A PROPOSAL FOR COALITION NETWORKING IN DYNAMIC OPERATIONAL ENVIRONMENTS

Randall Atkinson
Extreme Networks
NC, USA

Manish Lad
University College London
London, UK

Saleem Bhatti
University of St. Andrews
St. Andrews, UK

Stephen Hailes
University College London
London, UK

ABSTRACT

At present, military communications within battlefields are very restricted, both by policy and due to technology limitations. In Southwest Asia today, there are needlessly long and complex communications paths, often involving multiple relays and use of constrained-bandwidth MILSATCOM back-haul outside the theatre, when nearby forces could communicate directly via existing interoperable radios. This is a current problem for NATO and Coalition forces. The current Internet Protocol suite lacks core support for mobility, scalable support for multi-homed nodes, and does not provide the capabilities needed for optimal communications in forward operating areas.

We propose a coalition-based, multi-homed approach leveraging both local-area and wide-area connectivity, improving both the flexibility and robustness of communication, without conflicting with the security policy of sensitive communication. The Coalition Peering Domain (CPD), is a distributed, self-configuring architecture that supports the secure, collaborative networking relationships needed to provide this flexibility and robustness. The CPD facilitates the inter-connection of cooperating, but administratively separate, network segments. The CPD exploits multi-homed and multi-path communication to better-utilise all available connectivity. The Identifier-Locator Network Protocol (ILNP) provides native support for improved scalability in multi-homing and mobility, while easing use of network layer security and allowing inter-operation across different administrative domains. Our approach is compatible with current work in Mobile Ad-Hoc Networking (MANET).

ILNP has excellent compatibility with IPv6: existing IPv6 backbone networks do not require any modification to carry ILNP traffic natively. There are practical, realistic and deployable engineering solutions to realise the CPD and ILNP within the framework of IPv6.

I. INTRODUCTION

Battlefield environments today would benefit greatly from increased flexibility in communications, especially where there are coalition forces using administratively distinct networks that need to work together. We present here a scenario that highlights the problem space, captures the salient high-level requirements (based in part on the current operational environment in Southwest Asia) and outlines a possible solution based on our ongoing work.

In our scenario, there are coalition forces from many different countries. Further, each country may have units deployed from different military services. Typically, a survey of the communications infrastructure of such a collection of forces is likley to show a huge heterogeneity in technology, including a mixture of short-range, long-range, and back-haul communication. Most of the communication is radio-based: long-distance wired (or fibred) communications infrastructure is limited or non-existent in many potential operating theatres, including Southwest Asia.

In battlefield scenarios, continuous operation in the face of mobility is an imperative, of course. The heterogeneity of the underlying network is typified by communication between highly mobile units (for example, infantry with low powered, short-range devices or cavalry in armored vehicles) operating in the vicinity of mobile nomadic units that may provide back-haul and relay services (e.g., mobile command vehicles with higher powered equipment and a wider range of operating frequencies, potentially from HF through VHF/UHF to SATCOM).

To compound this heterogeneity, even within a single country's forces, different members of military coalition forces operate and maintain their own networks, each built on potentially different technologies, thus involving the use of multiple different radio systems and infrastructure. Even if the interoperability vision of the Joint Tactical Radio System (JTRS) programme is achieved, there are likely to be other coalition members who are not using interoperable JTRS radios. Enabling secure communications between coalition forces often requires the use of MILSATCOM links (which normally are bandwidth-constrained) between the operating theatre (e.g., Southwest Asia) and some other fixed location (e.g., Western Europe) where communications can be relayed between different countries or different services participating in coalition operations.

Communications back-haul outside the theatre is often expensive, complex, and creates a potential critical single point of failure into the deployed communications network. There would be significant gains if operational units were

able to auto-configure securely their communication networks according to predefined policies, when operating in proximity of each other. This would allow coalition forces to co-ordinate their activities directly, within theatre, making use of multiple local relay systems between different radio technologies simultaneously. This increases the robustness of the communications network, having the potential to improve traffic load distribution across the network and to make more efficient use of existing radio capacity. Of course, this does not preclude back-haul links outside the theatre for inter-area communications, if required.

So, our key high-level requirements are for a network architecture that:

1) Allows different coalition forces to communicate directly, across multiple radio systems in the field, without any back-haul through remote relays, using only local relays operated by the units themselves.
2) Enables coalition forces to discover each other and to auto-configure relevant network parameters (addressing, routing etc.) to allow interoperation with minimal manual intervention.
3) Is robust against network faults, including to loss of nodes, not just link and transmission problems.
4) Provides efficency in use of available radio capacity through the use of network multi-homing and communications using multiple network paths.
5) Supports application of policy constraints to control the operation of the communication network.
6) Can easily incorporate security policy and security protocols that are currently in use within existing tactical and strategic networks.

Our proposed architecture, the *Coalition Peering Domain (CPD)*, has implications for the user plane, the control plane and the management plane of comunications. For the user plane, we are concerned with packet handling: addressing, routing and forwarding in the core and edge network. In the control plane, we are concerned with the protocols and handshakes required to establish, configure and maintain a CPD. In the management plane, we are concerned with issue of overall policy and how it would affect both the user plane and control plane. In this paper, we focus on the user plane and control plane. However, we believe that the management plane could be implemented by applying exsiting policy systems within the CPD context, and in our design of the CPD we do not place any constraints on policy mechanisms that might be used.

In the remainder of this paper we outline a proposal, based on our ongoing work, that can meet these requirements. In Section II, we describe the *Coaltion Peering Domain (CPD)*, a network entity that is formed when adminsitratively distinct network domains (e.g., distinct coalition force units and distinct coalition forces themselves) merge or *peer* dynamically (on demand) at the network level to share network resources. In our discussion, we distinguish between the *coalition* of armed force refering to forces on the battlefield and the Coalition Peering Domain (CPD) that refers to the collaborative inter-networking relationships formed between those coalition forces' network-capable devices. In Section III, we list the key challenges in the current IPv4 and IPv6 network architecture that make it difficult to realise the CPD. In Section IV, we describe the features of the *Identifier Locator Network Protocol (ILNP)*, an IPv6-compatible network protocol that provides underlying networking functions to help address these deficiencies directly. We then analyse security considerations in Section V. In Section VI we discus related work, and we conclude in Section VII including a brief statement of the next steps in our ongoing work.

II. The Coalition Peering Domain

In this section we present a new dynamic network entity that enables networks to merge or to *peer* (in the network sense) together in novel ways. In traditional peering, the rigid administrative and policy boundaries prevent fully-effective resource allocation. However, in our new model, several different networks combine to form a single *Coalition Peering Domain (CPD)* [1] that more effectively shares network resources, while allowing each administrative entity to retain control of its own network nodes. We present initially the abstract structure of the CPD and then give an example of how it may be applied.

A. Outline CPD architecture

Fig. 1 illustrates a number of collaborative relationships or 'local peering agreements' between pairs of devices within proximity of each other. These bi-lateral peerings might be either simple links interconnecting different pairs of devices or more complicated associations controlled by local policies specified by the device owners. As the numbers of such local peering agreements begin to increase and different devices form more complex interconnection topologies, we refer to the formation of a *Coalition Peering Domain (CPD)*.

Each *Coalition Member (CM)* represents either an individual device or a local-area/personal-area network. Local peering agreements are negotiated and maintained directly between CMs. The agreements enable each CM to declare the resources that they are willing to offer, to accept or reject the resources that are being offered by potential peer-CMs, and to amend any of the peering parameters after the local peering agreement has been formed.

Coalition members who have wide-area connectivity (or more generically, connectivity outside the CPD) act as

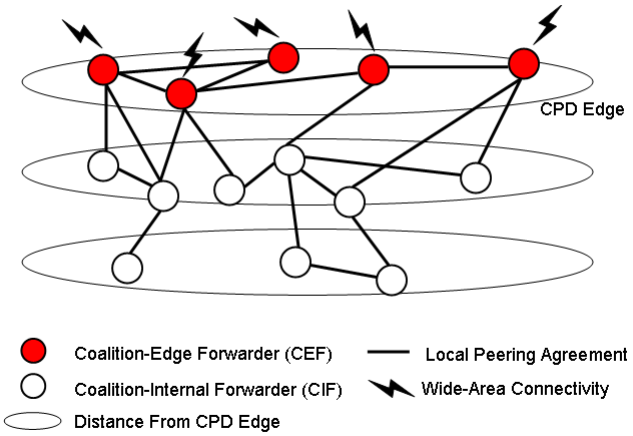

Fig. 1. CPD Architecture

Coalition-Edge Forwarders (CEF). These CEFs are the ingress–egress points for the CPD, allocating some proportion of their exterior bandwidth for this purpose. CEFs collectively form the logical border or *edge* of the CPD, and their aggregate external connectivity provides the CPD with a higher *potential* exterior data rate than any individual CM is able to achieve. To utilise this efficiently for CPD-egress traffic, CEFs forward a proportion of outgoing packets on their own CPD-egress links, but forward the remaining egress traffic by 'spraying' (distributing) the traffic across the CPD edge, via their local peering agreement links to neighbouring CEFs. The neighbouring CEFs do the same. Such distribution of outgoing traffic across multiple CEFs enables the aggregation of CEF egress links, providing higher CPD-egress data rates and greater robustness in connectivity through multiple connections.

Coalition members who do not have connectivity outside the CPD, or who choose not to make their exterior capacity available to other CMs, act as *Coalition-Internal Forwarders (CIFs).* CIFs contribute to the CPD by forwarding traffic to other CMs, including to/from CEFs. They need simply to forward CPD-outbound traffic by directing it towards their nearest CEF for CPD egress. Of course, CIFs may also use mechanisms for load balancing and also take responsibility for spraying directly to multiple CEFs, depending on the physical connectivity of the CPD and any applied policy constraints.

When nodes seek to join an exisiting CPD, or if a node wishes to initiate a new CPD, there is a secure handshake to establish trust. This handshake uses existing security protocols and security mechanisms including public-key exchanges for authentication. The result of the handshake is the exchange of sufficient information (including network parameters) to allow network layer connectivity for the new node as the node now becomes a CM. Upper-layer protocols (e.g., transport protocols and applications) are excluded from this handshake, deliberately, so as to provide isolation of operation for security and simplicity between network operation and the operation of upper-level protocols (both transport and application levels).

When two networks or two CPDs wish to communicate between themselves, one or more nodes initiate the merging of the two networks or CPDs by forming a new CPD which is the union of the two networks. Of course, this approach will work for many networks wishing to peer, not just for a pair of networks.

B. Example of CPD in the battlefield

Here we present an example of how the CPD may be applied within a possible battlefield scenario.

Consider an individual armored unit that is part of a coalition force, in which all vehicles are equipped with common short-range radios. The command vehicle within the unit might also have other longer-range radios for communication back to base camp or outside the theatre. In this model, each vehicle may act as a CIF and the command vehicle carrying the longer-range radios may act as a CEF. Not only can each vehicle forward traffic for other vehicles (CIFs) but ti can also forward traffic to and from the long-range radio (CEF) for communication directly back to headquarters or even outside the theatre (e.g., if SATCOM is available on the command vehicle). If several such units were deployed together, the separate unit-networks could form a CPD, allowing soldiers from all units to communicate with one another, and allowing soldiers from any unit to make use of all longer-range radio links (CEFs).

A real and immediate threat comes from radio jamming equipment that may be operated by opposition forces. The robustness in connectivity that is provided by the CPD helps to alleviate this because even if some radio frequency ranges are disrupted, thus disabling some CEF uplinks, other CEF uplinks may remain unaffected (e.g., different frequency radios, higher-powered wireless communications or wired connections), and so allow the CPD as a whole to maintain wide-area connectivity. Indeed, the CPD architecture can be used in complement to existing physical-layer and link-layer anti-jamming techniques.

Similarly, if one of the longer-range radios is broken or destroyed, this will be discovered automatically and the CPD will reconfigure dynamically to reflect the loss, thereby ensuring that all the soldiers still have long range connectivity from the remaining longer-range radios. As long as multiple longer-range radios are functioning, all are used simultaneously to give load distribution and to increase the robustness of the warfighters' network.

Of course, the vehicles in the above example could be replaced by any combination of people, tanks, ships, boats, armoured personnel carriers, cars, trucks, or aircraft.

C. CPD Control and Signalling

We present here a summary of our work so far on the control plane of the CPD, to allow end-systems to join and leave a CPD.

As part of local peering agreement formation, a simple three-way handshake is employed between potentially peering CMs, to agree both CPD membership and peering parameters.

A coalition member CM_A that wishes to form a local peering agreement with a neighbour CM_B, transmits to CM_B a *CPD Peering Request (CPD_PREQ)*. This CPD_PREQ contains the security and resource parameters that CM_A is willing to offer for peering. If CM_B is not willing to agree to a local peering agreement formation, it may either transmit a *CPD Peering Response (CPD_PRESP)* rejecting the request, or simply ignore the CPD_PREQ and allow CM_A to timeout awaiting response.

However, if CM_B wishes to agree to the formation of a local peering agreement, it then transmits back a CPD_PRESP accepting the request. Most importantly, this response message contains a number of parameters to determine the scope and operation of the local peering agreement:

- **Proposed CPD_ID:** an identifier for the CPD within which the local peering agreement is to operate. This CPD_ID may be either a newly generated identifier (thus forming a new CPD), or the identifier for a CPD within which CM_B already is a member (thus adding the new local peering agreement to the existing CPD).
- **Resource Parameters:** a set of routing, capacity and quality of service metrics that CM_B is willing to offer in return.
- **CPD Edge Proximity:** a metric to indicate CM_B's relative distance from the CPD-edge. For CIFs, this provides a means to determine the best routes for CPD-outbound traffic. For CEFs, this provides a means for CEFs to determine peering CMs to which they may spray CPD-outbound traffic.
- **Security and Trust Parameters:** a set of parameters to evaluate security between the peering CMs [2], [3].

On receiving a CPD_PRESP from CM_B, CM_A transmits a *CPD Advertisement (CPD_ADV)* message if it wishes to continue the local peering agreement formation with CM_B. Alternatively, if any of the parameters conflict with its own local security policy, it may either respond with a CPD_PRESP that rejects the local peering agreement formation, or simply ignore the CPD_PRESP from CM_B and allow CM_B to timeout awaiting an advertisement.

Each local peering agreement is then sustained through the regular exchange of CPD_ADV messages between peering CMs. This regular exchange provides additionally the mechanism to modify dynamically any of the aforementioned parameters.

This is effectively a soft-state approach. An obvious implementation strategy is to implement the handshake as a separate protocol, for example using multicast IP. However, it may also be possible to implement the CPD in other ways, for example, by piggy-backing some of the CPD parameters onto existing routing exchanges; through the use of special ICMP or DHCP messages; or by use of a modified form of the increasingly popular IETF ZeroConf [1] set of protocols.

III. Key network architecture challenges

Although the CPD uses terminology related to existing network standards and protocols, the particular architecture presents some specific challenges that cannot easily be met by currently deployed systems based on existing IPv4 and IPv6 standards. Here we discuss some of the problems of current network architecture, based on IPv4 and IPv6.

A. Distributed Administrative Responsibility

A key characteristic of our architecture is that Administrative responsibility is *distributed* across all Coalition Members. Thus, a CPD does not represent a single Administrative Domain (AD) that is under the control of a single organisation or entity, but rather a *collaborative* group of such entities. This is a particularly useful network characteristic for military coalition forces who need to retain full control of their own network, resources, policies and configurations. Existing mechanisms for network addressing and routing are not optimal in this context because they have been designed for environments in which administrative responsibility is hierarchic in nature. Although one could argue that such hierarchic administrative responsibility is wholly appropriate for any communications between military coalition forces, and that traditional intra-domain ad hoc routing mechanisms [4] provide an ideal solution, this is unsustainable within the highly dynamic environments that exist in a battlefield theatre.

B. Multi-homing and multiple network path routing

The CPD model is that of a multi-homed inter-network that has more than one inter-domain uplink, allowing multiple network path routing in to and out from the CPD. With these assumptions, the CPD has multiple egress/ingress points that are shared transparently by Coalition Members (CMs). Current approaches to IPv4 multi-homing have significant adverse impact on the upstream, global inter-domain routing

[1] http://www.zeroconf.org/

table; most Internet Service Providers (ISPs) believe that the current network multi-homing approaches do not scale adequately. Initially, IPv6 standards used the same approach to network multi-homing as IPv4. Because of scalability concerns from ISPs, the IETF has re-opened that portion of the IPv6 standards and begun work to develop a better approach and update the IPv6 standards [5]. Without scalable network multi-homing, the potential for robustness and load balancing of the CPD is probably lost.

The effects of multiple network path routing used for CPD egress/ingress traffic may have a noticeable impact on real-time communications especially. Such routing could particularly disrupt the operation of higher-layer protocols and congestion control mechanisms that assume specific behaviour (e.g. single-path routing) from the underlying routing infrastructure. The use of standard play-out buffering techniques at receivers may alleviate some of the problems that arise from delayed and mis-ordered data packets. Also, it may be possible to use adaptive forward error correction techniques, such as redundant encoding, to cope with some mis-ordering as well as some loss. However, this comes at the cost of additional play-out delay. Alternatively delayed and mis-ordered data packets may need to be dropped to maintain better continuity of real-time communications.

C. Mobility

Support for mobile nodes has become increasingly important for IP and not just for military use. Despite widespread interest, mobility support in both IPv4 and IPv6 remains poor. Although there has been much research and also much standards development in this area, including some support for Mobile Ad-hoc NETworking (MANET), real implementations and working systems are not deployed widely. For IPv4, mobility support is a retrofit and has some architectural drawbacks e.g., incompatibility with commonly deployed Reverse Path Forwarding (RPF) access control checks [6]. Although Mobile IPv6 has some improvements over Mobile IPv4, it is optional to implement, is not widely implemented at present, and remains an add-on feature rather than being a native property of the network protocol. The absence of any large scale deployments of either Mobile IPv4 or Mobile IPv6 creates uncertainty about how effective those standards really are. Simulations are not a substitute for actual large-scale deployments and operational experience in this regard.

D. Dynamic routing domain changes

In IP routing, the network is structured as a set of distinct routing domains and sub-domains, so as to provide hierarchical routing. Hierarchical routing has many advantages for scaling and engineering of the network. However, current IP routing standards allow neither the dynamic formation and identification of domains, nor the merging of domains as required for the operation of the CPD. Existing inter-domain routing protocols, such as BGP, are also too complex for direct use wihtin a CPD and it is not clear that there is benefit in adapting them for use within a CPD.

Additionally, part of the routing and CPD configuration control mechanism requires the use of discovery protocols that are not currently part of either BGP or interior routing protocols that assume wired/fixed network topologies (e.g. OSPF, ISIS, RIP).

At the time of writing, MANET protocols do support some desired dynamic features of routing and discovery, but lack the ability to have dynamic address management (see below) and do not always easily support network multi-homing. However, the wealth of research in MANET offers many mechanisms and ideas that can be adapted to enable the implementation of the CPD.

E. Dynamic address changes

In IP, the address has topological significance. That is, the IP address can be seen as a *locator* for an end-system communication interface within a given IP network. As such, it is tied very closely to the routing of packets, and changes to the address affect the routing of packets. In turn, it is clear that when CMs move, this sometimes will require that network addresses of end-system interfaces might need to change as the network topology changes. This has extremely important implications for the operation of upper-layer protocols (e.g., TCP, UDP, SCTP), and also for security protocols (e.g., IPsec, Secure Sockets Layer(SSL) or Transport Layer Security(TLS)).

F. IP Security

At present, military versions of IPsec are being deployed rapidly, both in the US DoD networks and in coalition partners' networks. It would seem reasonable to use IPsec for IP level security for an IP based communication network. However, the absence of a topology-independent identifier within the current Internet architecture means that IPsec is forced to use the IP address as an identifier in its Security Associations. So, when the IP address of a node's interface changes, as is often required today for mobile nodes and as would be required in the CPD architecture, the identity of that end-system changes and any existing IPsec security associations are broken. If a topology-independent identifier existed in the network architecture, then IPsec Security Associations (SAs) could bind to that in lieu of the IP address. In such a case, mobility (or even Network Address Translation - NAT) would be transparent to, and would have no adverse affect upon, the IP Security protocols.

G. Upper-layer Protocols

Unfortunately, the topology-dependent IP address is also used as part of the session state information in higher-level protocols, at the transport-layer (e.g. TCP, UDP, SCTP) and sometimes even at the application-layer (e.g FTP). This means that if there are changes to the IP address, for example because the node moves from one location to another, then that session state becomes invalid and so the session must be reinitiated. Very few applications can cope gracefully with such disruption to the communication session. Solutions offered in IPv4 and IPv6 to address this usually require the use of additional 'dummy addresses' in session state information plus an extra layer of processing to maintain correct mappings between the address in the session state and the real IP address currently in use over the air or on the wire.

IV. The Identifier Locator Network Protocol

The previous section might make it appear that it is impractical to implement and deploy the CPD. However, that discussion is presented specifically to highlight the extremely challenging nature of this problem space. Our analysis shows that it would be possible to implement the CPD in a number of ways, of which the most elegant architectural proposal enables the operation of the CPD at the packet level (the user plane) in the form of the *Identifier Locator Network Protocol (ILNP)* [7].

A. Dual-use IP addresses

We make a key observation that emerges from the discussion in the previous section: the fact that currently, in both IPv4 and IPv6, the IP address has the dual role of being both a (topological) *locator* to allow routing as well as being an *identifier* used in end-system session state. More specifically, the same bits of the address that are used for the locator are also used for the identifier. We see that for multi-homing, mobility and end-to-end security using IPsec, the overlap of the identifier bits and the locator bits of the address causes upper-layer protocol sessions to break. We believe that this observation is key to presenting an elegant architectural solution to the problem space.

B. A modified usage for the IPv6 address

ILNP is derived in large part from IPv6. In fact, they share virtually the same packet header and can use the same IP-layer options. The key difference between IPv6 and ILNP is that we propose a new concept for the 128 bit address space of IPv6. We propose that the most-significant 64 bits of the IPv6 address (the IPv6 *address prefix*) are used only to name a single sub-network,[2] while the least-significant 64 bits are used only as a node identifier. It is important to note that this is specifically as a *node* identifier, not as an *interface* identifier (as is currently the case with both IPv4 and IPv6). With ILNP, the high-order 64 bits of the IP address are called the *Locator*, *L*, and the low-order 64 bits are called the *Identifier*, *I*. We propose further that only *L* is used for routing in the core network; meanwhile, only *I* is used for session state in higher-level protocols (e.g. TCP, UDP, FTP), and for IPsec Security Associations.

This simple separation of the bits used for Locator and Identifier functions means that packets are routed through the core network (everything but the final hop) using only *L*. Local packet delivery of the packet over the last-hop link still uses 128 bits (both *I* and *L*) via the existing IPv6 Neighbour Discovery protocol [8]. So, it is clear that core routing protocols, forwarding of packets, and local (last-hop) packet delivery are all unchanged with respect to IPv6.

As only *I* is used for the session state, *L* can change, for example to support mobility or due to CPD address prefix changes as new CPDs are formed. Changes in values of *L* change only the routing, which is desired for mobility or for the operation of the CPD. However, as *I* remains unchanged when a node's routing prefix changes, then sessions are not affected when *L* changes. This also provides scalable network multi-homing: a system could transmit and receive packets with the same value of *I* but with different values of *L* for a twin-homed network or CPD. Unlike current approaches to network multi-homing, this approach has no impact on the Internet's inter-domain routing table and meets the expressed multi-homing objectives of ISPs.

This requires that the value of *I* be unique within the the scope of the local IP network or CPD, at least. The value of *I* could be opaque and requires no structure. However, there are various ways that such a 'unique' 64-bit value could be generated. The most obvious of these is to take one or more IEEE MAC addresses and use them to form values of *I* in EUI-64 format.

We currently recommend the following structure. The IEEE standard for EUI-64 identifiers defines a *local-scope* bit and a separate *multicast-bit*. If the local-scope bit is 0, then the the value of *I* is derived from an IEEE MAC address and hence is very likely to be globally unique. If *I* is not unique within the scope of the CPD (e.g. because of MAC address clashes) then link-layer communication (e.g. IEEE bridging) will be disrupted before the network-layer clash with ILNP becomes significant. If the multicast bit is 0, then the *I* names a specific node. If the multicast bit is 1, then the *I* names a specific multicast group. For a multicast packet, the destination Locator, *L*, usually contains a routing prefix for a multicast router that is a Rendezvous Point for the named multicast group. This multicast approach makes

[2] This does not preclude the use of a mask of /48 (the top 48 bits) for core routing leaving 16 bits for sub-netting.

ILNP compatible with existing multicast routing protocols, such as PIM [9] or CBT [10].

C. Impact of ILNP

ILNP has been carefully derived from IPv6. So, no changes are required to an IPv6 network core (to IPv6 core routers) or to local IPv6 Neighbour Discovery to support ILNP. The essential enhancements are focused on upper-layer protocols (e.g. TCP, UDP) and the Domain Name System (DNS). This strategy enables incremental deployment of ILNP within an IPv6-capable network segment by modifying only those *end-systems* that require the use of ILNP over an IPv6 core. Indeed, it should be possible to use ILNP and IPv6 concurrently on the same network.

For the upper-layer protocols, the pseudo-header checksum needs to be modified to use only *I* in the calculation, omitting *L*. Also, some additional network-layer enhancements are required to enable ILNP in hosts, specifically changes to improve security and provide assurance for the binding between a given Identifier and its associated Locators. As these changes are only in the ILNP *end systems*, existing IPv6 core routers can forward ILNP traffic without modification.

For the DNS, we replace the **AAAA** (or **A6**) records with a pair of new DNS resource records, **I** records for the Identifier(s) and **L** records for the Locator(s) of a given node. We also replace the existing DNS pointer record with the new **PTRL** and **PTRI** records. The reverse lookup process is slightly different than at present, moving from a one-step procedure to a two-step procedure. In the new procedure, one first makes a PTRL request for a given Locator value, to learn the authoritative DNS server(s) for that Locator value. One then sends a PTRI query to an authoritative DNS server specified in the PTRL response to obtain the fully-qualified domain name of the node specified by that L value and that I value. The existing IETF standard for Secure Dynamic DNS Update [11] is used by mobile or multi-homed ILNP nodes to update their L and I records in their DNS server(s). For added robustness, Multicast DNS (mDNS) could be used in the CPD if required.

As an optimisation, and to facilitate incremental deployment of the new CPD and ILNP networking architecture, we also propose to provide an enhanced networking API that might gradually replace the existing BSD Sockets networking API. This new API provides greater data hiding and a more suitable networking abstraction. This should reduce the time and effort required for a programmer to network-enable an application, while the improved abstraction should enable the new API to work well for hosts that implement IPv4, IPv6, and/or ILNP.

However, in the interim, we believe that it should be possible to make modifications to the existing IPv6 BSD Sockets API and so port a large number of existing applications to ILNP quickly.

With respect to the overall impact of CPD deployment on existing network infrastructure, only the DNS servers acting for CPD end-systems need enhancements with no requirements to update the entire population of deployed DNS applications in order to deploy the CPD.

V. Security Implications

For the military application domain, communications security is a key requirement. Our new network architecture is compatible with a range of network security approaches. This includes network security mechanisms currently deployed in military networks with concepts such as the Pink Network Architecture [12] originating at NRL, and the Black Core Network proposed by ASD/NII. We believe HAIPIS-compliant network encryptors that support IPv6 should be able to also support ILNP without modification. Both the CPD and ILNP operate at the network-layer and so affect only control-plane mechanisms. This leaves upper-layer protocols and applications unaffected, allowing them to continue using existing security mechanisms such as Secure Sockets Layer (SSL), Pretty Good Privacy (PGP), Secure/Multimedia Internet Mail Enhancements (S/MIME) and IPsec.

Further, when IPsec is used with ILNP, IPsec SAs are bound to only the Identifier and not the full IP address as is the case with IPv4 or IPv6. This enhances greatly the effectiveness of IPsec, which as a result shall operate seamlessly through NATs as well as between mobile hosts and across mobile networks.

Meanwhile, specific security policies or rule-sets may be applied by prospective CMs during the CPD handshake process to determine whether the level of security is sufficiently acceptable for establishment of a local peering agreement. Additionally, the structure of the CPD has an improved resistance to some Denial of Service attacks such as Sybil [13] and routing black holes. This is a direct result of both the handshake process used during CPD formation, and the multi-homed, multiple-path routed nature of the CPD.

VI. Related Work

There exist currently a number of systems and community-area initiatives that may provide useful tools and mechanisms that could be used to improve communication between coalition forces on battlefields, e.g., [14], [15], [16]. However, the high levels of security, flexibility and local control needed to protect the integrity of communications, the safety of coalition forces and their ability to carry out

their mission without being compromised, mean that such systems are unable to provide a complete solution.

One specific system that provides a close fit to enable better communication between coalition forces within a battlefield scenario is the HDNet system [17]. This focusses on a highly dynamic multi-hop wireless network model in which clustering is used to allow higher powered 'mobile base stations' to forward data on behalf of lower powered 'mobile hosts'. HDNet explores in some detail the aspects of relative mobility between the higher powered mobile base units and the lower powered mobile hosts. However, the gateway function provided by the mobile base stations make the underlying assumption of single-path routing, so multiple such units deployed together are unable to communicate and share resources effectively.

VII. Conclusion and further work

We have presented a proposal, based on our ongoing work, for a network architecture that can provide robust, flexible and efficient communication in the the battlefield. The system could be used across different coalition forces, discovering other friendly resources and self-configuring so as to build a Coalition Peering Domain (CPD), a single network entity leveraging as many of the available communication resources as policy allows. The system is built on the Identifier Locator Network Protocol (ILNP), and can be deployed incrementally as it has excellent compatibility with IPv6.

There are design and engineering choices to be made for some parts of the architecture. The biggest impact of the system would be the modifications that would be required for the DNS. Although these modifications are non-trivial, they are tractable and do not hinder the ability to deploy.

The integration of the CPD and the ILNP offers a significant advantage for communication within battlefields.

Our future plans include building prototype implementations of ILNP and the CPD, integrating those implementations together, and using the prototypes to begin experimental validation of the ideas presented in this paper. We contemplate use of an existing IPv6 research wide-area network as part of those experiments, both to demonstrate that the core of an existing IPv6 network does not need modification to support CPD and ILNP, and also to provide wider geographic reach for our experiments

VIII. Acknowledgements

The authors acknowledge support from the IST-RUNES project. Atkinson's work on ILNP has been supported by his employer. The idea of coalitions for netwokring originates in part from the work of Defence Research & Development Canada (DRDC) on coalition-based dynamic VPN infrastructures. Various discussions and deployment activities for DVC were undertaken at UCL during 2003-2004.

References

[1] M. Lad, S. Bhatti, S. Hailes, and P. Kirstein, "Enabling Coalition-Based Community Networking," in *Proc. London Communications Symposium 2005 (LCS 2005)*, 8–9 September 2005.

[2] D. Quercia, M. Lad, S. Hailes, L. Capra, and S. Bhatti, "Trusted Bandwidth Sharing in Mesh Networks," *BT Technology Journal*, vol. 24, no. 3, 2006.

[3] D. Quercia, M. Lad, S. Hailes, L. Capra, and S.Bhatti, "STRUDEL: Supporting Trust in the Dynamic Establishment of peering coaLitions," in *Proceedings of The 21st Annual ACM Symposium on Applied Computing (SAC 2006)*, 23-27 April 2006.

[4] "IRTF RRG Ad hoc Network Systems Research Subgroup," http://www.flarion.com/ans-research/.

[5] IETF, "Site multihoming by ipv6 intermediation (shim6)," http://ietf.org/html.charters/shim6-charter.html.

[6] F. Baker and P. Savola, "Ingress Filtering for Multihomed Networks," RFC 3704 (Best Current Practice), Mar. 2004. [Online]. Available: http://www.ietf.org/rfc/rfc3704.txt

[7] R. Atkinson, "A Revised Overview of the Identifier-Locator Network Protocol (ILNP)," 28 October 2005, UCL Department of Computer Science Research Note 05/26.

[8] T. Narten, E. Nordmark, and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)," RFC 2461 (Draft Standard), Dec. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2461.txt

[9] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification," RFC 2362 (Experimental), June 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2362.txt

[10] A. Ballardie, "Core Based Trees (CBT version 2) Multicast Routing – Protocol Specification –," RFC 2189 (Experimental), Sept. 1997. [Online]. Available: http://www.ietf.org/rfc/rfc2189.txt

[11] B. Wellington, "Secure Domain Name System (DNS) Dynamic Update," RFC 3007 (Proposed Standard), Nov. 2000, updated by RFCs 4033, 4034, 4035. [Online]. Available: http://www.ietf.org/rfc/rfc3007.txt

[12] R. Cole, D. Kallgren, R. Hale, and J. R. Davis, "Multi-Level Secure Mixed-Media Communication Networks," in *Proc. IEEE MILCOM 1989, Volume 1, Boston, MA, USA*, October 1989, pp. 117–121.

[13] J. R. Douceur, "The Sybil Attack," in *Proc. 1st International Workshop on Peer-to-Peer Systems*. Springer-Verlag, March 2002, pp. 251–260.

[14] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee, "MAR: a commuter router infrastructure for the mobile internet," in *MobiSYS '04: Proceedings of the 2nd international conference on Mobile systems, applications, and services*. ACM Press, 2004, pp. 217–230.

[15] M. Papadopouli and H. Schulzrinne, "Connection Sharing in an Ad Hoc Wireless Network among Collaborating Hosts," in *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, June. 1999, pp. 169–185.

[16] "Champaign-Urbana Community Wireless Network," http://www.cuwireless.net/.

[17] R. Sanchez, J. Evans, and G. Minden, "Networking on the Battlefield: Challenges in Highly Dynamic Multi-hop Wireless Networks," in *IEEE Military Communications Conference*, Atlantic City, NJ, Oct. 1999.