Problem Statement and Architecture for Information Exchange Between Interconnected Traffic Engineered Networks

draft-farrel-interconnected-te-info-exchange-00

A.Farrel, J.Drake, N.Bitar, G.Swallow, D.Ceccarelli

CCAMP WG, IETF 86th Orlando

DRAFT GOAL

- "This document sets out the problem statement and architecture for the exchange of TE information between interconnected TE networks in support of end-to-end TE path establishment"
- Inter-domain (area / AS) signaling exists but corresponding routing problem has never been addressed
- Need to provide an abstract representation of a TE domain to other TE domains (TE Reachability)
- Provide mechanisms that allow CSPF in one domain to compute a loose ERO to an egress node in another domain
- Yes, includes Overlay Model Framework

Reachability vs TE-Reachability

- Reachability

In an IP network, reachability is the ability to deliver a packet to a specific address or prefix. That is, the existence of an IP path to that address or prefix.

- TE Reachability

TE reachability is the ability to reach a specific address along a TE path

  - Unqualified TE reachability: helpful in determining a path to a destination that lies in an unknown domain
  - Qualified TE reachability (by TE attributes): TE metrics, hop count, available bandwidth, delay, shared risk.

USE CASES

- "For the purposes of this document, a **domain** is considered to be any collection of network elements within a common sphere of address management or path computational responsibility. Examples of such domains include IGP areas and Autonomous Systems."
- 1. Peer Networks. E.g:
  - Two interconnected TE domains w/ multiple attachment points
  - Mesh of interconnected TE domains w/ multiple attachment points
- 2. Client-Server (Overlay) Networks:
  - Same addressing space vs VPNs in Overlay Context
  - Multiple server layer domains
  - Dual Homing

Peer Networks

- Two interconnected TE domains w/ multiple attachment points

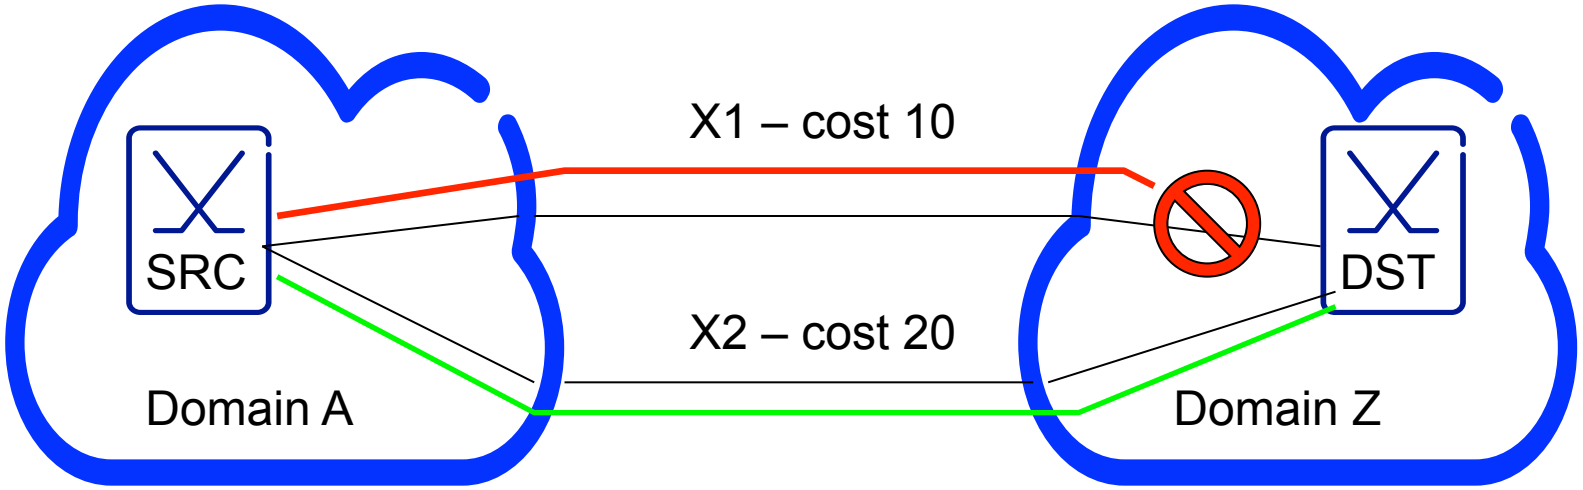

Peer Networks

- Mesh of interconnected TE domains w/ multiple attachment points

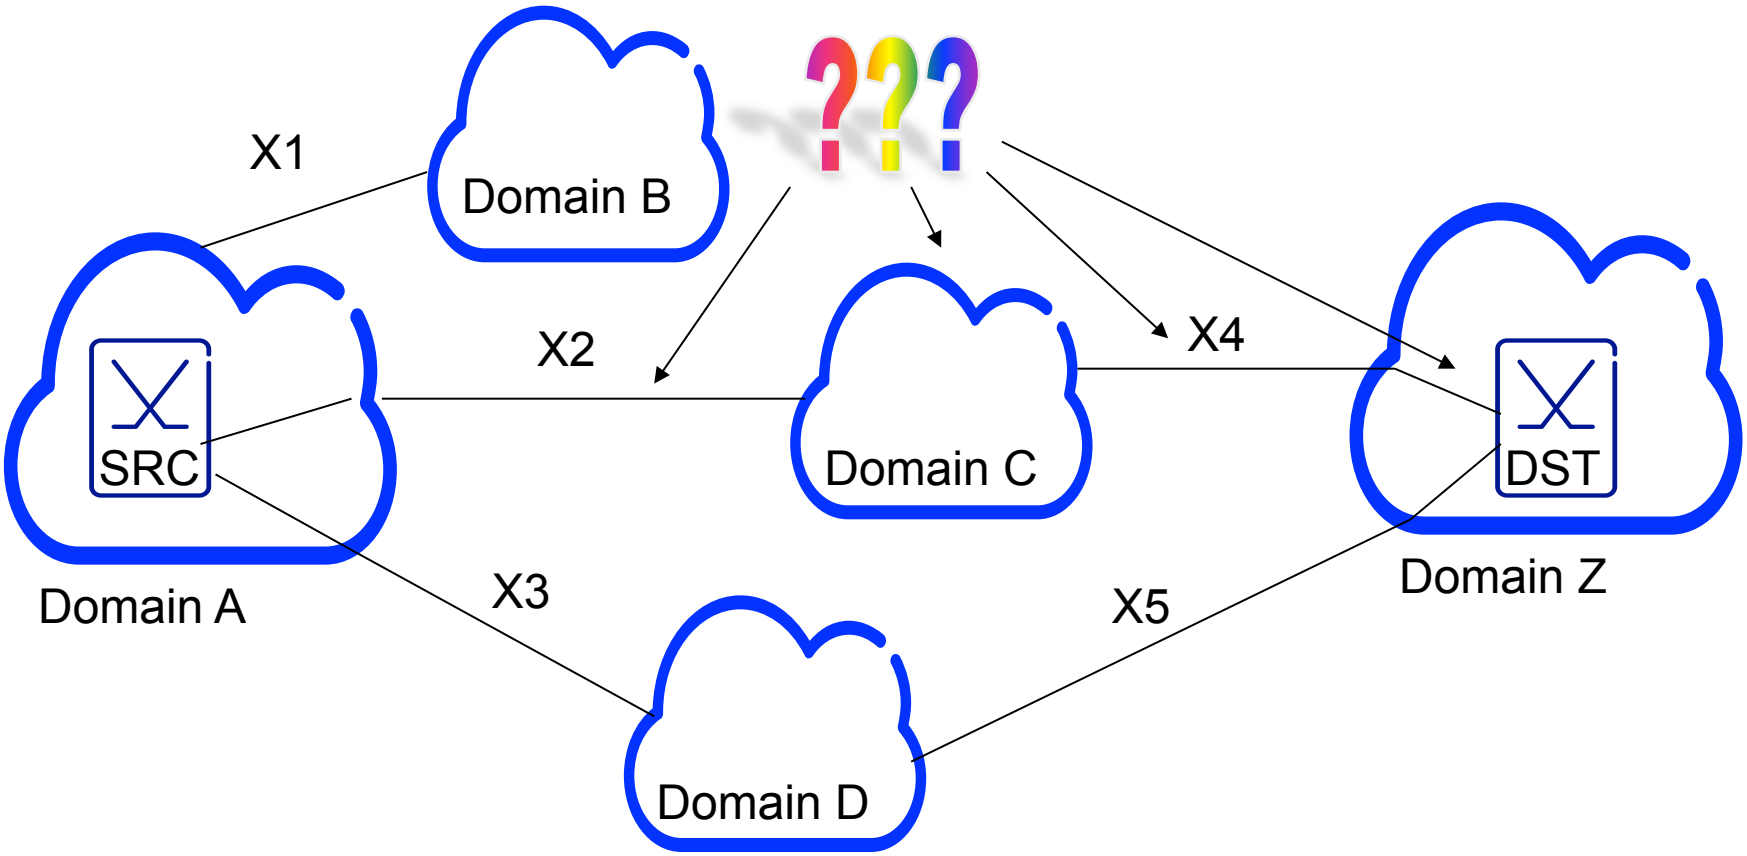

Overlay Networks

- Domains belonging to one network are connected by a domain belonging to another network (same address space)
  - Once connections are performed across the lower layer network, the domains of the upper layer network can be merged into a single domain by running IGP adjacencies over the tunnels.

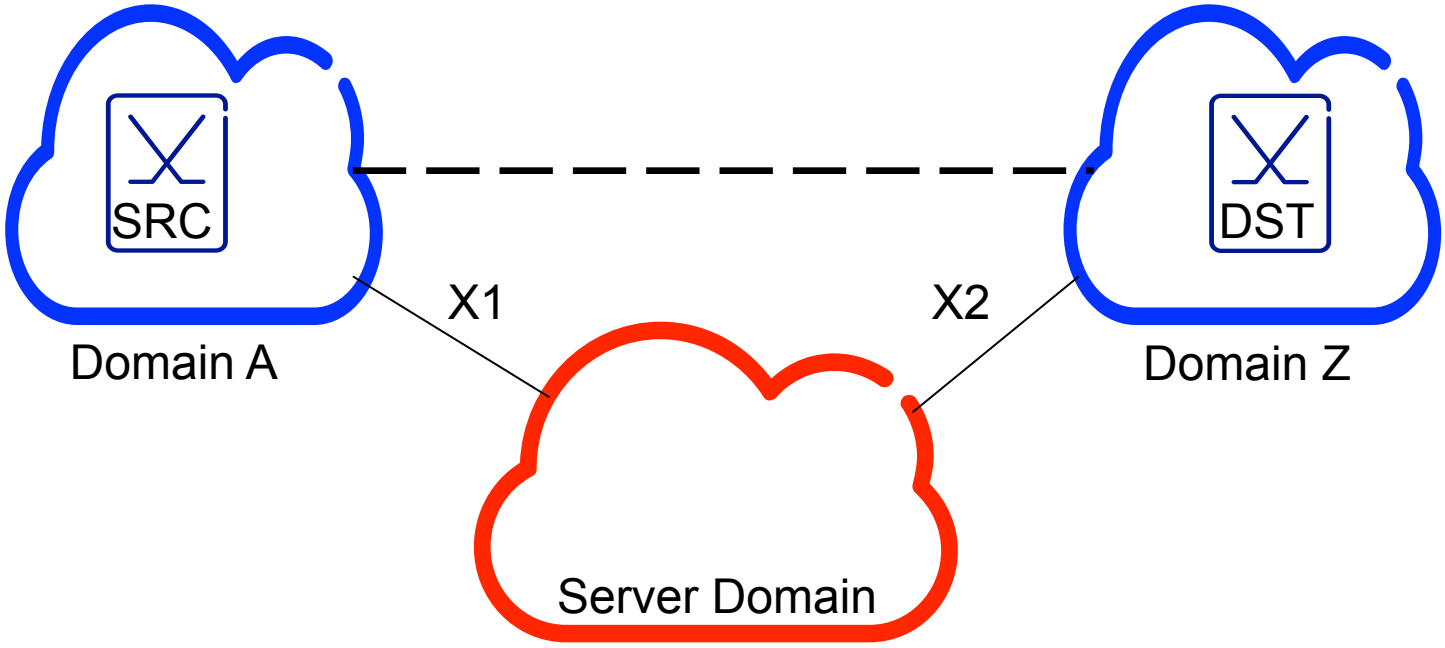

Overlay Networks for VPN

- Client network has a different address space that of the server layer (non-overlapping not guaranteed)
  - VPN sites comprise a set of domains interconnected over a core domain.

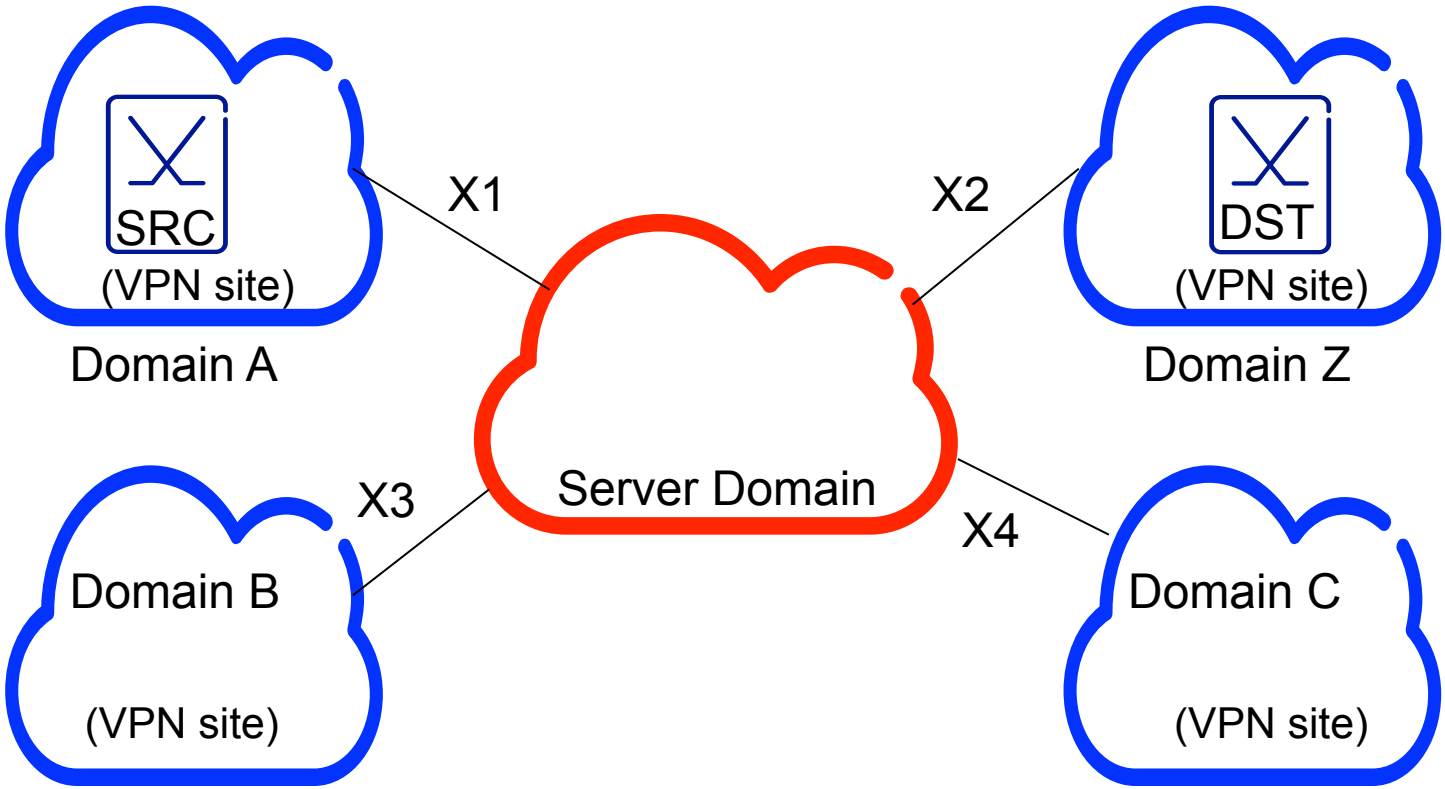

Overlay Networks-Multiple servers domains

- Connectivity between higher layer domains is provided by a sequence or mesh of lower layer domains.

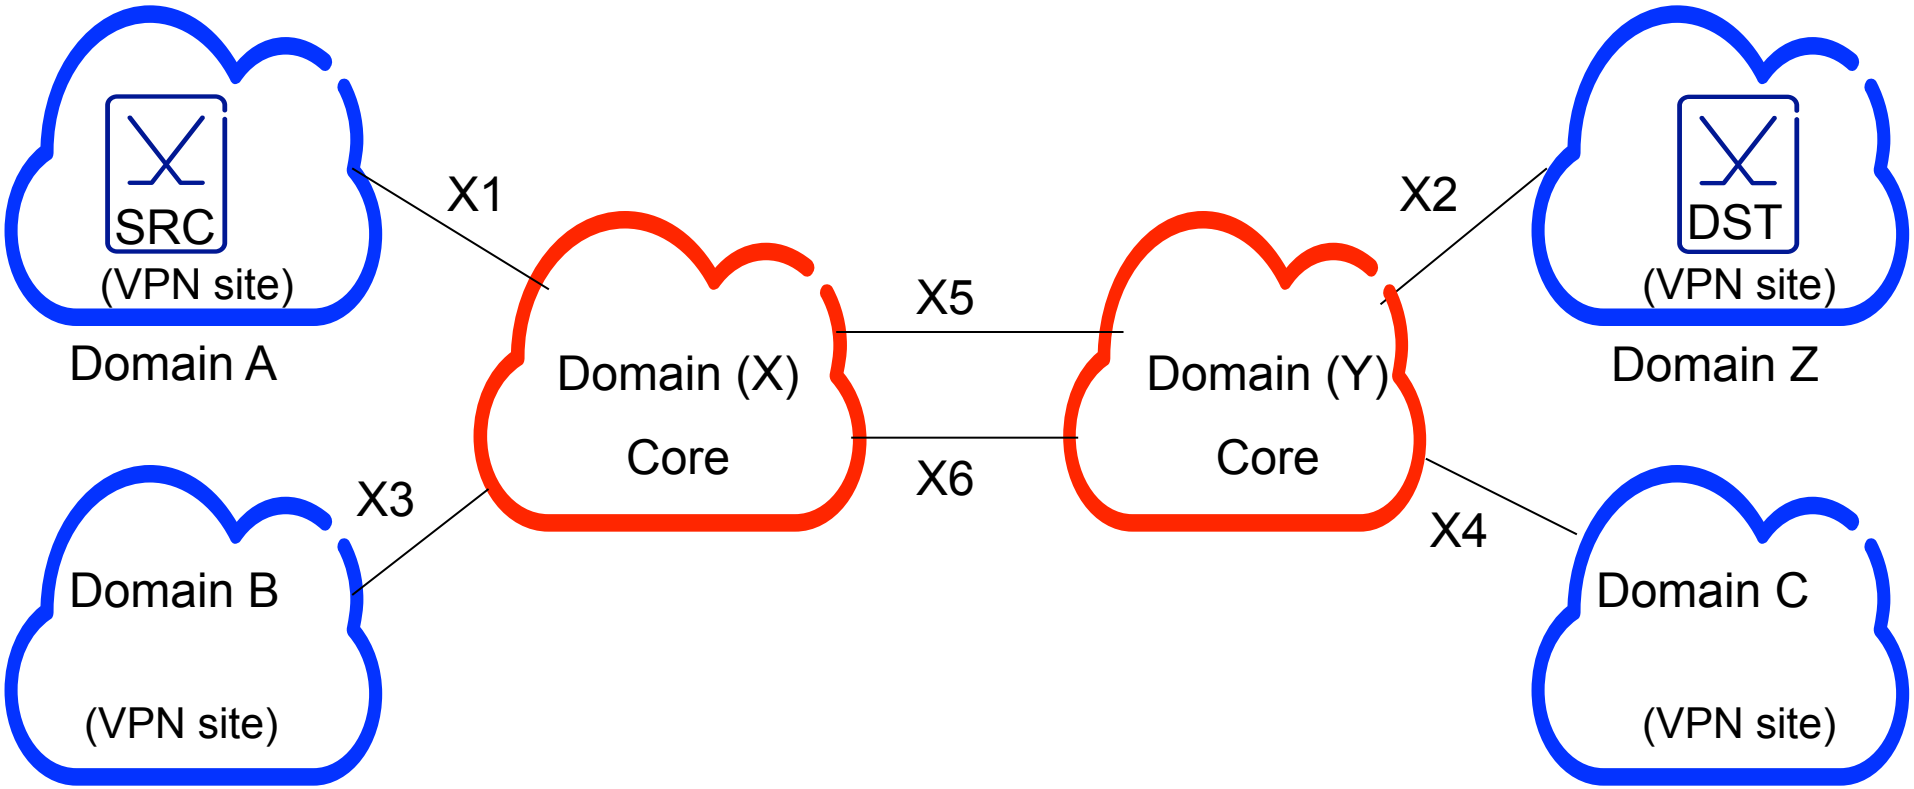

Dual-Homing

- Further complication added to the client-server relationship: client domain attached to more than one server domain

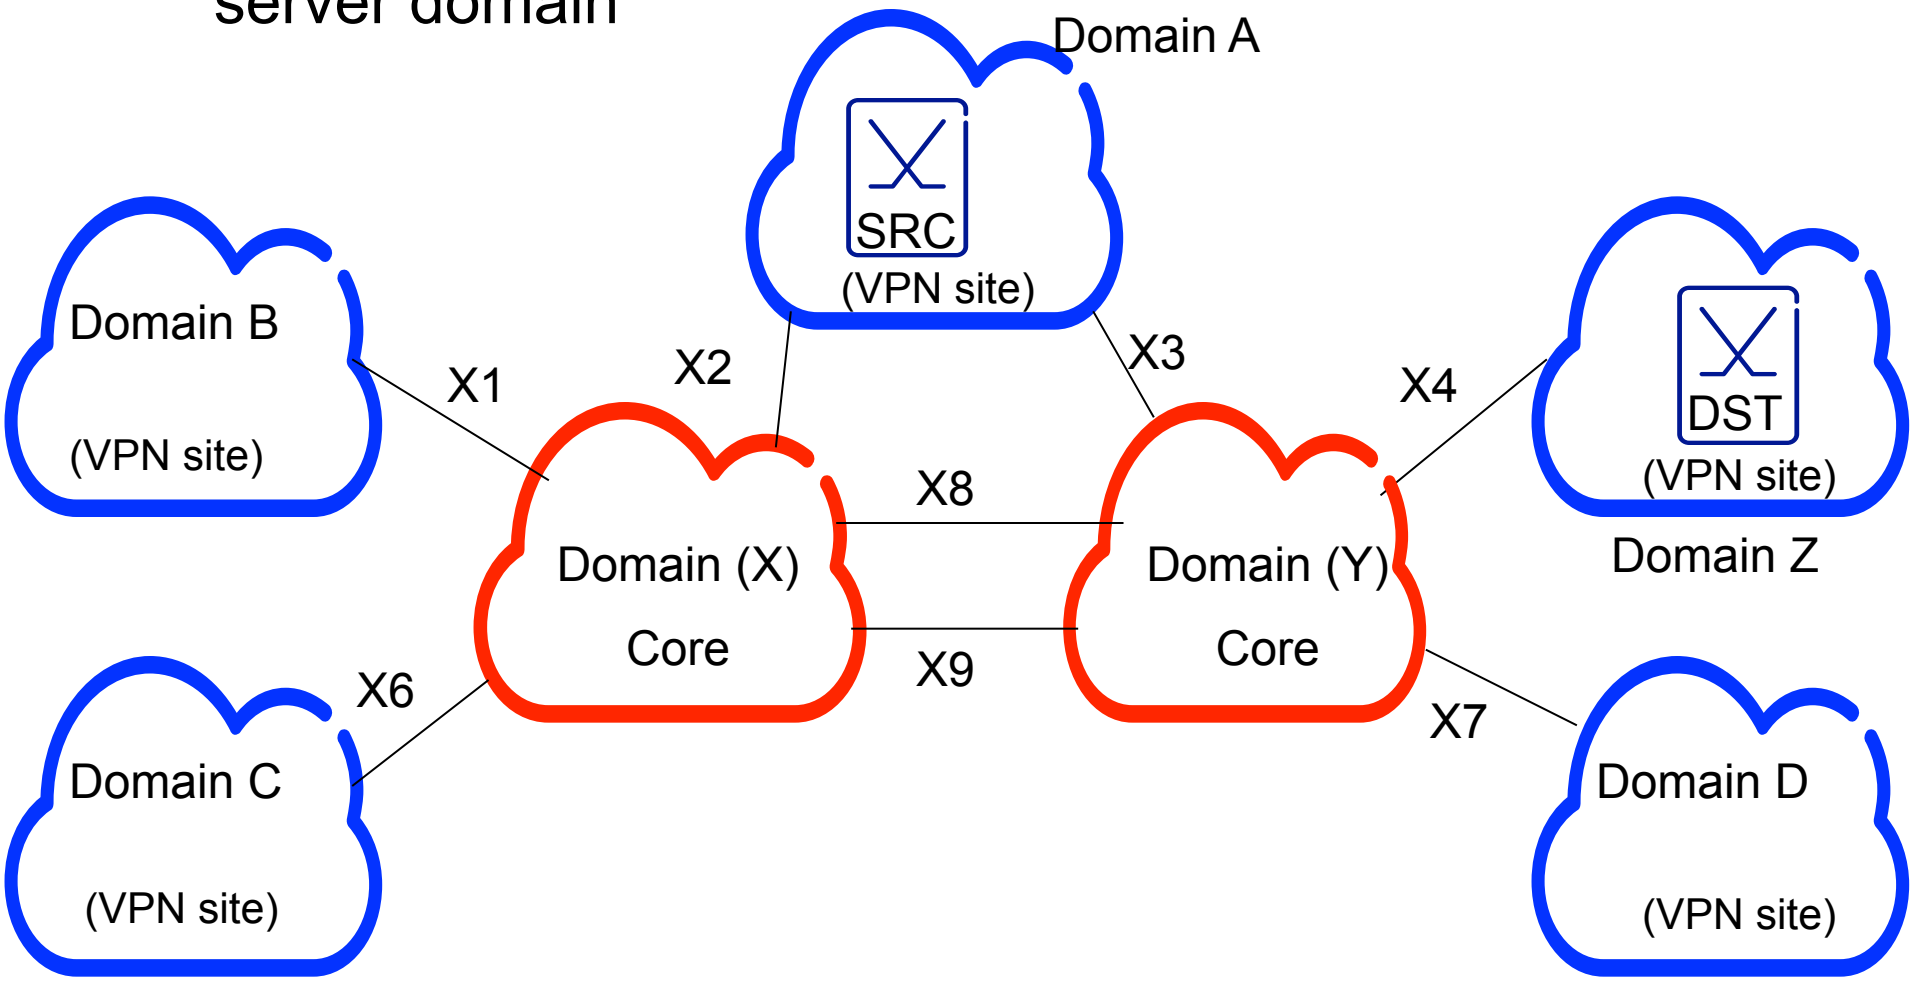

Issues

- Problem statement: A mechanism is required that allows path computation in one domain to make informed choices about the exit point from the domain when signaling an end-to-end TE path that will extend across multiple domains.
- Policy and filters
  - Control what information is exported by a TE domain and information that is imported by a TE domain
  - Confidentiality: keep details of a TE domain safe from prying eyes
- Information overload and churn
  - Control the volume of TE information distributed within an inter-connected set of TE domains (e.g. periodic intervals or significant change in resources)
  - Virtual Network topology (node and links): "is made up of links in a network layer. Those links may be realized as direct data links or as multi-hop connections (LSPs) in a lower network layer. Those underlying LSPs may be established in advance or created on demand."

Aggregation models

- Virtual Node
  - TE domain or subsets of it are advertised as a single node
  - Blocking fabric difficult to represent efficiently, subject to churn as LSPs are established and released
  - Basically devolves in representing virtual links (see below) internal to the virtual node
- Virtual links
  - A set of nodes and links in a TE domain are advertised as a single edge-to-edge link
- TE domain need to be advertised as a combination of real/virtual nodes and links
  - E.g., Real nodes at TE domain edges, virtual links across the TE domain, real TE LSP endpoints

Next Steps

- Architectural concepts
  - Basic Compoenents: Peer interconnection and Overlay interconnection
  - Abstraction not aggregation:
    - Abstract Links and abstract Nodes
    - Abstraction In Peer Network and in Overlay Networks
    - Considerations for Dynamic Abstraction
    - Requirements for advertising abstracted links and nodes
- Building on existing protocols
  - BGP, IGP, RSVP-TE