ACP-WG I / IP01

15 July 2014

International Civil Aviation Organization

INFORMATION PAPER

AERONAUTICAL COMMUNICATIONS PANEL (ACP)

SEVENTEENTH MEETING OF WORKING GROUP - I

Montreal, Canada 15 – 16 July 2014

Agenda Item: Enhancements to MIPv6 defined in DOC 9896

LISP – A Multi-Homing and Mobility Solution for ATN using IPv6

(Presented by Bernhard Haindl)

SUMMARY

This paper introduces LISP (Locator/Identifier Separation Protocol) as novel approach in order to enhance the ATN/IPS-based network layer. It focuses on usage of LISP technology in a multi-homing environment in order to achieve mobility and high availability for safety critical communication over aeronautical networks between aircraft and ground infrastructures. We show how the LISP mechanisms can be used to apply fast convergence, provide session continuity, enable secure dynamic location updates and enable a make-before break handover in multilink environment. Finally, we discuss how the LISP mobility system can be enhanced to include application type or QoS specific information into the data link selection.

LISP – A Multi-Homing and Mobility Solution for ATN using IPv6

LISP Basic Elements

LISP Mapping System

The LISP mapping system separates identity ('who') from the location ('where') of an IP device. At the border between EID (End system Identifier) and RLOC (Routing Locator) space, LISP capable routers perform encapsulation and de-capsulation of IP traffic traveling across. Thereby LISP dynamically establishes unidirectional tunnels without explicit configuration of tunnel endpoints.

The core functionality of LISP is provided by the mapping system together with the LISP border routers. Neither the IP end systems located in the EID space nor the IP routers in the RLOC space need to be changed in order to benefit from LISP. In general, LISP scales much better then IP destination based routing. Typically, in order to deliver traffic to a certain location a corresponding entry identifying the destination has to be entered into a router's routing table. Whenever a new network appears at a router the IP routing system has to push that information to adjacent routers. Such operation takes some time and may result in network instability during the time the topology changes. This effect is worsening with a growing size of the routing tables.

LISP utilizes a pull model. Whenever there is a need to communicate to an EID, the mapping system is queried in order to provide the actual RLOC. This is very similar to DNS, which holds billions of symbolic names in a distributed database. Experience shows that serving the DNS customer with a pull model is much more scalable. LISP basic elements are depicted in Figure 1.

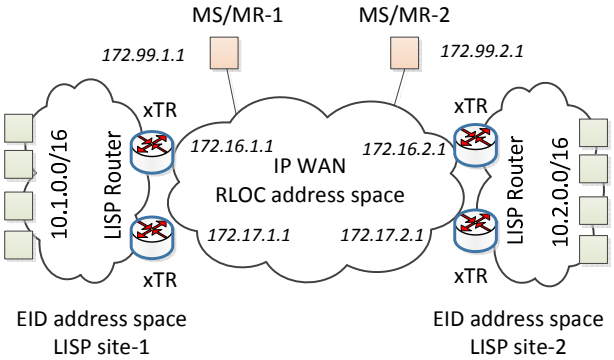

Figure 1. LISP Basic Elements

LISP Control Messages

As soon as a LISP router knows about an EID and has access to it, it registers this EID at the mapping server MS (time t_1 in Figure 2) with the corresponding RLOC(s). Map-Registration messages are transported in UDP and

integrity-protected by security credentials known only to the LISP routers of a given site and the mapping server.

MAP-Registration messages are acknowledged by the MS with Map-Notify (t_2). Additionally Map-Notify messages are used by LISP routers to discover other LISP routers of the same LISP site. The RLOCs could be a static or a dynamic IP addresses given from or assigned by the ISP(s).

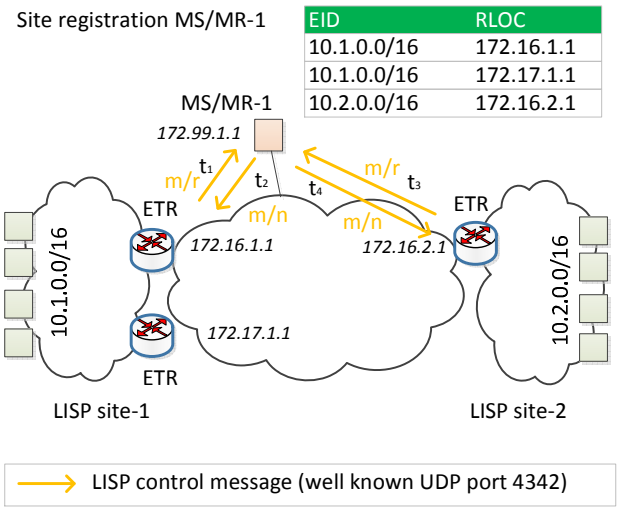

Figure 2. LISP Map-Registration

ETRs (Egress Tunnel Router) responsible for a given site have the authority about the currently available RLOCs. In analogy to DNS, the ETRs administer their zone file and give an authorized answer, whereas, the MS/MRs are just a redirection to find a responsible ETR. The mapping between RLOC and EID is stored in the mapping database of the ETR.

Let us consider what happens if traffic for a foreign EID appears at Site 1 (t_0 in Figure 3). The ITR will look up its mapping cache. In case of the first IP datagram for that destination, the mapping cache will be empty and the ITR (Ingress Tunnel Router) will query its mapping resolver (MR) via Map-Request for that EID (t_1). The Map-Request will be passed on to one ETR responsible for that EID based on the registration information stored in the MR. The up-to-date knowledge about the current situation is stored in the corresponding ETRs mapping database. The ETR will send a Map-Reply to the asking ITR containing all RLOCs which can be used to reach the EID (t_4). The result is stored in the mapping-cache of the asking ITR.

This example considers only one RLOC for EID Site 2. If there were more RLOCs at site 2, the ETR will respond with priority and weight of every RLOC in the Map-Reply. A priority value is used by the ITR to decide which RLOC to use – the lower the value the higher the priority. In equal priority configurations the weight is used in order to perform load-balancing among the RLOCs.

For instance, weight of 50 means equal round robin among all RLOCs. Therefore the ETR of site 2 is able to control incoming traffic using priority and weight.

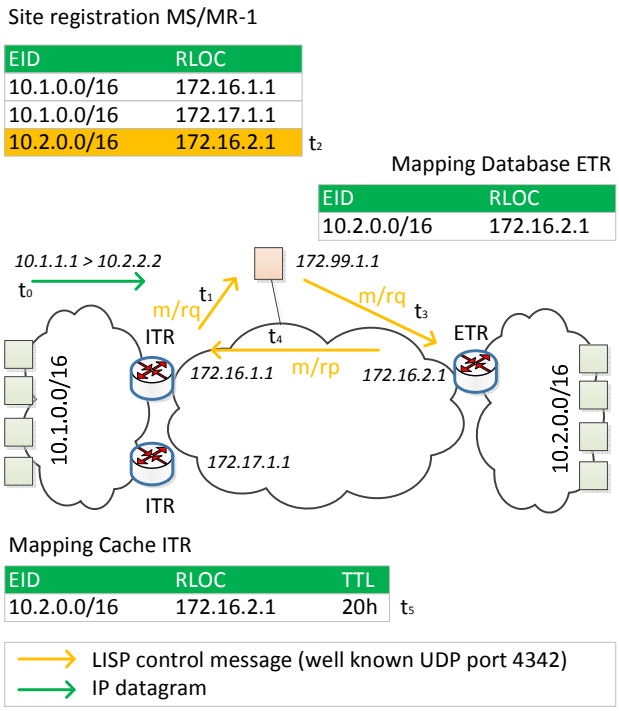

Figure 3. LISP Map Request

All control messages are UDP based and not integrity-protected. A nonce is used to overcome spoofing attacks from attacker's off-path. Missing Map-Replies can only be detected via timeouts by the ITR.

An ITR may select another MR if the primary MR is offline. The MR always forwards the request to the ETR last heard based on the registration procedure. The Map-Request is sent in a special encapsulated format to ease the computation at the mapping server. After finding a responsible ETR in the registration the mapping resolver can easily pass on the message.

Load balancing of traffic sent from an ITR to different RLOCs of a given LISP site is based on hash values. The ITR produces a hash including IP source and IP destination EID addresses or the 5-tupel also including source and destination ports (UDP, TCP) and IP protocol-type. This hash value is used as index to one of the RLOCs. So traffic of a given source-destination address pair (or flow) will always address the same RLOC. After the first IP datagram locates a corresponding RLOC (t_0-t_4), subsequent IP datagrams (t_6 as in Figure 4) find an entry in the mapping-cache of the ITR. In addition, the IP datagrams are encapsulated by the ITR, transported across the IP-WAN (t_7) and de-capsulated by the ETR at the destination site (t_8).

Every Map-Reply contains a TTL (time-to-live) field for a certain EID. As long as the TTL is not zero, the ITR will use the mapping-cache entry without re-querying the mapping system for that EID. Traffic for that EID is now sent without further control messages as long as the RLOC is valid and available. In a worst case scenario RLOC changes at a LISP site can lead to "black-holes" for a reasonable long time. The usage of so called "remote probing Map-Requests/Replies" for entries found in the mapping-cache of an ITR may overcome this problem. Such probes are directly exchanged between LISP routers without mediation of the LISP mapping system.

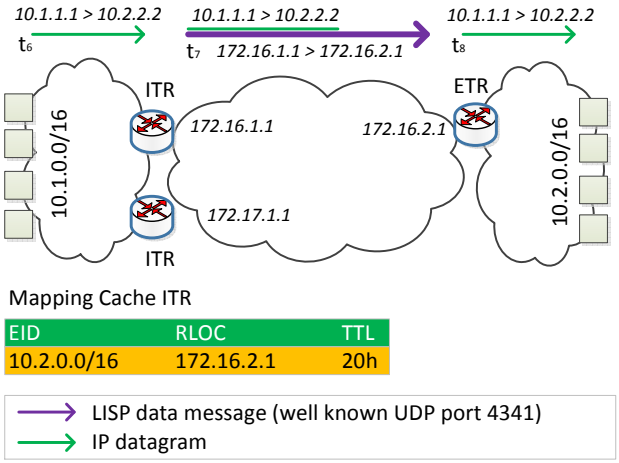

Figure 4. LISP Dynamic Tunnel

Requirements arising from mission critical communication may not allow delays in an initial call setup. LISP requires traffic that either triggers establishing a LISP tunnel (the first IP datagrams will get lost until the mapping-cache is filled) or keeps a connection alive. Such base traffic is typically required by mission critical applications in order to test reachability of remote peers (e.g. keep-alive or heart beat messages) and thus mostly avoids any delay caused by tunnel establishment.

LISP Security Features

Air Traffic Management (ATM) is a safety-critical domain where great care must be taken to ensure security and timeliness of the information exchange. This is a non-trivial design challenge, given that the producers and consumers of information, as well as the information itself, frequently reside in different domains, necessitating some form of cross domain solution. In order to deploy ATM communication networks relying on LISP we introduce two approaches that secure network traffic between LISP sites, a base Virtual Private Network (VPN) configuration and a combination of LISP and Group Encrypted Transport (GET) VPN.

LISP provides an inherent capability to provide base VPN functionality among LISP sites as EID addresses are not routed in the RLOC address space (i.e. the IP WAN). In addition, registration of EIDs in the mapping system needs knowledge of security credentials for a LISP border router. Therefore non-LISP sites or any other devices in the RLOC space cannot communicate with EID sites as the routing system has no information how to route traffic to an EID address, as illustrated in Figure 5.

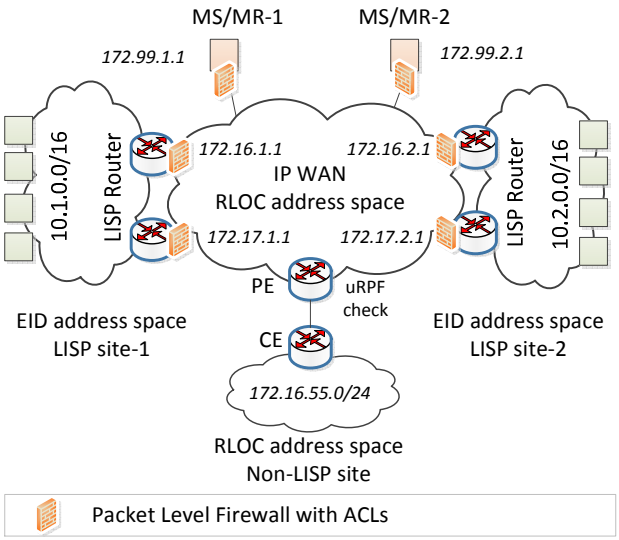

Figure 5. LISP Basic VPN (ACL)

In a default configuration, LISP data messages and control messages are sent in clear-text over the IP WAN allowing a potential attacker to gain knowledge about the identity of EIDs behind a LISP site and what RLOCs can be used to get to that EID site. An attacker may direct a Denial of Service (DoS) attack against this base VPN established by LISP even if he is not able to communicate from a non-LISP site with an EID device. In order to prevent such attacks access lists (ACL) on the incoming interface of any LISP router should be used allowing only RLOCs and MS/MRs and LISP data/control message flows. Combining ACL with Unicast Reverse Path Forwarding (uRPF) checks towards all non-LISP site networks overcomes the problem of source address spoofing where an attacker spoofs valid RLOC addresses.

In the case of having full control over the IP WAN or a trusted relation with the network provider then the base VPN functionality described here might be sufficient. However, it is understood that maintaining ACLs increases administrative overhead and uRPF check requires coordination with the IP WAN service provider. These measures therefore reduce the plug and play capabilities of LISP especially for deployable systems and therefore alternate approaches should be investigated. Combining LISP with GETVPN provides such an alternate approach.

The base idea of GETVPN is crypto-graphical protection of packets transported over a network (see [6]). That protection can offer integrity/authenticity for such packets and optionally privacy by means of encryption. GETVPN is based on IPsec principles but uses a group key for all legitimate members in order to protect traffic. Traditional IPsec technology is based on point-to-point security associations and requires a separate key for each of such association which leads to scalability issues. Joining a GETVPN group and distribution of key material is accomplished by secure communication between GETVPN group members (GM) and a GETVPN key server (KS).

Figure 6 shows how GETVPN can be used in a LISP environment. A GETVPN key server is placed in the RLOC space. All LISP routers and mapping servers are acting as GETVPN group members at the RLOC side. By identifying packet flows that need security processing and defining a policy for these flows - known as crypto map – it is possible to select traffic to be protected on both, the RLOC interface of the LISP routers and the interfaces of the mapping servers. The easiest setup of such a policy is to protect all LISP packets (control and

data) in the same way by using the actual group key. One task of the key server is periodic rekeying hence a distribution of new keys is necessary just in time before the actual key times out.

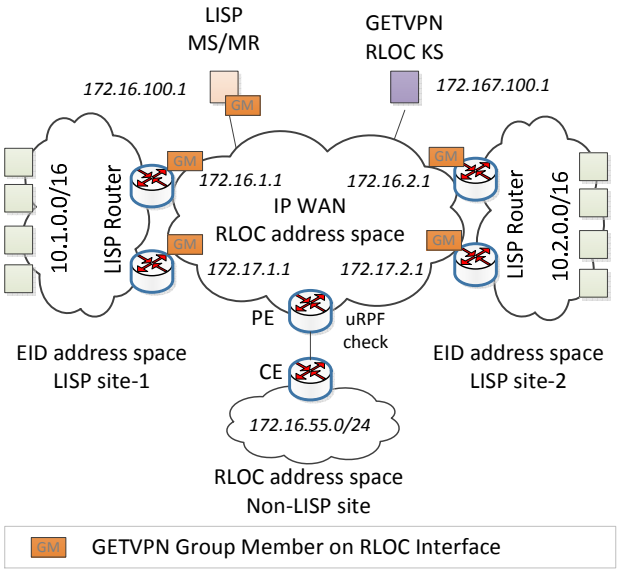

Figure 6. LISP GETVPN

This approach solves two security issues: LISP control messages can be integrity-protected by GETVPN hence you need no additional care on maintaining ACLs. LISP data messages carrying user traffic between LISP sites can be privacy-protected. EID addresses are then not sent in clear-text when passing the IP-WAN which might be an issue if pure GETVPN was used without LISP technology.

More sophisticated scenarios can be handled in a similar manner. For instance multiplexing of several VPNs over the same IP WAN with different IPsec policies for each individual VPN is possible. However, the important point is, that combining the stateless tunneling technique of LISP network together with the stateless IPsec technology of GETVPN, scales a magnitude better than any traditional methods like MPLS-VPN and traditional point-to-point IPsec.

Even with GETVPN a security issue remains, which is protection against DoS attacks. Especially key servers need such a protection as they are not covered by LISP separation technique - they require direct connection to the IP WAN. Generally speaking, protection against DoS cannot be solved by a single measure moreover, it requires special observation and monitoring capabilities to react to such a threats. Typical examples are business critical webservers residing in the public Internet. More information about usage of LISP for mission critical communication can be found in [7].

Multilink Concept in ATN

In the joint Action Plan 17 (AP 17) the concepts of ATM have been analyzed from an operational point of view and the expected technical requirements have been formulated, also for services which are not yet deployed but are expected to be deployed in the future. The results in the COCR provide information for all operational services with respect to their periodicity, volume and technical requirements.

The technology investigations in AP 17 led to the following three proposals for new air/ground data link communications system:

1. L-band Digital Aeronautical Communication System (LDACS) as a future ground based communication system;

2. A future satellite based communication system;

3. Aeronautical Mobile Airport Communication System (AeroMACS) as a future specific system for the surface communications in airports with high density traffic;

While for AeroMACS (aeronautical mobile airport communications system), a specific existing standard has already been identified, the other two technology developments require further activities to finalize the selection and standardization of the proposed systems. The new link technologies will complement the existing ones, like ACARS, VHF Digital Model (VDL) Mode-2, and will coexist with them for a long time period to ensure a realistic transition. Figure 7 shows a ground and airborne infrastructure, as well as the different A/G data links of a possible Future Communication Infrastructure (FCI).

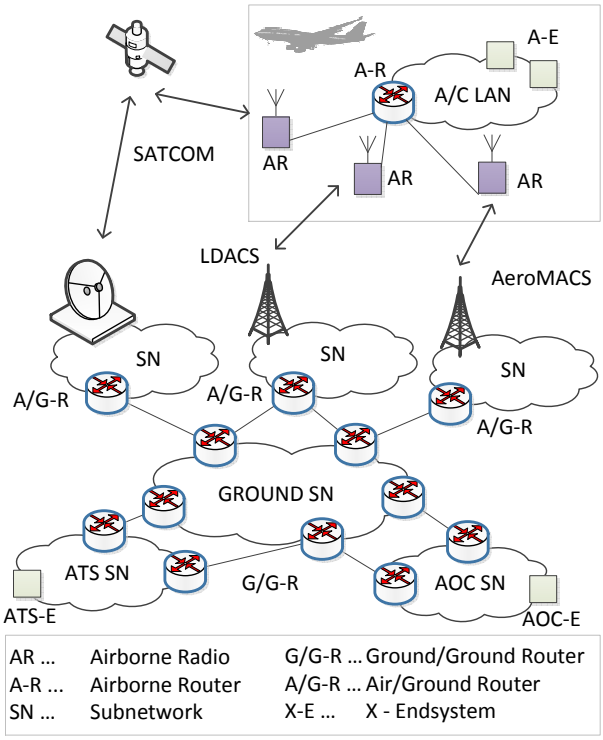

Figure 7. Future Communication Infrastructure

The currently defined new operating concepts, anticipate that multiple new datalink systems will be required to meet availability and quality requirements. It is assumed that only these future links, or a combination of them, will be able to provide adequate QoS support for services of the future operational concept. Such a combination of independent A/G datalinks, with the objective to provide robust and highly available data link connectivity, was identified as new functionality within the FCI and is called “Multlink” concept. This concept addresses the selection of the most appropriate radio link for a given data link service and defines the requirements for handovers between radio links according to various operational contexts. Currently no IPS Multilink architecture exists, which provides solutions for the four main areas:

- Security
- Addressing among A/C and address resolution
- QoS management for mixture of ATM-over-IP traffic comprising categories with different criticality and priorities
- Multi-homing and handover

The current ATN/IPS protocols and solutions defined in ICAO Doc.9896 provide QoS, security and mobility solutions for a future aeronautical data network system. But it doesn't consider a "Multlink" concept or Multi-homing environment.

LISP has the capability to provide a transparent multi-homing solution for the end systems which allows load sharing between the different radio technologies dependent on the available QoS. It also solves the network mobility problem of the aircraft network and together with GETVPN it provides a maintainable security solution.

LISP and ATN Use Cases

It is envisaged that future ATN system will comprise of heterogeneous systems (i.e. short-range and long-range terrestrial as well as satellite access technologies). Hence it would be important to monitor and control these various systems efficiently and also provide means to handover active sessions between these different radio technologies

Currently, the different radio access networks work independently and have limited or in most cases no interaction amongst themselves. To support efficient control and seamless vertical handovers, the airborne radios (AR) have to communicate with the airborne router (A-R) to forward availability and QoS parameters. Figure 8 shows a possible FCI, where the A-R and the G/G router of the ATS or AOC ground sub networks (G/G-R) are LISP capable router, performing the encapsulation/de-capsulation and IP address mapping (see Figure 1). This allows another kind of mobility which supports seamless communication even without any interruption in case of local failures.

The example in Figure 8 shows an aircraft in the continental or domestic airspace that is equipped with the future SATCOM and LDACS radio. Each interface of the A-R acts as RLOC to one or more ground radio sites of the LDACS or SATCOM Sub Network. Since more than one radio link is active, we have a multi-homing scenario which can be used either for load-balancing or for selecting the best radio link based on the operational flight domain and the given data link service. The appropriate radio link will be determined by the radio parameters received from AR. In our scenario we have given by usage of appropriate LISP router configuration LDACS priority 1 and SATCOM priority 2. In the current situation only SATCOM is available therefore only this RLOC is registered in the mapping system.

Communication from other LISP sites will show this RLOC in the mapping cache. As soon as LDACS becomes available, the LISP mobility system can be used to switch session flow over to LDACS and the mapping caches of all LISP sites will adapt to the new LDACS RLOC. For the session this vertical handover is a move without break.

Site registration MS/MR-1

EID	RLOC	PRIO
10.1.0.0/16	172.15.1.1	
10.1.0.0/16	172.16.1.1	
10.2.0.0/16	172.19.1.1	
10.2.0.0/16	172.20.1.1	
10.3.3.0/28	172.17.1.1	2
10.3.3.0/28	172.18.1.1	1

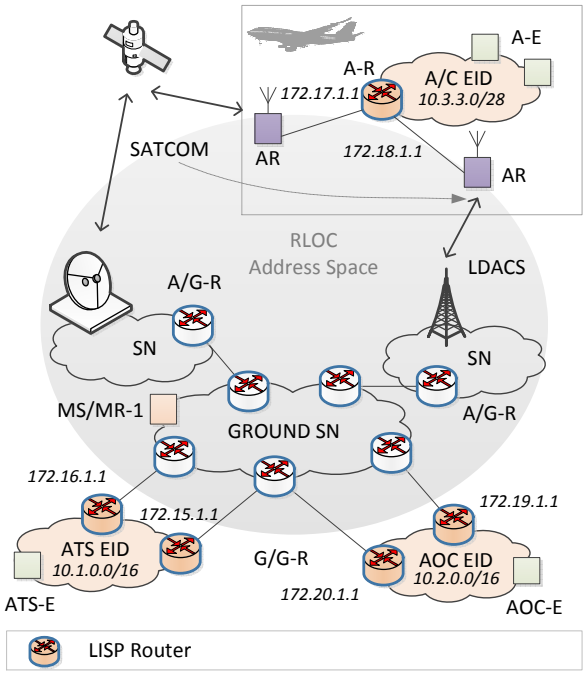

Figure 8. Mobility concept with LISP

If the AR of the LDACS system reports loss of redundancy the A-R could adjust the priority for the RLOCs in the mapping database (e.g. by using the EEM) then it would be possible to use the LISP mobility system to switch back to more reliable RLOC. That means even if the aircraft has connectivity via several link technologies, a constrained based move without a break can be done.

If the aircraft is considered an operational flight domain, where messages of different transactions can be sent over different radio links, then LISP can provide load balancing based on the application type. In addition, outgoing traffic engineering based on the application type may be configured on the A-R router. Figure 9 illustrates such a scenario.

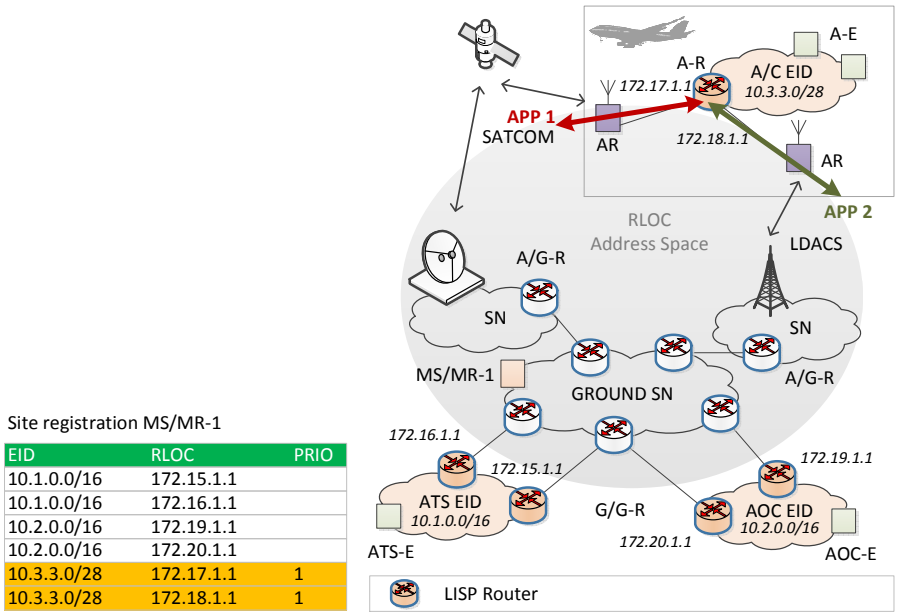

Site registration MS/MR-1

EID	RLOC	PRIO
10.1.0.0/16	172.15.1.1	
10.1.0.0/16	172.16.1.1	
10.2.0.0/16	172.19.1.1	
10.2.0.0/16	172.20.1.1	
10.3.3.0/28	172.17.1.1	1
10.3.3.0/28	172.18.1.1	1

Figure 9. LISP Load Balancing Scenario

Conclusion

The usage of LISP in the network transport system has several benefits compared to traditional methods.

First, it provides a clear demarcation point between user sites and WAN provider network for routing and security. That makes distribution of responsibilities and reduces the control complexity between the user and WAN provider.
Second it allows enhancement of the network concerning new sites connected to the WAN provider or multi-homing of existing sites or even changing or adding of WAN providers without any disturbance of existing ongoing communication. Such changes are automatically detected and integrated in a seamless way.

Third, network security especially integrity and authentication of messages as well as optional encryption can easily be added if necessary (e.g. in case the WAN network is located in a shared or non-trusted security domain) in a scalable and seamless manner by usage of GETVPN technology available in commercial off the shelf IP routers.
Fourth, multi-homing and mobility are an inherent functionality of LISP hence support of multilink concept for communication to mobile aircrafts is available from Day-One by using just COTS LISP routers. With appropriate configuration of these COTS LISP routers “Move before Break”, “Seamless communication” and “Constraint-based Routing depending on available QoS” implementations are possible.

The LISP technology and associated concepts as outlined in this contribution can be easily deployed at an individual Air Navigation Service Provider level (ANSP) or a group of ANSPs to support communication up to the aircraft in order to facilitate seamlessly shared operation and management of the air traffic management network infrastructure.