Network Working Group
Internet Draft
Expires: April 2001

R. Callon
Juniper Networks
B. Gleeson
Nortel Networks
Eric Rosen
Cisco Systems
Chandru Sargor
Jieyun (Jessica) Yu
Cosine Communications

Outline for A Framework for Network Based Virtual Private Networks

<draft-callon-nbvpn-outline-00.txt>

Status of this Memo

Abstract

The work of the IETF NB-VPN working group will be aided by a working group framework document. Input for such a document needs to come from multiple sources, including companies involved in the major proposals being developed by the working group.

We are intending to produce a framework document based on this outline. The resulting document will discuss technical issues for Network Based Virtual Private Networks (NB-VPNs). It is the intent to produce a coherent description of the significant technical issues which are important in the design of network based VPN solutions. Selection of specific approaches, making choices regarding

engineering tradeoffs, and detailed protocol specification, are outside of the scope of a framework document.

1. Introduction

NOTE: This is a rough draft for an outline for a working group VPN framework document. It is expected that this outline will be updated during the process of completing the framework document. However, we also expect that agreement will be easier if we first agree on the general format and most of the content for the outline, and then undertake to fill in specific sections.

The text included in this outline is intended for the purpose of giving a general idea regarding what subjects may be discussed in each section. It is expected that the text included herein is likely to be re-written and/or taken from other documents. Note that in many cases the text in this outline will consist of only brief bullet items, which will list the general topics which are likely to be discussed in each section.

With the exception of brief introductory material, the scope of this outline is limited to Network Based Layer 3 VPNs. This implies VPNs for which the provider network participates in layer 3 forwarding, and in routing and management of the VPN, as defined below. CPE-based VPNs are outside the scope this document. This scope is selected to match the anticipated scope of the IETF NB-VPN working group. If the scope of the working group is wider than anticipated, then the scope of this outline may be extended accordingly.

This document does not make choices, but rather describes issues, technology, and the possible solutions to each problem. We will therefore describe multiple possible solutions which may be used.

This document does not describe any specific complete solution. Note that any specific solution will need to make choices, and will need to make tradeoffs between flexibility and conciseness. Also note that the requirements for VPNs will vary between different applications, and therefore there may be a need for multiple different solutions for use in different situations.

1.1 Overview of Virtual Private Networks

The intention of this and the following sections is to introduce the main characteristics of VPNs, and to specify what is in/out of scope for this document (which is intended to match the scope of the WG). The intention is to only briefly discuss each of the items listed below, and refer to later sections where appropriate.

VPNs

- private networks
- interconnected over a public infrastructure (note: To some extent this is true of frame relay, ATM, and even ADM networks; Thus it is not clear how "unique" this definition really is.).
- show picture. Define "provider edge (PE)" and "customer premise edge (CPE)".
- Private addresses in the private network implies encapsulation/tunneling (of some sort, unless there are separate physical links)
- Need for security
- Need for QoS
- Isolation between VPNs - separate routing/forwarding tables per vpn (at the PE gear, the input port implicitly creates restrictions regarding where the packet can be delivered).

1.2 Types of VPNs

- Many types. It is not up to this document to decide between different types of VPNs. There are tradeoffs.

1.2.1 CPE-Based vs Network-Based VPNs

- CPE-Based VPNs characterized by tunnels between CPE gear. Tunnels could either be link layer connections (e.g. ATM, FR) or IP in IP (IP/IP, GRE, IPSEC). Provider network unaware of the operation of IP (or other network layer protocol) in the private network, just sees ATM/FR frames or IP packets.

- Network Based VPNs - provider network participates in reachability distribution and tunnel establishment (optionally also other functions). Does not preclude CPE to CPE tunnels, but these are set up with the involvement of the provider network.

- Note that these definitions actually could overlap (if the provider manages CPE gear). A clearer distinction is whether a VPN is layer 3, as discussed in section 1.2.4 below.

1.2.2 Types of Network Based VPNs

Different types of network based VPNs may be distinguished by the service offered.

- Multi-site Layer 3 service - provider forwards packets based on layer 3 information (in scope)
- Multi-site Layer 2 service - (transparent LAN service) - provider forwards packets based on MAC address (out of scope)

- Link Concatenation (VLL) - provider forwards packets on the basis on the incoming link on which the packet was received (out of scope)

1.2.3 Network Based Layer 2 VPNs

- Network is aware of VPN, but does only level 2 forwarding. Options include forwarding based on MAC addresses, use of pt-to-pt link layer connections, multipoint-to-pt (e.g merged MPLS LSPs), pt-to-multipoint (e.g. ATM VCCs) etc.

1.2.4 Network Based Layer 3 VPNs

- What this is: Network layer forwarding in the carrier (specifically in PE gear). Network is aware of the VPN (for example, it forwards L3 packets for the private network, and may participate in routing, configuration / discovery)

- How this differs from CPE based VPNs and network based layer 2 VPNs

- Given that PE gear needs to forward packets directly from the private network, using the private network's address space, this implies that PE gear needs to participate in some manner in routing for as many private networks as the PE gear supports (refer to later sections). Basically this moves some functions from the private network to the public network.

- "Network Connectivity Service" versus "Full Network Service". Former provides constrained connectivity at layer 3. Latter may include other network services such as firewalls, user authentication and address assignment (e.g. RADIUS, DHCP) etc.

- Note: As an example, a layer 3 VPN may support constrained connectivity, so that a single site may be in multiple VPNs, so that to go from site 1 to site 3 you might have to traverse site 2. This would for example make sense where the firewall is in site 2. We don't intend to talk about how the firewall works, but will talk about how a VPN could support constrained connectivity. Similarly there could be NAT boxes between overlapping private address spaces in different private networks, which would most likely provider constraints similar to the firewalls, in that routing between two sites may need to traverse a third site.

1.3 Terminology

1.4 Acronyms

BGP	Border Gateway Protocol

CE	Customer Edge
CoS	Class of Service
CPE	Customer Premise Equipment CR-LDP ??
FEC	Forwarding Equivalence Class
GRE	Generic R?? Encapsulation
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol (eg, RIP, IS-IS and OSPF are all IGPs)
IP	Internet Protocol (same as IPv4)
IPsec	Internet Protocol Security protocol
IPv4	Internet Protocol version 4 (same as IP)
IPv6	Internet Protocol version 6
IS-IS	Intermediate-System to Intermediate System routing protocol
L2TP	Layer 2 Tunneling Protocol
LAN	Local Area Network
LDP	Label Distribution Protocol
LSP	Label Switched Path
MIB	Management Information Base
MPLS	Multi-Protocol Label Switching
NBMA	Non-Broadcast Multi-Access
NMS	Network Management System
OSPF	Open Shortest Path First routing protocol
P	Provider equipment
PE	Provider-Edge equipment
PHP	Penultimate-Hop Popping
PPTP	Point-to-Point Tunneling Protocol
QoS	Quality of Service
RFC	Request for Comments
RIP	Routing Information Protocol
RSVP	Resource Reservation Protocol
RSVP-TE	Resource Reservation Protocol with Traffic Engineering extensions
VLAN	Virtual Local Area Network
VPN	Virtual Private Network
VR	Virtual Router
VRF	Virtual Routing and Forwarding instance

2. A Brief Overview of Requirements

Note: Generally, we expect that detailed requirements should go into a separate document, and our understanding is that there is a separate effort in the IETF VPN working group to define requirements. This section will therefore be very brief.

- reference requirements document in progress
- list requirements quickly, this includes:
  - ease of management (this is somewhat subjective, but is

important)
- tunneling (to support private addressing)
- multiplexing (multiple tunnels; implicit if over IP; explicit if over MPLS)
- security / privacy
- scaling
  - in size of each VPN, number of VPNs, and in bandwidth; Number of VPNs and Size of primary concern
  - of the public network
  - of each private network
    <Note: The intention is to discuss scaling implications of each of the methods in appropriate sections later, without passing judgment on the methods.>
- QoS support and SLA support
- intranets and extranets

Note: There is an issue regarding how much we want to say in this section. There is a lot which could be added. However, it is probably more appropriate to keep this very short, on the basis that the more complete description of requirements will be in a separate document.

3. Functional Components of a VPN

Basic functional components:

1. A mechanism to discover and distribute VPN membership/capability information
2. A mechanism to tunnel traffic among VPN sites
3. A means to exchange and maintain the private routes pertaining to the VPN sites connected to it and reachability information for the public backbone to be able to forward data from the VPN sites over the backbone.

- Control plane (for setting up VPNs)
- Routing plane (for routing within a VPN)
- Data plane

Tunnels for data might or might not also be used for routing.

Note: We are not sure whether or not we will need to add a functional decomposition internal to a PE. If this is needed to aid discussion in the Routing or other sections to follow, then this can be added here.

4. Customer Interface - Services and Protocols

This section to discuss the service and protocols required at the CE/PE boundary. Includes the protocols used, and what the provider

part of the network looks like in routing terms to the customer.

4.1 Customer view of Routing in the Private Network

- Uses normal IP routing. On a basic level, for a classic level 3 VPN, the customer does not see the VPN at all -- it looks like normal network equipment. This implies that the customer sees a bunch of routers, which need to be interconnected just like any old routers. The customer expects that the quality of service and routing capabilities of the VPN will be the same or very similar to other network equipment. EXCEPT for the n-squared issue (see below).

- Virtual Forwarding Instances (brief introduction).

4.1.1 Options for Routing for Intranets

- Options for routing are therefore the same as is found in any private network: One area IGP (RIP, OSPF, IS-IS, or proprietary); Hierarchical IGP. IGP within each site, and BGP or static routing between sites.

- Is PE router (or VFI within PE router) a part of the site? If an IGP is used within the site, and static route between sites, is the static route between CE and PE (probably). If BGP, maybe and maybe not.

4.1.2 Options for Routing for Extranets

<this section is tbd>.

4.1.3 Customer Edge and Provider Edge equipment

- Above discussion is from the perspective of how routing is done on an overall basis in a private network (eg, between sites).

- Is the PE box in a site or not?
  - for one-area solution, the point is mute (everything is in the same area)
  - for multi-area or multi-domain issue, could be done either way
- If PE (or VFI) is part of site: More work for provider, less work for private network
- If CE is border router: more control for customer.

4.1.4 Routing Across a Full N-Squared Mesh

Note: This section looks at routing from the perspective of the customer network. If the customer has "n" sites, then from the customer's perspective the n sites need to be interconnected and

routing has to work between the n sites. Solutions which are specific to PE to PE operation within an VPN solution will be discussed in section 7.

There are a range of possible customer views of routing in the private network. With one approach the customer only sees routing within a single site, and a link (or a small number of links) to a public network. With this approach there is no issue in the private network with a view of "n-squared" links between n sites. Similarly, in many cases the connectivity between sites may be limited. For example, there may be a small number of core sites (one or more), and each other site might be attached only to one or two core sites. Finally, in many cases the number n of sites is small enough that n-squared is still a moderate number. There are therefore a number of cases in which there is no problem with the appearance of n-squared links between n sites.

When n sites are fully connected between a large number of sites, then it will look to the customer as if the topology is very richly branched across the VPN links. How is this handled? How do you route efficiently over the n-squared mesh?

Note that this same issue comes up in other networks. For example, where multiple routers are interconnected over a frame relay or ATM subnet. Standard IP routing protocols have therefore developed ways to deal with this issue.

- discuss how this works with BGP, OSPF, and IS-IS.

- clarity PE versus CE issue (if entire topology is seen in private network, all are routers and it doesn't matter. Else there probably is no n-squared issue).

- Scaling may vary with routing approach -- since different private networks have different sizes, this is one of many reasons why multiple approaches to VPNs are needed.

4.2 Quality of Service

- QoS / SLAs

4.3 Visibility into the Provider Network

- The provider hosted part of the network may be opaque or transparent. (May appear as a separate domain with no visibility into internal structure, or customer may be able to log into all "virtual" routers and modify parameters, add routes etc)

4.4 Carrier's Carrier

We have not yet decided whether a "carrier's Carrier" service should be included. If so, then a discussion may go here.

5. VPN Establishment and Maintenance

VPN establishment and maintenance is a very important part of any solution for VPNs, and therefore is suggested as the first section of the "how does the carrier actually solve the problem" discussion.

This section covers the issues and mechanisms used for the establishment and maintenance of VPNs. There are two aspects of this - the information distributed, and the mechanisms used to distribute the information.

5.1 VPN Control Plane

Describe what the problem is and what we are trying to accomplish. Finding other parts of the VPN. Applies to both intranet and extranet case.

Information distributed may include

- Membership information (which nodes have members in which VPNs - used to establish the control plane topology / neighbor discovery)

- Tunnel end point information (used to establish tunnels for control and/or data) (we need to determine what this is other than membership info)

- Topology information (full mesh / hub & spoke)

- Reachability scheme to use (e.g. per-vpn instance or shared instance, and which protocol)

Mechanisms include

- Use BGP

- Use IGP (IS-IS or OSPF)

- Use IP multicast

- Use a VPN server / directory where nodes register and query

- Use network management system / MIB

Discuss how the mechanisms used for the control, routing and data planes may be combined in different ways, and how different schemes may carry things out in a different order. e.g. 2547: combined control and routing plane + separate data plane;

VR: separate control plane and combined routing and data plane; also tunnels before routing for VR, tunnels after routing for 2547.

5.2 VPN Membership

Which devices need information for which VPNs? This information needs to be dynamically distributed. Allows for neighbor discovery. Different schemes may use this information in different ways.

Need to deal with

- Static users / sites - permanently attached to a PE via a dedicated connection
- Dynamic users - may appear at any PE (e.g. dial users) and attach to a VPN

Schemes may include

- advertising membership/interest in a vpn to peer nodes (e.g. piggybacking on a routing protocol)
- registering and querying a directory or server with membership information

5.3 Controlling VPN Route Distribution

Common problem that all schemes must address - cannot have all VPN routes on all PEs. VPN membership information can be used to control which devices have the routes for a particular VPN.

Mechanisms:

- route filtering (2547)
- controlling tunnel establishment (VR)

5.4 Data-Plane Topology

Each vpn has a data plane topology which consists of a set of nodes interconnected via tunnels over which customer data traffic is sent. This may range from a full mesh to a hub and spoke topology, or anything in between. Different topologies may be needed for policy or scaling reasons.

- Discuss information / mechanisms that can be used to automate construction of the desired topology.

Tunnels across the backbone may be either

- shared between multiple VPNs
- dedicated to a specific VPN

Discuss scaling / QoS issues regarding this.

<maybe this last issue should be in the tunneling section?>

5.5 VPN Capabilities negotiation

Advertising and selection / negotiation of common routing / tunneling mechanisms

5.6 More detailed discussion on specific control plane mechanisms

- BGP
- LDAP
- Network Management Systems

<do we need per-mechanism sections here - can be useful to group things together in this way but could lead to duplication>

6.0 VPN tunneling

6.1 Encapsulation

- possible formats for encapsulation, IP in IP, GRE, IPsec, MPLS (L2TP is not in scope)

- overhead and control mechanisms may vary.

- when a packet arrives it needs to be determined which VPN it belongs to. In many cases any one tunnel will be associated with a single VPN. The method of making this mapping will therefore depend upon the method of encapsulation which is used. This could use the MPLS label, the IP address (in the case of IP in IP encapsulation), IPsec security association, or a VPN ID. If the latter, then somewhere in a GRE header, IPsec header, or new form of encapsulation a place needs to be found to put a VPN ID.

- Somewhere here we need to discuss scaling, in terms of the numbers of tunnels. We probably mention the issue here, and give more detail below.

6.2 Tunnel Establishment

- Explicit signaling to determine multiplexing value, vs distribution

of multiplexing value without explicit signaling (e.g. piggybacking of MPLS label on some other protocol)

Different tunneling schemes use different methods:

- IPSec tunnels - advertise endpoint IP address and use IKE signaling to establish the tunnel.

- MPLS LSP - advertise the MPLS label. Could also advertise endpoint IP address and use RSVP-TE / CR-LDP to establish an LSP.

- IP/IP tunnels - no signaling, no multiplexing field - advertise outer IP address

- GRE tunnels - no signaling, multiplexing field?

6.3 Hierarchical Tunnels

This section would discuss shared vs dedicated per-VPN tunnels and the scaling issues involved. Some of the text below (6.4) may be moved here. Hierarchy applies to both MPLS and non-MPLS tunnels. Discussion of multipoint operation might also go here.

There might also be a new top level section "Hierarchical Tunnels" (or perhaps "Tunnel Scaling") that is independent of tunnel maintenance, since scaling and maintenance are separate issues. Use of hierarchical tunnels appears to be primarily about scaling, but may also have other features.

6.4 Tunnel Maintenance

- Generally, for each tunnel, we need to set it up, and over time make sure it is still up. There may optionally be a way to remove the tunnel in the rare chance that we are done.

- Maintenance also related to routing model used - with VR there's a routing instance running over the tunnel so no tunnel specific mechanisms are needed, with Aggregated this isn't the case .

- Tunnel maintenance may be explicit or implicit. For example, for IP in IP encapsulation, if you have been told that you can encapsulate in address "w.x.y.z" for a particular VPN, you might just send a packet there. Similarly in some cases MPLS tunnels can be setup by simply advertising the label to use for specific sets of traffic. Advantages of implicit (null) maintenance: Don't need n-squared protocol exchanges. Disadvantage: Tunnel might not work due to a network failure. There might have been some other way around the failure. Note that for some approaches n-squared adjacencies might

need to exist for the routing protocol anyway (see routing section).

- Method for tunnel maintenance will vary with encapsulation.

6.4.1 Maintaining IP-in-IP and/or GRE Tunnels

We probably want to mention that tunneling over IP might be implicitly multipoint to point. If you have "n" tunnel endpoints, then at any one place you need only "n" pieces of information (the IP address to tunnel to that point). If we ignore routing issues (see section 8 below), then this implies that the scaling is O(n).

6.4.2 Maintaining LSPs as tunnels

- Label Exchange

- How do you decide which VPN is mapped to which LSP?

For this section, I think that piggybacking labels on BGP simply makes BGP one possible signaling protocol for LSPs, although perhaps one that in some usages becomes useful only in the specific case that you know that you have a PE to PE tunnel, and you want to multiplex multiple VPN-specific tunnels inside the PE to PE tunnel.

With point to point tunnels between PEs, or per-VPN between PEs, the scaling would be O(n-squared). If a single level of tunnel is used, each VPN-specific, and if tunnels are point to point, this could be a lot of tunnels. This would be a problem for large networks (hundreds of PEs). This problem can be solved through two compatible mechanisms:

- Use hierarchical LSPs: VPN-specific tunnels can be multiplexed inside PE-PE tunnels.

- Multipoint to point tunnels: The PE to PE tunnels can be multipoint to point. This implies that if we have "n" PEs, then there are only "n" LSPs within the core of the network. While each LSP could be relatively complex (in that it has multiple branches), nonetheless on each link there are only a maximum of n LSPs, and at each node in the network the total amount of information is proportional to n times the number of direct neighbors (ie, number of links * n is an upper bound on information).

- VPN-specific tunnels which are multiplexed inside the PE to PE tunnels can also either point to point or multipoint to point. If the latter, then the amount of information for each ingress PE is proportional to the number of VPNs that it supports times the number of places that each VPN needs to send traffic. The amount of

information for each egress PE is proportional to the number of VPNs that it supports. This is of course a minimal amount of information which is needed by any solution. (note, also see routing scaling, below).

6.4.3 Maintaining IPsec Associations as Tunnels

6.5 Multiplexing

<this might be a separate section, or might be discussed in each section above.>

7. Routing for VPNs Across the Public Network

- This refers to carrying of Private VPN routing information across the public network.

7.1 Virtual Forwarding Instances

- PE routers may end up supporting a large number of VPNs, and therefore a large number of routing instances. This makes scaling hard in PE routers. On the other hand, the resource load on a particular PE is largely linearly proportional to the number of VPNs that the PE router supports, and to the size of the VPNs.

- Note that with any Network Based VPN, the PE gear is involved in routing for the private networks that it supports. This implies that scaling of the PE gear will by definition be no better than proportional to the number of VPNs which the PE supports times the average size of the VPNs.

7.2 Virtual Routers

- Route across the tunnels, and treat them as normal interfaces.
  - as point to point tunnels
  - as an NBMA network
  - as a broadcast/multicast network
- Refer to overlay routing
- Since we are routing across the tunnels, failure of the tunnels is detected by the routing protocols.

7.3 Aggregated Routing Model

Aggregated routing means one routing instance is used to carry routes of many or all the VPNs supported by the PEs. This implies that routing needs to be separated from data forwarding (the tunnels are used for data forwarding). This model routing may be piggybacked on a common routing protocol used for multiple VPNs, possibly involving

BGP.

The common routing protocol can be either the same protocol and instance used for SP network routing, or a different routing instance.

- Could use IGP or a BGP. Problems with IGP (OSPF or IS-IS implies flooding all information throughout area). BGP allows separation of information.

- Since we are not routing across the tunnels, other means are needed to ensure that if the tunnels fail then either the tunnels are rapidly re-constructed, or routing within the private network responds.

7.3.1 Aggregated Routing with OSPF or IS-IS

- Link state protocols broadcast routing info throughout an area

- This is a problem (wrong way to distribute VPN routing information). Scaling implications.

7.3.2 Aggregated Routing with BGP

- Flexibility regarding where information goes.

7.3.3 Managing PE to PE Backbone Networks

7.3.4 Partitioning of Routing Information with BGP

- May have separate set of route reflectors for Internet and VPN routes

- May partition VPN routes among route reflectors

7.4 Inter-Domain Routing and Route Aggregation

This refers to dealing with VPNs which span multiple carrier routing domains, and the routing implications thereof.

7.5 Header Lookups in the VFIs

- VFIs may (under the most straightforward implementation) have to do more than one header lookup. Depending upon how the tunneling is done, there could be several. Ways to reduce this are in the following sections.

7.6 Penultimate Hop Popping for MPLS

- How PHP reduces header lookups.

- Situations in which you can or can't use this.

7.7 Demultiplexing to Eliminate the Tunnel Egress VFI Lookup

- How this might be done
  - Using MPLS
  - FECs for each LSP is mapped to the next hop
  - Requires more LSPs, but less forwarding lookups. this is a straightforward engineering tradeoff of resources (which resource would be rather use).

- This can also be done with other encapsulations. This uses more instances / values of the 'multiplexing' field, whatever that is.

8. QoS and IP Differentiated Services

8.1 QoS in the Public Network

- VPNs may use Diff Serve, as one way to obtain the QoS which is desired in the VPN.
- Bandwidth guarantees for LSPs

8.2 Mapping QoS from the Private to Public Network

- IP Diff Serve, as used in the VPN, may be mapped to IP Diff Serve across the carrier network.

9. Security Issues

Note: We need to think hard about this section: This is an important issue for VPNs. Again in a framework document we only discuss possible solutions and their tradeoffs, we don't pick any solution.

9.1 Security of User Data

- Need for authentication and/or encryption.

- Need to protect against spoofing (sending traffic which is alleged to come from inside the private network), and denial of service attacks. <question: what other types of attacks do we want to protect against? Do we actually want to mention them here, or will we be helping attackers? My suspicion is that some attackers already know as much or more than we are likely to include.>

- CPE to CPE (is this and also host to host outside of the scope of the working group? At least it should probably be listed as a

possibility)

- PE to PE protection (encryption and/or authentication) does not protect CE to PE link, but protects data between PEs.

9.2 Security of Routing Information

If overlay routing (with VRs and tunnels) then is tied to security of the tunnels, which is the same as the security of the user data. With aggregated model, is tied to security of a single instance of routing information. In both cases we also depend on the security of the PEs (assuming that if a PE is compromised then VRs within the PE will also be compromised).

For both models, routing information is exchanged across the CE-PE boundary. We need to consider whether this is another possible hole.

We should also discuss the impact of configuration errors. It is not clear a priori whether this is the same or different with the different approaches (we will need to work this out in detail).

9.3 Security of Membership Information

10. Interoperability

- It is likely that interoperability of any one VPN solution is based on the completeness of the standard, and as such is outside of the scope of the framework document.

- Interoperability between different VPN solutions might be discussed here.

11. Intellectual Property

TBD.

12. Authors' Addresses

Ross Callon
Juniper Networks
1194 N. Mathilda Avenue,
Sunnyvale, CA 94089
+1-978-692-6724
E-mail: rcallon@juniper.net

Bryan Gleeson
Nortel Networks
2305 Mission College Blvd

Santa Clara CA 95054
+1-408-565-2625
E-mail bgleeson@shastanets.com

Eric C. Rosen
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA, 01824
+1-978-244-8918
E-mail: erosen@cisco.com

Chandru Sargor
Cosine Communications
1200 Bridge Parkway
Redwood City, CA 94065
+1-650-637-2416
Email: Chandramouli.Sargor@cosinecom.com

Jieyun Jessica Yu
Cosine Communications
1200 Bridge Parkway
Redwood City
CA 94065
Tel: (650) 628-4881
Email: jyy@cosinecom.com

13. References

tbd.