

BESS Working Group Meeting Minutes

Session: BGP Enabled Services (BESS)

Time: Thursday, 09:00–11:00

Chairs: Matthew Bocci, Jeffrey Zhang

Scribe: Mankamana

1. Working Group Update

Presenter: Chairs

Scheduled time: 09:00–09:15

The chairs opened the meeting, reviewed the Note Well and meeting procedures, and asked participants to use the Datatracker onsite tool for attendance and the microphone queue.

The chairs reminded presenters that slides should be submitted several days before the meeting. A number of presentations had been submitted very late, making preparation of the session difficult.

The chairs highlighted the newly formed **BGP Directorate Review Team**. BGP Directorate reviews should be considered input to the whole working group, rather than comments that only document authors need to address.

Document Status

No new RFCs had been published since the previous IETF.

Several documents were in the RFC Editor queue. Authors/editors were asked to respond promptly to RFC Editor questions.

The BGP-based multihoming in VPLS document was discussed as being back in the publication process after previous delays and editorial changes.

EVPN BFD

The EVPN BFD document had WG consensus to publish but did not pass the implementation poll. BESS requires at least one implementation for Standards Track documents.

Publication as **Experimental** was discussed as one possible way forward.

A participant/co-author stated support for moving the document to Experimental because the technical comments had been addressed.

The chairs noted that an Experimental document should clearly define the experiment and what outcome would justify potential progression to Standards Track in the future.

Next step: Chairs to determine the appropriate publication path.

EVPN/MVPN PE Interoperability

Comments had been provided and an updated version was expected before the shepherd write-up.

An IPR-related issue discussed earlier in IDR was also raised. The understanding from IDR was that there

was no relevant IPR, but this needed to be clearly reflected on the BESS list.

EVPN Control Plane for Geneve

The document had WG consensus but failed the implementation poll.

It was being moved toward Experimental status. The document needs text defining the experiment and its success criteria.

Action: Authors to update the document accordingly.

BGP Link Bandwidth / eBGP DMZ

The document was in WG Last Call.

Ketan : A possible IPR issue had previously been raised. During the IDR meeting, the individual concerned stated that there was no IPR, and that information was forwarded to the BESS mailing list.

The chairs indicated that the document should therefore be able to progress.

Alexander (Sasha) Vainshtein, Ribbon Communications, noted that the L2 Gateway work had multiple implementations and deployment in a live network and supported repeating the WG Last Call.

Next step: Chairs may issue another WG Last Call.

The D-Path for EVPN interconnected networks had received a BGP Directorate review indicating that it was not ready. Authors were asked to address the comments.

BGP Directorate Reviews

Ketan : BGP Directorate perspective emphasized that Directorate reviews are intended as input to the entire working group. WG participants were encouraged to discuss, agree, disagree, and contribute to resolution rather than leaving all responses to document authors.

RFC 7432-bis and EVPN Interoperability Modes

Outstanding comments on RFC 7432-bis were discussed.

The chairs considered whether the EVPN interoperability-modes draft should be merged with RFC 7432-bis.

Jorge Rabadan, Nokia, commented that some parts could fit naturally into RFC 7432-bis because they concern EVPN broadcast-domain behavior, but other Layer-3 aspects are related to RFCs 9135 and 9136 and would not naturally belong in RFC 7432-bis. He therefore did not support simply merging the entire draft.

Himanshu Shah requested a summary of the differences between RFC 7432 and the bis document.

Jorge Rabadan, Nokia, explained that early versions had contained a change-log-style section, but it had later been removed because it was considered noise. He suggested that such information might potentially be placed in an appendix.

Another participant argued that because RFC 7432-bis is a very large document, a delta/change

summary would make it substantially easier to review.

Next step: Continue the discussion on the mailing list.

Keyur Patel: MUP Safi draft has implementation. We have not filed implementation report since BESS does not require such.

Expired WG Documents

The chairs noted that BESS has a large number of expired WG drafts, including documents that have been inactive for many years.

Authors were asked to keep active documents refreshed.

Some very old inactive WG drafts may be marked dead.

2. draft-ietf-bess-extended-evpn-optimized-ir

Presenter: Wen Lin, HPE

Scheduled time: 09:15–09:25

The presentation reviewed extensions to optimized ingress replication for EVPN, particularly for multihoming.

The goal is to reduce BUM replication overhead while avoiding implementation complexity associated with maintaining the original source IP address or ESI-label information during decapsulation and re-encapsulation at an assisted replicator.

The proposed procedure enables the assisted replicator to avoid replicating traffic toward leaves that are multihomed with the originating leaf. Replication toward those peers is instead handled using existing EVPN procedures.

A flag is introduced to signal the extended multihoming assisted-replication behavior.

The presenter reported an implementation and customer deployment and requested **WG Last Call**.

Discussion

Alexander (Sasha) Vainshtein, Ribbon Communications, asked whether assisted replication and the proposed extension could also be applied when EVPN uses an SRv6 underlay.

The presenter responded that the procedure could be extended for SRv6.

Jorge Rabadan, Nokia, noted that he and Wen Lin were already working on a proposal related to this question and expected to publish the work.

Another participant asked why point-to-multipoint LSP mechanisms could not be used instead of assisted replication and noted that EVPN solutions increasingly need to consider both SR-MPLS and SRv6 underlays.

Due to time constraints, further discussion was moved to the mailing list.

Next steps:

- Shepherd to review the draft.
- Progress toward WG Last Call.
- Continue SRv6 applicability discussion on the mailing list.

3. draft-ietf-bess-evpn-first-hop-security

Presenter: Krishna, Cisco

Scheduled time: 09:25–09:30

The presenter reviewed the motivation for EVPN First-Hop Security.

ARP/ND inspection and IP Source Guard rely on DHCP snooping information. With multihomed hosts, DHCP snooping state must be synchronized between multihoming PEs. Synchronization is also important when a host moves within the EVPN fabric.

The draft defines a DHCP synchronization route carrying information such as MAC address, IP address, DHCP binding state, lease time, and creation time.

The document is already a WG document and has authors from multiple vendors. The presenter stated that no technical issues remained open.

The authors requested **WG Last Call**.

Discussion

A participant noted ongoing work related to DHCPv6 Prefix Delegation and suggested investigating whether the synchronization information could be unified rather than defining too many separate routes.

The presenter agreed to follow up.

The chairs reminded authors to explicitly email the chairs when a document is ready for adoption or WG Last Call.

Next step: Authors to formally request WG Last Call.

4. draft-ietf-bess-evpn-umr-mobility

Presenter: Krishna, Cisco

Scheduled time: 09:30–09:35

The presenter reviewed the WG-adopted draft covering mobility procedures involving the EVPN Unknown MAC Route.

The document clarifies MAC mobility procedures, particularly in data-center interconnect scenarios, and adds text concerning symmetric IRB.

The presenter stated that the functionality is supported by multiple vendors.

The authors requested progression of the WG document.

No substantial mic discussion was recorded before moving to the next presentation.

5. draft-kriswamy-bess-evpn-perflow-df

Presenter: Krishna, Cisco

Scheduled time: 09:35–09:40

The presentation proposed a new **per-flow DF election algorithm** for EVPN multihoming.

Existing DF election algorithms primarily distribute traffic across VLANs or multicast groups. The proposal aims to provide load balancing for individual flows within a single VLAN and use a common DF mechanism for multicast, broadcast, and unknown-unicast traffic.

The proposal defines a new DF algorithm type carried in the existing DF Election Extended Community defined by RFC 8584.

For each packet, a hash is calculated and mapped against the ordered list of participating PEs to select the DF for that flow.

The draft does not mandate the hash fields, but all PEs participating in the Ethernet Segment must use consistent fields.

The presenter highlighted improved utilization of multihomed PEs and per-flow load distribution.

The authors requested **WG adoption**.

Next step: WG participants to review the new revision and provide comments.

6. draft-jot-bess-evpn-mcast-router-sync-00

Presenter: Jorge Rabadan, Nokia

Scheduled time: 09:40–09:50

The draft addresses synchronization of multicast-router state in EVPN networks.

Existing EVPN specifications provide mechanisms for synchronizing multicast host membership learned via IGMP/MLD across multihomed PEs. The missing case is where a multicast router rather than a host is multihomed.

The proposal introduces a **Multicast Router Discovery (MRD) route** to synchronize multicast-router information across PEs attached to the Ethernet Segment.

For PIM neighbors, additional PIM Hello information may be carried using a new sub-TLV in the Tunnel Encapsulation Attribute.

The proposal also reuses EVPN Route Types 7 and 8 for information triggered by PIM Join/Prune messages.

Three use cases were presented:

1. An external multicast querier multihomed into an EVPN broadcast domain.
2. EVPN Layer-3 multihoming with a multihomed PIM router.
3. OISM deployments with a PIM router multihomed into an Ethernet Segment.

Discussion

Mankamana : WG should carefully determine the **minimum PIM information** that needs to be synchronized. PIM has many Hello options, and unnecessarily synchronizing them all could create

significant complexity.

Prasad , Cisco, specifically questioned synchronization of the **PIM GenID**, noting that the GenID can change following a PIM process restart and asking why it must be synchronized.

Jorge Rabadan, Nokia, explained that without GenID synchronization, failover to a PE that did not receive the Hello directly might appear as a GenID change and trigger unnecessary PIM signaling.

The participants agreed that the exact PIM information that must be synchronized needs further discussion.

Next step: Continue discussion, particularly regarding minimum PIM state and GenID synchronization.

7. draft-li-bess-evpn-ead-multipath

Presenter: Zhiqiang Li

Scheduled time: 09:50–10:00

The draft described a potential issue when multiple PEs originate Ethernet A-D routes for the same Ethernet Segment using the same Route Distinguisher.

Under normal BGP best-path processing, the routes could become comparable, allowing a Route Reflector to select and propagate only one of them.

The presentation argued that suppression of other EAD routes could affect:

- ECMP/load balancing across multihomed PEs.
- Mass withdrawal after PE or segment failure.
- Availability of ES/EVI labels used for split-horizon processing.

The proposed solution is for BGP speakers to treat applicable EAD routes as multipaths and advertise, reflect, and install all valid EAD routes for an Ethernet Segment.

No new route type or capability is introduced.

The proposal also recommends configurable limits on the number of paths accepted per ESI as protection against excessive or malicious advertisements.

Discussion

Alexander (Sasha) Vainshtein, Ribbon Communications, questioned the underlying problem statement.

He pointed out that RFC 7432 specifies Route Distinguisher behavior designed to make the applicable EAD routes unique. If the RFC requirements and recommendations are followed, the routes should not become comparable and should therefore not be suppressed by normal RR best-path selection.

The chairs acknowledged that the unique-RD requirement/recommendation is central to determining whether the proposed problem actually occurs.

Next step: Continue the technical discussion on the mailing list, particularly whether the issue arises in RFC 7432-compliant deployments.

8. draft-many-bess-rfc9252-dual-sid

Presenter: Himanshu Shah, Ciena

Scheduled time: 10:00–10:10

The presentation addressed MPLS-to-SRv6 migration scenarios where both underlays coexist.

The stated motivation is to avoid duplicate service signaling during migration and thereby reduce control-plane overhead.

The proposal preserves the MPLS service label rather than making it Implicit NULL in certain RFC 9252 procedures and introduces a capability exchange so that peers can determine whether both sides support the new behavior.

The intention is to allow the ingress PE to choose either the MPLS or SRv6 underlay while maintaining interoperability with legacy MPLS PEs.

Discussion

Several participants raised concerns about the proposal.

Jorge Rabadan, Nokia, challenged the characterization of RFC 9252 behavior and noted that the RFC clearly specifies the existing service-label handling.

He also argued that although the proposal aims to reduce control-plane overhead, disabling or reducing use of transposition could reduce BGP update packing and therefore increase control-plane overhead.

Another participant raised concerns about IPv4-in-IPv6 next-hop behavior and the operational implications for migration from IPv4 BGP transport to IPv6 transport.

Alexander (Sasha) Vainshtein, Ribbon Communications, noted that some EVPN mechanisms have different applicability depending on whether the underlay is MPLS or SRv6. He used assisted ingress replication and point-to-multipoint delivery as examples and questioned whether a single advertisement can correctly represent both forwarding environments.

Due to time constraints, the chairs stopped further discussion.

Next step: Continue the technical discussion on the BESS mailing list.

9. draft-ietf-srv6ops-srv6-deployment

Presenter: Jakub Horn, Cisco Systems

Scheduled time: 10:10–10:25

The presentation described operational procedures used to migrate MPLS networks to SRv6.

The presenter described a phased approach:

1. Enable IPv6 routing and IPv6 loopbacks across the network.
2. Add Route Reflectors using IPv6 transport.
3. Enable SRv6 locators and dual functionality on selected PEs.
4. Advertise service routes through both MPLS and SRv6 control-plane paths.

5. Use BGP attributes such as Local Preference to select the preferred forwarding path.
6. Gradually migrate individual services and PEs.
7. Remove legacy forwarding after the migration is complete.

The presenter emphasized that the migration can be performed with fine granularity, potentially prefix-by-prefix or VPN-by-VPN.

Additional mechanisms such as SRv6-to-MPLS service gateways and other interworking procedures were also mentioned.

The presenter asked BESS for feedback on applicability of similar procedures to VXLAN environments.

Discussion

Jorge Rabadan, Nokia, said that the example presented in the slides was very clear and suggested adding such an example to the document.

He also recommended referencing BESS service-gateway specifications dealing with subjects such as gateway redundancy and propagation of route properties.

The presenter agreed.

Actions:

- Consider adding the migration example to the draft.
- Add references to relevant BESS service-gateway work.
- Continue discussion of VXLAN applicability.

10. draft-tantsura-bess-evpn-route-type-exp-use-00

Presenter: Jeff Tantsura, Nvidia

Scheduled time: 10:25–10:35

The draft proposes reserving a small range of EVPN route types for **Experimental Use**.

The motivation is to allow implementation and interoperability experimentation without squatting on unassigned EVPN route-type code points while waiting for RFC publication.

The proposal would update RFC 7432 and reserve a small range for experimentation according to the principles of RFC 8126.

The authors emphasized that these code points are intended for experiments, not production deployment.

Discussion

Jorge Rabadan, Nokia, asked how existing Route Reflectors that have not been upgraded would handle unknown experimental route types.

He noted that a Route Reflector may not understand basic information required to validate or compare such routes, such as the route key and associated structure.

He referenced previous BESS experience with vendor-specific route types, where minimum semantics were defined so Route Reflectors could safely process the NLRI.

Jeffrey Haas, HPE, raised additional security and operational concerns. If experimental code points are not sufficiently constrained, an experiment could escape its intended environment or collide with another experiment using the same value.

He suggested that the document consider:

- Explicit enablement.
- Limiting the blast radius.
- Capability negotiation.
- Potentially carrying version information so participants know which version of an experiment is in use.

Ketan Talaulikar, Cisco, speaking as a co-author, reiterated that the purpose of the draft is simply to provide legitimate code-point space for experimentation and avoid code-point squatting. It is not intended to define deployment procedures for production networks.

Keyur Patel, Arrcus, echoed concerns raised by Jorge and Jeff and recommended that the document contain very strict guidelines for implementers so experimental NLRIs do not inadvertently appear in production environments.

The presenter agreed that the comments were valid and requested further WG input.

Next steps:

- WG to review the proposed experimental range.
- Authors to strengthen operational/security guidance.
- Continue discussion of Route Reflector behavior, containment, capability negotiation, and potential versioning.

11. draft-smn-idr-inter-domain-ibgp

Presenter: Moshiko Nayman

Scheduled time: 10:35–10:45

The draft had previously been discussed in IDR and was brought to BESS to determine whether BESS is a more appropriate home for the work.

The document describes compartmentalizing a large service-provider network into multiple domains while retaining a single AS number.

Domain Border Routers connect the domains and selectively advertise routes between them. Depending on the deployment model, next-hop-self behavior may be required.

The presenter stated that the approach does not require new BGP protocol extensions and has been deployed operationally for approximately two years.

The work describes behavior similar to inter-AS VPN Options A, B, and C, but within a single AS.

Discussion

Jeffrey Haas, HPE, noted that IDR already has an adopted document covering the normative BGP next-hop-self behavior at Route Reflectors.

He observed that once the core BGP mechanism is handled in IDR, much of the remaining work in this draft concerns VPN deployment scenarios, which may make BESS the more appropriate WG.

Gunter Van de Velde, Nokia / Routing Area Director, noted that the document primarily describes inter-AS-style VPN options and deployment guidance rather than new BGP protocol behavior. He observed that this weakens the case for handling the complete document in IDR under the current IDR charter.

A participant requested that, if BESS decides to take the document, it be handled in an expedited manner because the work has already been under discussion for approximately two and a half years and had progressed through IDR adoption consideration.

The chairs agreed that this was a reasonable request.

Next step: Continue discussion on the mailing list and determine the appropriate WG ownership and processing path.