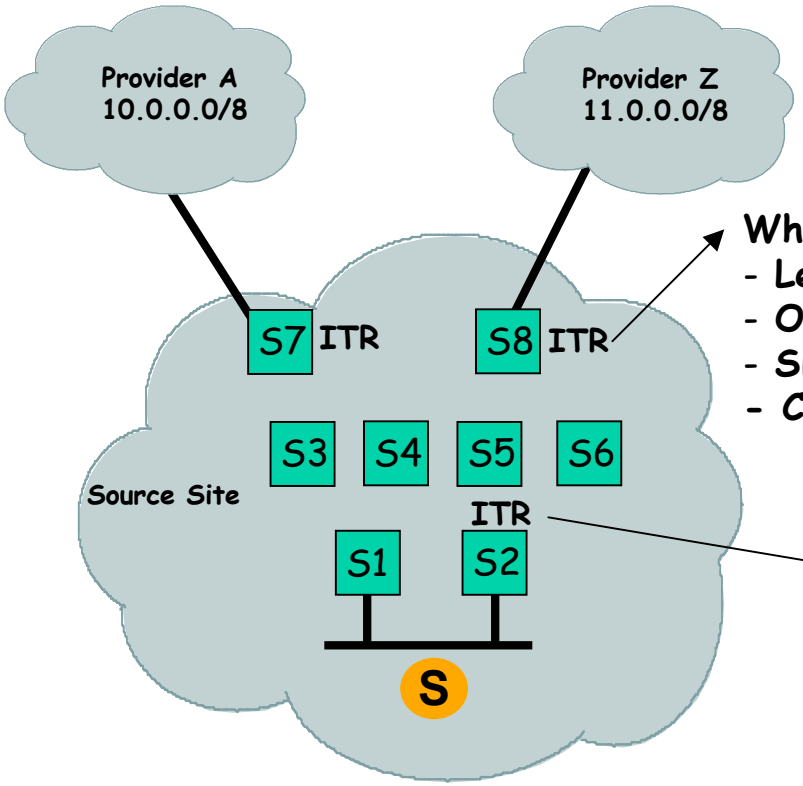

When you want LISP caches centralized for the site:
- Less granular EID-prefixes
- One or two Locators for entire EID-prefix
- Single points of failure
- Considerably less control messaging

When you want the LISP distrubuted:
- More granular EID-prefixes
- More locators
- Spread across cheaper memory
- Path diversity
- More control messaging

Deployment Scenarios

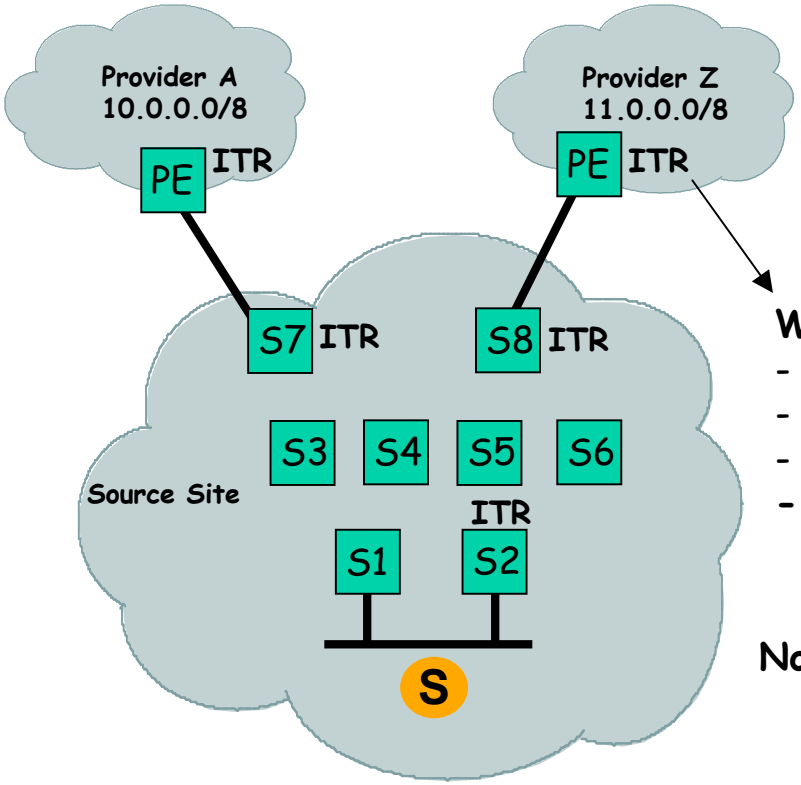

When the ISP wants control with no site participation:
- Less granular EID-prefixes
- One or two Locators for entire EID-prefix
- Single points of failure
- Considerably less control messaging

Note: ITRs and ETRs can be placed independently

Deployment Scenarios

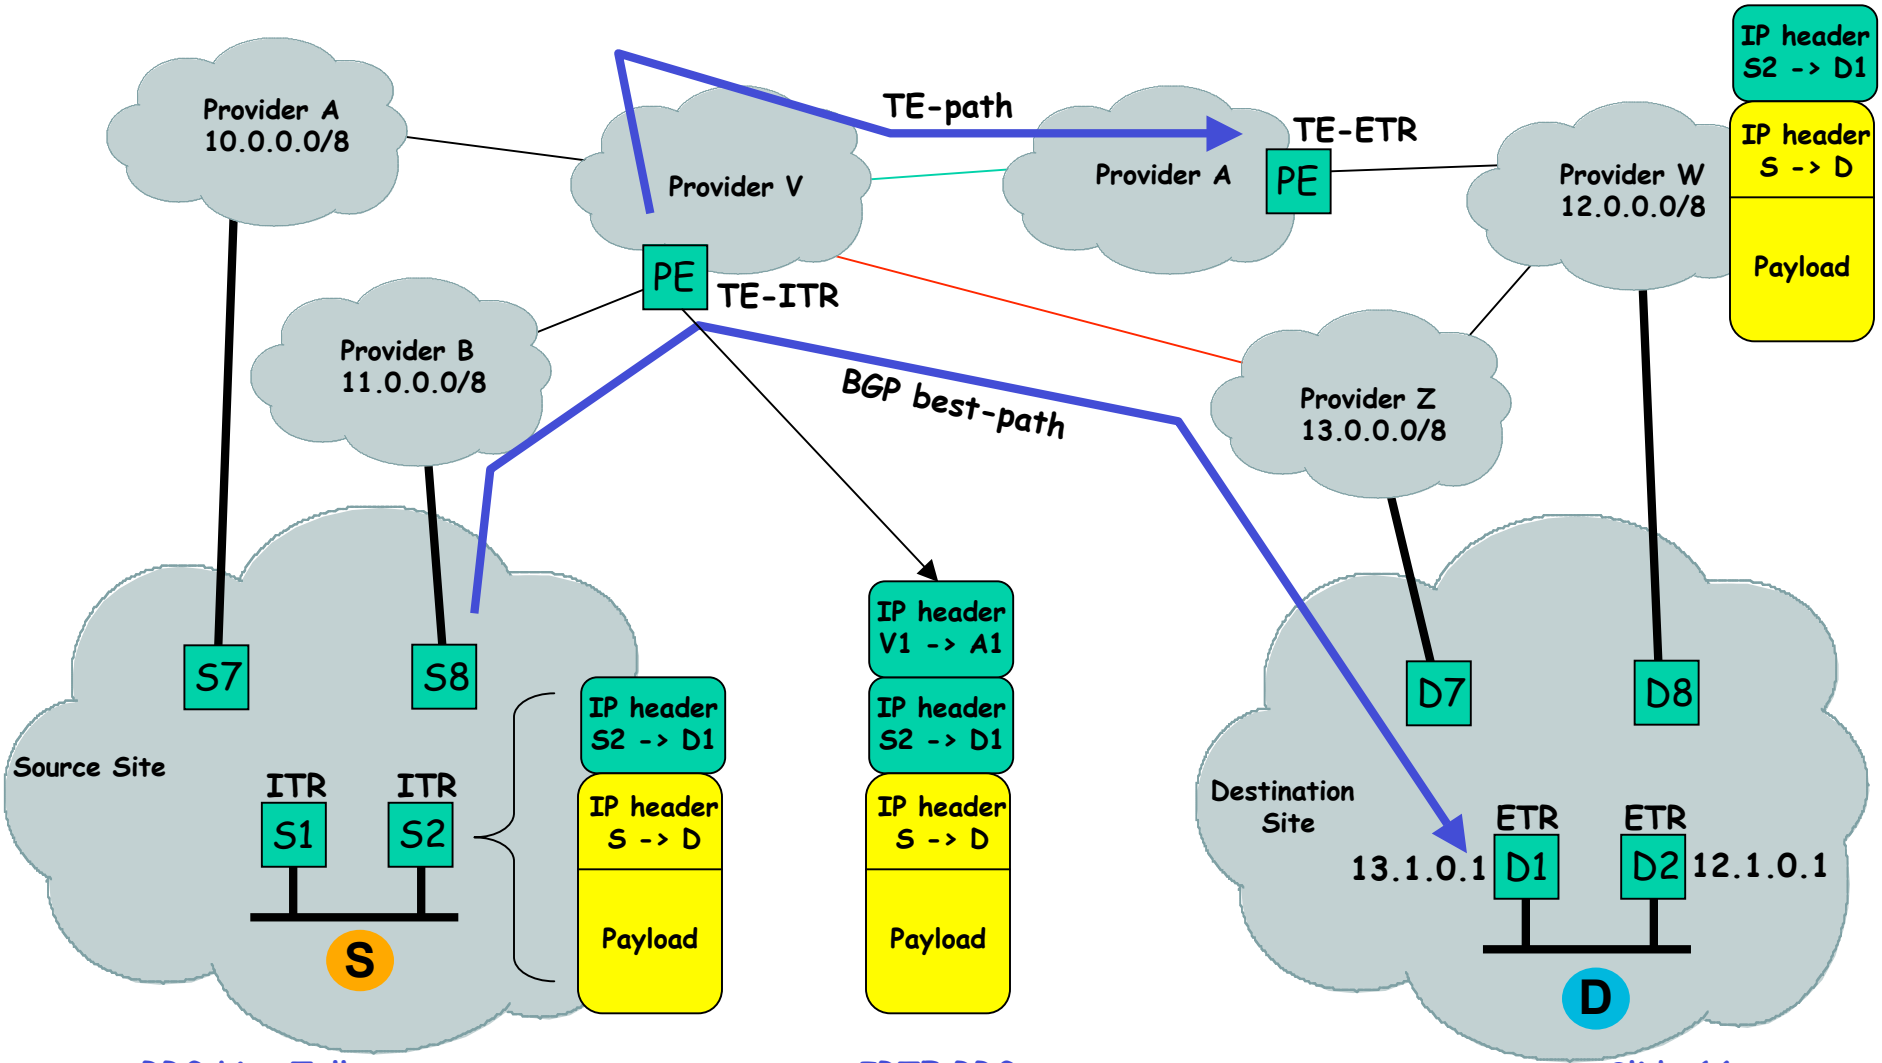

Locator Reachability

- ICMP Unreachable Messages
  - Will tell you about Locator unreachability
- ICMP EID-to-RLOC Replies
  - Tell you only about mappings
  - Reachability is assumed
- ICMP EID-to-RLOC Requests
  - Can test for reachability
- Return packets confirm reachability

Locator Selection

- Need both ITR and ETR side control
  - Use of priorities and weights
  - Priority enables a locator from an ETR side
  - Weights indicate how traffic is balanced across enabled Locators, when 0, ITR can decide
- Large content providers said they need to glean
  - Simply swap Locators for returning packets
  - Don't want to store large caches of clients (even if aggregated into EID-prefixes)

Prototype Schedule

- Prototype a software forwarding version
  - Dino the code maven
- Recruit multiple vendors for prototype interoperability testing
- Provide eval platforms for lab testing
  - Vince/Dave M tests internally at cisco
  - Jason/Chris, Dorian/Peter, Ted/Peter test externally
  - That would be UUnet/VB, NTT/Verio, and Sprint
  - Lixia and Geoff could test in research labs
- Report on prototype status at Suimmer IETF
- In parallel, determine hardware requirements

Sneak Preview of -01 Draft

- Add anti-spoofing support and EID-to-RLOC hijacking avoidance
  - Nonces to protect against response spoofing
  - Use public/private keys
  - Not relying on PKI
- Use AH for ICMP messages
  - Gets through firewalls
- Changes from prototype experience

Ready…Fire…Aim

Ducking for cover…