TRILL Weiguo Hao
Yizhou Li
Donald Eastlake
Internet Draft Huawei
S. Hares
Hickory Hill Consulting
Muhammad Durrani
Brocade
H. Zhai
ZTE Corporation

Intended status: Informational May 20,2014
Expires: November 2014

Analysis of Active-Active Connection Solutions
draft-hao-trill-analysis-active-active-02.txt

Status of this Memo

Copyright Notice

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

Draft [TRILL-Active-PS] lists basic problems which any active-active solutions should address, these problems include frame duplications, loop, MAC address flip-flop and unsynchronized information among member RBridges. For each problem, there may be multiple ways to deal with it. Some solutions solve all or most of the problems

listed, and at the same time introduces extra issues. This draft tries to analyze and compare the different solutions for each of the issues, gives a brief summary on the pros and cons, and/or the applicable scenarios.

Table of Contents

1. Introduction

The IETF TRILL (Transparent Interconnection of Lots of Links) [RFC6325] protocol provides loop free and per hop based multipath data forwarding with minimum configuration. TRILL uses IS-IS [RFC6165] [RFC6326bis] as its control plane routing protocol and defines a TRILL specific header for user data.

Classic Ethernet(CE) devices typically are multi-homed to multiple edge RBridges which form an edge group. All of the uplinks of CE are bundled as a Multi-Chassis Link Aggregation (MC-LAG). An active-active flow-based load sharing mechanism is normally implemented to

achieve better load balancing and high reliability. A CE device can be a layer 3 end system by itself or a bridge switch through which layer 3 end systems access to TRILL campus.

Draft [TRILL-Active-PS] lists the following problems which any active-active solution should address:

```
                                   +------+
                                   | CEx  |
                                   +------+
                                       |
                                    +------+
                                   |(RBx) |
                                   +------+
                                       |
                               -------------------
                              /                    \

                             |                       |
                             |   TRILL Campus        |
                             |                       |
                              \                     /
                               --------------------
                                    |     |    |
                                --------      |     --------
                               |              |             |
                             +------+      +------+      +------+
                             |(RB1) |      |(RB2) |      | (RBk)|
                             +------+      +------+      +------+
                              |              |   |          |
                              |   -----------|   |------    |
                              |   |LAG1          LAG2   |   |
                             +------+                 +------+
                             |  CE1 |                 | CE2  |
                             +------+                 +------+
```

Figure 1 TRILL Active-Active Access Scenario

1. Frame duplications

2. Loop

3. Address flip-flop

4. Unsynchronized information among member RBridges

For each problem, there may be multiple ways to deal with it. And some solutions solve all or most of the problems listed, and at the same time introduces extra issues. This draft tries to analyze and compare the different solutions for each of the issue, gives a brief summary on the pros and cons, and/or the applicable scenarios. The co-authors believe such analysis is helpful to design a more completed solution in future.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT","SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

The acronyms and terminology in [RFC6325] is used herein with the following additions:

BUM - Broadcast, Unknown unicast, and Multicast.

CE - Refer to [CMT]. The device can be either physical or virtual equipment.

CMT - Coordinated Multicast Trees [CMT].

Edge group - a group of edge RBridges to which at least one CE is multiply attached using MC-LAG. When multiple CEs attach to the exact same set of edge RBridges, those edge RBridges can be considered as a single edge group. One RBridge can be in more than one edge group.

LACP - Link Aggregation Control Protocol.

LAG - Link Aggregation, as specified in [8021AX].

3. Frame duplications

Problem:

Frame duplication may occur when a remote host sends multi-destination frame to a local CE which has an active-active connection to the TRILL campus.

Solution:

To avoid local CE receiving multiple copies from a remote RBridge, the designated forwarder (DF) mechanism should be supported. DF

election mechanism allows only one port in one RB of a MC-LAG to forward multicast traffic from TRILL campus to local access side for each VLAN. The basic idea of DF is to elect one RBridge per VLAN from an edge group to be responsible for egressing the multicast traffic.

Each RB in an edge group elects a DF using same algorithm which guarantees the same RB elected as DF per MC-LAG per VLAN. [draft-hao-trill-dup-avoidance-active-active-00] describes the detail DF mechanism and TRILL protocol extension for DF election. The RB that is elected as a DF for a given VLAN will forward multi-destination traffic in the egress direction towards the CE. All non-DF RBs drop multi-destination traffic in the egress direction towards the CE. All edge RBs, including DF and non-DF, can ingress the traffic to TRILL campus as usual. As DF election is based on VLAN, DF ports for different VLANs can be on different edge RBs. Thus egress bound multicast traffic can be load balanced among multiple edge RBridges in an edge group on per VLAN basis.

4. Loop

Problem:

If a CE sends a broadcast, unknown unicast, or multicast (BUM) packet through DF port to a ingress RB, the RB will forward that packet to all or subset of the other RBridges that only have non-DF ports for that MC-LAG. Because BUM traffic forwarding to non-DF port isn't allowed, in this case the frame won't loop back to the CE.

If a CE sends a BUM packet through non-DF port to a ingress RB, say RB1, then RB1 will forward that packet through TRILL campus to DF RBridge for the MC-LAG. In this case the frame will loop back to the CE.

Solution:

A traffic split-horizon filtering mechanism should be used to avoid looping back among RBridges in a edge group.

Split-horizon mechanism relies on ingress nickname to check if a packet's egress port belongs to same MC-LAG with the packet's incoming port to TRILL campus. The following sections describe different nickname allocation schemes:

4.1. Independent nickname allocation

Each ingress RBridge allocates a unique nickname for each MC-LAG independently. It is not required that the nickname provisioned on all involved edge RBridges remains the same for one corresponding MC-LAG.

When the ingress RBridge receives BUM traffic from an active-active accessing CE device, the traffic will be injected into TRILL campus, ingress nickname is the allocated unique nickname on ingress RB.

When an egress RBridge receives the BUM traffic from the TRILL campus, it checks the ingress nickname in the TRILL header and filters out the traffic on all local interfaces connected to the same CE. Each egress RBridge should track the nickname(s) associated with the other RBridge(s) with which it has a shared multi-homed LAG. The solution has limited nickname allocation scalability issue, because each RBridge needs allocate per nickname per MC-LAG.

4.2. Consistent nickname allocation per MC-LAG

Edge RBridges forming an MC-LAG in an edge group are assigned a globally unique pseudo-nickname. If multiple MC-LAGs exist, edge BRridges for each individual MC-LAG should be assigned such a pseudo-nickname. It should be guaranteed that pseudo-nickname provisioned on all involved edge RBridges remains the same for one corresponding MC-LAG.

When a ingress RBridge receives traffic from a active-active accessed CE, it performs TRILL encapsulation with the pseudo-nickname as ingress nickname. When the traffic comes to each egress RBridge, the egress RBridge checks ingress nickname in TRILL header and filters out the traffic on all local interfaces connected to the same CE. Each egress RBridge relies on the pseudo-nickname to filter out the frame on all local interfaces connected to the same CE.

4.3. Consistent nickname allocation per edge group RBridges

```
                  +-----------+
                  |   (RB4)   |
                  +-----------+

                   |     |    |
           --------      |     --------
          |              |             |
        +------+     +-------+      +------+
        |(RB1) |     | (RB2) |      | (RB3)|
        +------+     +-------+      +------+
          *   |         *  |  ^        * |  ^
          *   |         *  |  ^^^^^^^^^^^^^^^^
          *   ----------*--------------*-|     ^
          ****************************** |        ^
  MC-LAG1 *                      MC-LAG2 | MC-LAG3   ^
      +------+                       +------+    +------+
      |  CE1 |                       | CE2  |    | CE3  |
      +------+                       +------+    +------+
```

Figure 2 Consistent nickname allocation per edge group RBridges scenario

An edge group forming one or multiple MC-LAGs is assigned a globally unique pseudo-nickname. All MC-LAGs corresponding to the edge group share same pseudo-nickname to save nickname space. It should be guaranteed that pseudo-nickname provisioned on all involving edge RBridges in an edge group remains same.

In above figure 2,CE1 and CE2 are active-active accessed to RB1,RB2 and RB3,CE3 is active-active accessed to RB2 and RB3. Globally unique pseudo-nickname of p-nick1 is assigned to the edge group which contains RB1,RB2 and RB3, p-nick2 is assigned to the edge group which contains RB2 and RB3. P-nick1 is used for MC-LAG1 and MC-LAG2, p-nick2 is used for MC-LAG3.As only one pseudo-nickname is assigned for MC-LAG1 and MC-LAG2, so nickname consumption is lower than the consistent nickname allocation method per MC-LAG.

If one or more CE's uplinks occur link failure, the CE will connect to new edge group RBs. At this time, the CE will use new pseudo-nickname corresponding to the new edge group as ingress nickname.

Take the topology shown in figure 2 as example. If the link between CE1 and RB1 fails, CE1 will connect to the edge group which contains RB2 and RB3 only. Then p-nick2 will be used as ingress nickname for CE1. If RB1 encounters node failure, both CE1 and CE2 will connect

to the rest edge RBs which are RB2 and RB3. Then p-nick2 is used as ingress nickname for all of CE1, CE2 and CE3.

To enhance network convergence, access link failure and edge node failure should be detected by each edge RBridge in a edge group as fast as possible.

4.4. Comparison

Solution	Independent Allocation	Consistent Allocation per MC-LAG	Consistent Allocation per Edge Group
Nickname consumption	High	Medium	Low
Scalability	Low	Medium	High

5. Address flip-flop

MAC learning in TRILL can be performed either in data plane or control plane. When a local host h1 attaches to multiple edge RBridges, learning at the remote host for h1 may have MAC flip-flop problem.

There are different ways to avoid this for data plane learning and control plane learning scenarios.

5.1. Data plane learning mode

Problem:

For data plane learning mode, to avoid mac address flip-flop on remote RBs, a pseudo-nickname [TRILLPN] solution was proposed. The basic idea is to use a virtual RBridge of RBv with a single pseudo-nickname to represent an edge group that MC-LAG connects to. Any member RBridge of that edge group should use this pseudo-nickname rather than its own nickname as ingress nickname when it injects TRILL data frames to TRILL campus. The use of the nickname solves the address flip flop issue by making the MAC address learnt by the remote RBridge bound to pseudo-nickname.

If DF-election mechanism is used for frame duplication prevention, access ports on an RB are categorized as three types: non mc-lag, mc-lag DF port and mc-lag non-DF port. The last two types can be called mc-lag port. For each of the mc-lag port, there is a pseudo-nickname associated. If consistent nickname allocation per edge group RBridges is used, it is possible that same pseudo-nickname associated to more than one port on a single RB. A typical scenario is that CE1 is connected to RB1 & RB2 by mc-lag1 while CE2 is connected to RB1 & RB2 by mc-lag 2. In order to save the number of pseudo-nickname used, member ports for both mc-lag1 and mc-lag2 on RB1 & RB2 are all associated to pseudo-nickname pn1.

On the other hand, pseudo-nickname introduces another issue, which is incorrect packet drop by RPF check failure. Due to edge RBridges which use a pseudo-nickname other than own nicknames as the ingress nickname (Eg. Nick-Y) when the RBbridge forwards BUM traffic from local CE, the traffic will be treated by an RBridge (RBn) sitting between the ingress RB and distribution tree root as traffic whose ingress point is RBv. If same distribution tree is used by these different edge RBridges, the traffic may arrive at RBn from different ports. Then the RPF check fails, and some of the traffic receiving from unexpected ports will be dropped by RBn.

Solutions:

To overcome the RPF check failure issue, the following three solutions have been proposed: CMT, centralized replication and tunneling among edge RBs. For local replication behavior on the ingress RBridge, CMT, centralized replication and tunneling among edge RBs solutions should consider all the above access ports type and may be different. The following subsections will give more details.

5.1.1. CMT

CMT [CMT] solution allows edge RBridges to specify different distribution trees to forward BUM traffic from a connecting CE device by using a new IS-IS Affinity sub-TLV. Remote RBridges calculate their forwarding tables and derive the RPF for distribution trees based on the distribution tree association advertisements. The BUM traffic injected to TRILL campus by ingress RB will not return to ingress RB again.

When an ingress RBridge of RB1 receives BUM traffic from an active-active accessing CE1 device, local replication behavior on RB1 is as follows:

1. Local replication to non mc-lag ports as per RFC6325.

2. Local replication to the ports associated with the same pseudo-nickname as that associated to the incoming port as per RFC6325.

3. Local replication to the mc-lag DF port associated with different pseudo-nickname as per RFC6325. Do not replicate to mc-lag non-DF port associated with different pseudo-nickname.

The above local forwarding behavior on the ingress RB of RB1 can be called CMT local forwarding behavior.

In this solution, it's required to establish multiple distribution trees in a TRILL campus, i.e. if a CE is active-active accessed to 4 edge RBridges, at least 4 distribution trees are required. No hardware upgrade is needed for RBridges in the TRILL campus, only software upgrade is needed.

5.1.2. Centralized replication

The solution has all ingress RBs send BUM traffic receiving from local active-active connecting CE to a centralized node via unicast TRILL encapsulation. When the centralized node receives the BUM traffic, it decapsulates the traffic and forwards the BUM traffic to all destination RBs using a distribution tree established via the TRILL base protocol. To avoid RPF check failure on a RBridge sitting between the ingress RBridge and the centralized replication node, some change of RPF calculation algorithm is required. RPF calculation on each RBridge should use the centralized node as ingress RB instead of the real ingress RBridge of RBv to perform the calculation. The BUM traffic injected to TRILL campus by ingress RB will return to the ingress RB via distribution tree established as per TRILL base protocol. [draft-hao-trill-centralized-replication-00] describes the detail centralized replication solution.

When the ingress RBridge of RB1 receives BUM traffic from an active-active accessing CE1 device, one copy of the traffic is forwarded locally to other CE devices connecting via MC-LAG ports that share same pseudo-nickname with the port connecting to CE1, another copy of the traffic will be sent to a centralized node via unicast TRILL encapsulation. Then it is replicated and forwarded to all destination RBridges including RB1 itself along TRILL distribution tree established as per TRILL base protocol. When RB1 receives the

TRILL multicast traffic, it will decapsulate TRILL encapsulation and forward it to all local CE devices except CE1, if these CE devices connect to RB1 via non-MC-LAG ports and MC-LAG DF ports. For other CE devices which are connected to RB1 via MC-LAG non-DF ports, the traffic will be dropped and will not be forwarded to these CEs.

In summary, local replication behavior on RB1 is as follows:

1. Local replication to the ports associated with the same pseudo-nickname as that associated to the incoming port as per RFC6325.

2. Do not replicate to mc-lag port associated with different pseudo-nickname.

3. Do not replicate to non mc-lag ports.

The above local forwarding behavior on the ingress RB of RB1 can be called centralized local forwarding behavior it is different from CMT local forwarding behavior.

If ingress RB of RB1 itself is the centralized node, BUM traffic injected to TRILL campus won't loop back to RB1. In this case, the local forwarding behavior is same with CMT local forwarding behavior.

In this solution, it's required to consume more network bandwidth between ingress RB and distribution tree root node than CMT solution. Both hardware and software upgrade are required on edge RBs participating in active-active connection and the distribution tree root node. This solution doesn't require multiple distribution trees in TRILL campus.

5.1.3. Tunneling among edge RBs

This solution allows only a selected edge RBridge in an edge group participating in active-active access to be responsible for forwarding BUM traffic from connecting CE to TRILL campus along distribution tree per TRILL base protocol. All other edge RBridges in the edge group send BUM traffic from connecting CE to the selected edge RBridge through unicast TRILL encapsulation. When the selected edge RBridge receives unicast TRILL traffic from RB1 in a same edge group, the selected RBridge decapsulates the unicast TRILL packet. Then it forwards the BUM traffic through TRILL multicast encapsulation to TRILL campus along distribution tree established as per TRILL protocol.

The traffic will reach all destination RBridges and will loop back to ingress RBridge of RB1 similar to the above centralized

replication solution, so local forwarding behavior on RB1 is same with the centralized local forwarding behavior.

If ingress RBridge of RB1 is selected RBridge, the BUM traffic that is injected into TRILL campus won't loop back to RB1, the local forwarding behavior is same with the CMT local forwarding behavior.

In this solution, it's required to consume more network bandwidth among edge RBs. Both hardware and software upgrade are required on edge RBs participating active-active connection. This solution only needs one distribution tree in TRILL campus.

5.1.4. Comparison

Data Plane Mode:

Solution	CMT	Centralized replication	Tunneling among edge RBs
Dist tree required for N-active scenario	N	1	1
Network bandwidth consumption	Low	High	High
Local forwarding behavior on ingress RB	CMT	Ingress RB is the centralized node: CMT	Ingress RB is selected RB: CMT

```
      |                        |         | Other ingress RB:        |Other
ingress RB:           |

      |                        |         |  centralized             |
centralized                |

      +------------------------+---------+-------------------------
+----------------------------+

      |  Software upgrade      | All RBs |        All RBs           |    root
and edge nodes     |

      +------------------------+---------+-------------------------
+----------------------------+
```

      | Hardware upgrade | No | root and edge nodes | root
and edge nodes |

      +------------------------+---------+-------------------------
+---------------------------+

5.2. Control plane learning mode

If a CE device is multi-homed to multiple edge RBs in active-active mode, each edge RB should announce the MAC of its attached end systems to all other RBs through ESADI-like control protocol. Remote RBriges will learn the MAC association with different ingress RB nicknames and generate multiple MAC forwarding entries in ECMP mode. All edge RBs should disable the data plane MAC learning function. MAC to nickname association should be learned only through the control plane.

Pseudo-nickname mechanism was basically designed to avoid MAC address learning flip-flop when a MAC address could be learnt to more than one RBridge. With control plane MAC learning, pseudo-nickname is not required since multiple mac to nickname entries can be leaned for the same MAC. The problem of RPF check failure for multicast frame caused by pseudo-nickname mechanism is not an issue here.

In the control plane MAC learning solution, if an edge RB participating TRILL active-active access receives BUM traffic from connecting CE device, it uses its own nickname as ingress nickname instead of pseudo-nickname to ingress data frame into a TRILL campus.

This method requires hardware and software changes.

6. Unsynchronized information among member RBridges

Problem:

A local Rbridge, say RB1 in MC-LAG1, may have learned a VLAN and MAC to nickname correspondence for a remote host h1 when h1 sends a packet to CE1. The returning traffic from CE1 may go to any other member RBridge of MC-LAG1, for example RB2. To avoid always flooding for unicast traffic on RB2, MAC address should be synchronized among the edge RBridges in a edge group.

To ensure DF election consistency, dynamic joined VLAN through VLAN registration protocol (VRP) (GVRP or MVRP) and multicast group through IGMP or MLD protocol should be synchronized among all RBridges in a edge group.

Solution:

Synchronization mechanism should be provided to ensure information consistency among all edge RBridges in a edge group. Three synchronization solutions as follows are provided.

6.1. RBridge channel based communication protocol

RBridge channel based communication protocol among all RBridges in a edge group is introduced to implement synchronization. The communication protocol is restricted to RBridge nodes in each edge group, other RBridges in TRILL campus needn't involve. A new type of RBridge Channel message should be given by a Protocol field in the RBridge Channel Header to indicate synchronization information in the payload. RBridge channel message is forwarded through TRILL data plane. Transmission delay is relatively low.

6.2. TRILL LSP extension

TRILL LSP can be extended to implement synchronization among all edge RBridges. Synchronization information is conveyed through new TLVs or sub-TLVs in TRILL LSP. Because TRILL LSP is flooded to all RBridges in TRILL campus, so it may cause campus wide fluctuation. TRILL LSP is forwarded through control plane. Transmission delay is relatively high.

6.3. ESADI extension

TRILL ESADI can be extended to implement synchronization among all edge RBridges. Currently ESADI only support MAC synchronization, it doesn't support VLAN and multicast group information synchronization. Similar to the solution of RBridge channel based communication protocol, ESADI message is forwarded through TRILL data plane. Transmission delay is relatively low.

6.4. Comparison

Solution	RBridge channel based	TRILL LSP extension	ESADI extension
Flooding scope	Edge group	Campus wide	Edge group
Forwarding	Data plane	Control plane	Data plane

7. Solution summary

The possible mechanisms for each individual problem listed in [TRILAA] are described and compared in this document. The readers can compile a complete solution from these mechanisms.

If there are multiple mechanisms for an individual problem, the readers can picked up the most appropriate one based on the scenario. For example, to solve MAC address flip-flop problem, if control plane learning is not possibly supported, pseudo-nickname mechanism via data plane MAC learning should be used.

When a mechanism is used to solve an individual problem, other additional issues may be introduced and a complete solution should be carefully designed to solve those non-generic issues. For example, when pseudo-nickname mechanism is used to solve MAC address flip-flop problem, RPF check failure issue is incurred. Three mechanisms, CMT, centralized replication and tunneling among edge RBs, can be used to solve the RPF check failure issue. If any one of them is used, local forwarding behavior on ingress RBridges should be carefully designed to ensure BUM traffic not duplicated or looped to ingress RBridge's local connecting CE devices.

In summary, the candidate mechanism for each of the problem is listed as follows.

```
   +----------------------
+------------------------------------------------------------------+

   |       Problem        |
Mechanisms                            |

   +----------------------
+------------------------------------------------------------------+

   | Frame duplication    |                          DF
election                            |

   +----------------------+---------------------------------------
+-------------------------+

   |         Loop         |      Data plane MAC learning          |
Control plane     |

   |                      |                                       |       MAC
learning      |

   |                      |---------------------------------------
+-------------------------+

   |                      |  CMT | Centralized  | Tunneling
|                         |

   |                      |      | replication  | among edge RBs
|                         |

   +----------------------+---------------------------------------
+-------------------------+

   |  Address flip-flop   | Independant alloc|  Consistent alloc  |
Consistent alloc     |

   |                      |                  |    per LAG         |      per
Edge Grp       |

   +----------------------+------------------------+--------------+--
+----------------------+

   |  Unsynchronized      |                        |
|                      |

   |   information        | RBridge channel based  |  LSP extension  |   ESADI
extension    |

   +----------------------+------------------------+----------------
```

+----------------------+

8. Security Considerations

This draft does not introduce any extra security risks. For general TRILL Security Considerations, see [RFC6325].

9. IANA Considerations

This document requires no IANA Actions. RFC Editor: Please remove this section before publication.

10. References

10.1. Normative References

[1] [RFC6165] Banerjee, A. and D. Ward, "Extensions to IS-IS for Layer-2 Systems", RFC 6165, April 2011.

[2] [RFC6325] Perlman, R., et.al. "RBridge: Base Protocol Specification", RFC 6325, July 2011.

[3] [RFC6326bis] Eastlake, D., Banerjee, A., Dutt, D., Perlman, R., and A. Ghanwani, "TRILL Use of IS-IS", draft-eastlake-isis-rfc6326bis, work in progress.

10.2. Informative References

[4] [TRILAA] Li,Y., et.al., " Problem Statement and Goals for Active-Active TRILL Edge ", draft-ietf-trill-active-active-connection-prob-03, Work in progress, May 2014.

[5] [TRILLPN] Zhai,H., et.al., "RBridge: Pseudonode Nickname", draft-hu-trill-pseudonode-nickname, Work in progress, November

2011.

[6] [CMT] [CMT] Senevirathne, T., Pathangi, J., and J. Hudson,

"Coordinated Multicast Trees (CMT)for TRILL", draft-ietf-trill-cmt-03.txt Work in Progress, April 2014

[7] [RFC7178] - D. Eastlake, V. Manral, L. Yizhou, S. Aldrin, D. Ward, "Transparent Interconnection of Lots of Links (TRILL):

RBridge Channel Support", RFC7178, May 2014.

[8] [RFC6439] Perlman, R., Eastlake, D., Li, Y., Banerjee, A., and F. Hu, "Routing Bridges (RBridges): Appointed Forwarders", RFC 6439, November 2011.

[9] [ESADI] H. Zhai, F. Hu, et al, "TRILL (Transparent Interconnection of Lots of Links): ESADI (End Station Address Distribution Information) Protocol", draft-ietf-trill-esadi-05.txt, February 2014, working in progress.

Authors' Addresses

Weiguo Hao
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China
Phone: +86-25-56623144
Email: haoweiguo@huawei.com

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China
Phone: +86-25-56625375
Email: liyizhou@huawei.com

Susan Hares
Hickory Hill Consulting
7453 Hickory Hill
Saline, CA 48176
USA
Email: shares@ndzh.com

Muhammad Durrani
Brocade communications Systems, Inc
mdurrani@Brocade.com

Hongjun Zhai
ZTE Corporation
68 Zijinghua Road
Nanjing 200012 China

Phone: +86-25-52877345
Email: zhai.hongjun@zte.com.cn