

Secure & Agile Wide-Area Virtual Machine Mobility

Saleem N. Bhatti
University of St Andrews, UK
saleem@cs.st-andrews.ac.uk

Randall Atkinson
Cheltenham Research, VA, USA

Abstract—Global Information Grid (GIG) operations would benefit greatly from improved support for virtual machines (VMs) that can migrate not only between physical devices within a datacentre, but also between physical devices located on different continents, while maintaining their existing IP communications sessions. Such VM migration can enable improvements with: CPU load-balancing, network traffic-engineering, distributed denial of service (DDoS) mitigation, fault-tolerance, and resilience. Existing migration approaches often require complex network configuration and management, may often require use of expensive proprietary technologies, and also often require active cooperation from upstream service providers. We describe a VM mobility approach that enables datacentre operators to directly and unilaterally provide and control intra-site and wide-area VM mobility. We present several use cases with different degrees of location transparency. Our mechanism is based on a new naming approach which has been recommended for progression within the IETF.

I. INTRODUCTION

DISA and other DoD components operate multiple very large (e.g. measured in CPU-acres) geographically distributed data centres to support warfighter operations. These data centres already have achieved significant improvements in hardware utilisation and overall efficiency from widespread use of commercial virtualisation technology (e.g. VMware, Xen) so that multiple logical hosts (“virtual machines”) are supported concurrently by a single physical device [1]. Similarly, virtual machine (VM) migration from one server within a data centre to another server within the same IP subnetwork (“bridged layer-2 network”) has been described previously in the research literature [2] and are available commercially. [3]

An important capability to achieve further improvements in availability, efficiency, resilience, and utilisation is the ability to migrate ‘live’ virtual machines (VMs) from a physical device in one location to a different physical device in a remote location. In turn, this requires that existing IP sessions remain up and running despite the migration of the VM from one device to another, possibly across continents. So there is specific current military interest in technologies that can improve these “VM migration” capabilities, in order to provide more comprehensive support to the warfighter [4].

Prior IETF work on Mobile IP has not resulted in a capability that is widely supported in commercial off-the-shelf products used in wide area VM migration deployments. So, at present, commercial approaches to VM migration require that all nodes in a single ‘migration group’ are members of

the same IP subnetwork. Within a single data centre, very large bridged Ethernet LANs often are used, which can create bridge table convergence time issues. For wide area VM migration, various technologies are used at present to create wide area bridged layer-2 networks, which are necessary to keep existing IP sessions live [5] [6]. Current approaches require specialised support from service providers (e.g. MPLS L2 VPNs) or from site border devices (e.g. MPLS L2 VPNs, GRE tunnels) and create additional encapsulation overhead (e.g. GRE tunnels, MPLS L2 VLNs). These current layer-2 approaches to LAN/MAN/WAN VM migration create significant capital, deployment, and operational costs. Also, there are concerns as to how such deployments will scale. A practical operational security issue is that such large layer-2 networks can be more brittle in the event of a cyber assault and can impede the operator’s ability to deploy security devices that provide domain separation and risk partitioning.

A. Requirements

The Department of Defense (DoD) *Cloud Computing Strategy* [7] outlines an information strategy based on the use of cloud services, supported by virtualisation of services. In support of this position, we believe VM mobility and migration approaches should:

- Enable data centre operators to maintain essential services without requiring specialised networking features (e.g. layer-2 VPNs) from their wide-area network providers or deployment of specialised network gateways.
- Enable wide-area VM mobility (e.g. between continents) across different routed IP networks, in addition to enabling local-area (e.g. within a data centre) VM mobility.
- Avoid interruptions to data centre services for critical GIG applications, services, and other capabilities.
- Avoid dependence on any specific network design, in order to enable adaptive data centre network designs that maximise resilience, fault-tolerance, and scalability.

For the first point, there are various constraints today. A common constraint is that commonly deployed VM mobility solutions require a large flat layer-2 bridged network today, rather than working well across routed IP subnetworks.

For the second point, both commercial and governmental data centre operators have expressed frustration that existing wide-area VM mobility approaches often do not work well, usually require expensive proprietary technologies that create

vendor lock-in, and might also impose significant operations and management overhead. Our approach is vendor-independent and is not even specific to VM mobility, but instead also can be used to provide generalised and integrated host and site mobility, host and site multi-homing, and other capabilities [8]–[11].

Today, the third point above can be achieved fairly well, with a small downtime experienced during the final stages of migration, given appropriate provisioning for the data centre platform and network connectivity.

For the last point, we take the position that any capabilities should be enabled and controlled only from the site-network. This position reduces the security exposure of the site-network, by not requiring trust relationships with, or administrative dependencies upon, third parties such as service providers, beyond providing basic network connectivity [10], [11].

B. Structure of this paper

This paper presents some background information and a problem description in Section II. After a description of the salient aspects of ILNP in Section III, this paper describes several general use cases, in Section IV. Then, three specific capability scenarios are presented: VM mobility in a single data centre, (Section V), VM mobility across sites with client transparency (Section VI), and VM mobility across sites without client transparency (Section VII). A short description is also given of additional capability when distributed applications support ILNP, in Section IX. A concluding summary is presented in Section X.

II. BACKGROUND AND RELATED WORK

We summarise briefly the background and related work to introduce the problem statement. However, we do not attempt to provide a comprehensive literature survey.

Earlier work has shown VM mobility to be possible with good performance [2]. The basic principle involves exploiting the dynamics of the operating system paging mechanism to ‘write’ pages of memory from the *current* physical device, where the VM is currently resident, to the *target* physical device of the VM migration operation.

That study was concerned mainly within a single site, but today there is great interest in wide-area migration of VMs, for example to provide operational resilience, e.g. [12]. However, such wide-area migration must maintain existing communication flows within the VM being migrated, and so the target host must be on the same IP subnetwork as the current host. For wide-area connectivity, this means that the network-site of the target hosts must be seen as part of the same IP subnetwork. This can be achieved today through various engineering techniques including the use of tunnels and proxies, e.g. [13]. This leads to complex and potentially brittle infrastructure, which requires careful configuration and management, and might not scale well to large numbers of VMs or physical hosts. Such engineering issues are examined in [13] and ‘extending the LAN’ remains the basis of commercial solutions today to provide multi-site and wide-area migration, e.g. Cisco VXLAN [6] and Juniper QFabric [5].

Our approach is based on the Identifier Locator Network Protocol (ILNP), an evolution of the current Internet Architecture which adds additional namespaces and has been recommended for standardisation within the IETF [14]. Other approaches for evolving the Internet include the Host Identity Protocol (HIP) [15] and the Locator Identifier Separation Protocol (LISP) [16]. While HIP also provides crisp separation of identity from location, HIP differs from ILNP in several ways. For example, HIP protocol specifications require the deployment and use of strong cryptography, even within protected enclaves. Ordinary software implementations of HIP incur a higher computational burden for each HIP packet due to the required per-packet cryptography, which can impair both application performance and network interface performance. This issue can be resolved by using cryptographic hardware accelerators, which increases the hardware cost. ILNP supports use of strong cryptographic protection, for example High-Assurance IP Encryption (HAIBE), but does not require it in all deployments, permitting deployment-specific cost/benefit tradeoffs. Unlike both ILNP and HIP, LISP relies on mapping IP addresses into a separate routing schema and encapsulating IP packets sent between LISP nodes. This increases the per-packet protocol overhead, which is a significant concern in low-bandwidth tactical links, and also increases the routing complexity of the resulting network deployment. Additionally, unlike HIP and ILNP, LISP requires upgrades to existing network infrastructure and requires additional management and control to support the new routing mechanism which is introduced into provider networks. Research prototypes of ILNP, HIP, and LISP exist today.

The position taken in the first point above – autonomous operation of the site-network – is explained fully in [10]. This position is important when considering the discussions later on with respect to certain capabilities, such as multi-homing: we do not claim that ILNP can uniquely provide such capabilities, only that it offers features that are not possible to realise today, namely that the work presented here (a) uses an end-to-end architecture that can be controlled from and by the site-network alone; and (b) can provide several important network capabilities in harmony, concurrently.

A. Problem Statement

The requirement for the same IP subnet for the target hosts stems from a fundamental problem in the use of IP addresses. In the right-hand column of Table I, we see a summary of the use of namespaces in IP today (both IPv4 and IPv6). The IP address is used for various purposes in the protocol stack. This creates implicit bindings between communication objects in the otherwise separate layers. For example, today all TCP and UDP flows are bound to a specific interface on a host, impacting mobility and multi-homing¹. If an application-layer protocol uses the IP address instead of a Fully Qualified Domain Name (FQDN) for its session state, then that application-layer session is bound to a specific IP address, which in turn

¹SCTP supports the use of multiple interfaces concurrently, but is not widely deployed. The IETF is studying Multi-Path TCP (MP-TCP) extensions.

is bound to a specific interface on a specific node, unless the application includes specific application-layer support for mobility. Such application mobility support is rare in practice.

TABLE I
USE OF NAMES IN ILNPV6 AND IP.

Protocol layer	ILNPv6	IP (v4 and v6)
Application	FQDN	FQDN, IP address
Transport	Identifier, <i>NID</i>	IP address
Network	Locator, <i>L64</i>	IP address
(interface)	dynamic binding	IP address

This paper contributes an architectural description of how ILNP can provide extremely flexible support for VM mobility for several common scenarios, by using an architecture with a cleaner set of namespaces, as shown in the second column of Table I. Also, we highlight security issues, as these are particularly important for the GIG.

An empirical performance evaluation of ILNP’s Locator Update mechanism, (described later) was presented in a previous MILCOM paper [11], based on the use of a lab prototype. However, the new ideas presented here have not yet been experimentally evaluated. Many ILNP features have been described previously in [8]–[11], [17]–[20]. Additionally, ILNP has been subject to public, expert scrutiny and debate within the IRTF Routing Research Group (RRG), and was recommended for progression within the IETF [14].

III. OVERVIEW OF ILNP

The *Identifier-Locator Network Protocol (ILNP)* is a set of backwards-compatible, incrementally-deployable extensions to IP.² ILNP permits harmonious integration of mobile nodes, mobile networks, localised addressing, node multi-homing, site multi-homing, network-layer end-to-end security, traffic engineering options, and multi-path transport-layer flows. With ILNP, host network software is upgraded and no router changes are required. For some use cases, however, an ILNP-enabled Site Border Router (SBR) may be beneficial, offering a centralised management mechanisms for certain capabilities.

The IRTF Routing Research Group report recommended that the IETF standardise ILNP [14]. The formal ILNP protocol specifications³ have been extensively reviewed within the Internet community and already have been approved for publication as Internet RFCs. At least two native open-source implementations of ILNP are underway. The University of St Andrews is implementing ILNP in FreeBSD, while a university in P.R. China is implementing ILNP in Linux.

We have previously described several military-relevant capabilities enabled by ILNP [8]–[11], [20]. Results from an overlay implementation of ILNP also were provided previously [11]. Here, we give only a summary of the main ILNP functions as they pertain to our current discussion. In this discussion, for simplicity, we discuss only ILNPv6, an instance of ILNP which is a set of IPv6 extensions.

A. ILNP – Architecture

The Identifier-Locator Network Protocol (ILNP) provides a cleaner namespace for the protocol stack as shown in Table I. Applications can use their own namespace, but default to using FQDNs, consistent with a long-standing IAB Recommendation [21]. Transport-layer protocols use only a *Node Identifier (NID)* (or just *Identifier*), which has no topological significance and names a *node*, rather than a particular interface of a node. For example, transport-layer end-system state includes only the *NID* value, rather than a whole IP address. The network layer uses topologically- significant *Locator (L64)* values only for routing and forwarding.

Unlike IP, ILNP supports dynamic bindings between *NID* values, *L64* values, and network interfaces. Nodes can choose to use multiple *NID* and *L64* values and use multiple interfaces simultaneously, and adjust dynamic bindings between them as required (e.g. to implement mobility and multi-homing). Dynamically binding transport-layer state (and application-layer state) to interfaces enables mobility as a first class function, and eliminates the need for the indirection and tunnelling of Mobile IP.

Table I summarises the differences in naming architecture. Well-behaved applications (e.g. those that use FQDNs as application-layer names or work through a NAT without specialised support) should not require modification to operate over ILNP. Any ILNP device can use multiple (different) *NID* values simultaneously, which can be assigned to logical, virtual, or physical nodes. As *NID* values are bound dynamically to *L64* values, this gives immense flexibility for mobility, including the mobility of virtual machines.

In summary, for ILNP, *end-to-end protocols bind to NID values*, which are used above the network layer only. *L64* values are, effectively, names of networks (e.g. a network prefix as used today), with dynamic bindings between *NID* values, *L64* values, and interfaces.

B. ILNP – Engineering

ILNPv6 can be deployed over existing IPv6 infrastructure without requiring router or routing system changes. *L64* and *NID* values are encoded into the IPv6 address space – see Figure 1. The top 64 bits are the *L64*, which has the same syntax and semantics as an IPv6 routing prefix. So, existing mechanisms (e.g., IPv6 Router Advertisements) can be used to determine its value. The lower 64 bits are the *NID*, which has the same syntax as the IPv6 Interface Identifier, but different semantics. A *NID* value names a node, rather than a specific interface of the node. *NID* values are not used for routing in the network core. The *NID* value is used only by the end-host, so only end-system code needs to be updated. The *NID* values can be chosen in the same ways as IPv6 Interface IDs can be chosen [22]. This can also leverage related IETF standards that can encode a different *NID* value. For example, nodes may use cryptographically generated *NID* values [23], or choose *NID* values to obfuscate node identity for privacy reasons [24]. Hence, ILNPv6 is backwards-compatible and incrementally deployable with IPv6.

²<http://ilnp.cs.st-andrews.ac.uk/>

³<http://tools.ietf.org/id/draft-irtf-rrg-ilnp>




Fig. 1. Comparison of IPv6 unicast format [25] with ILNPv6 unicast format.

C. Localised Addressing

As Locator values are not part of end-to-end state, then *Locator re-writing* optionally can be used with ILNP to dynamically change Locator values for packets (in a similar fashion to IPv6 prefix translation as in [26]). For example, Figure 2 describes an established TCP connection between a host (H1) and a correspondent node (CN) within a site network. CN’s end-to-end session state is expression (1), while H1’s state for the same connection is expression (2), where P values are port numbers. From a TCP viewpoint, CN and H1 see the same connection. However, the SBR re-writes the Locator value in packets to/from H1 with the value L_L (for ingress) and value L_1 (for egress).

$$\langle TCP : P_{CN}, P_1, I_{CN}, I_1 \rangle \langle ILNP : L_{CN}, L_1 \rangle \quad (1)$$

$$\langle TCP : P_1, P_{CN}, I_1, I_{CN} \rangle \langle ILNP : L_L, L_{CN} \rangle \quad (2)$$




Fig. 2. Localised addressing with ILNP. As the end system protocol state at the CN binds only to I values, changes to L values have no significance. This means that SBR1 can re-write value L_1 in packets to an internal value used only within the site network, e.g. value L_L .

This is ILNP’s equivalent of IP’s widely deployed *Network Address Translation (NAT)* mechanism. However, manipulation of Locator values also permits location privacy, multi-homing, mobility, and traffic engineering – without loss of end-to-end transparency [8]–[11], [20].

D. Network-Layer Security

A full description of end-to-end network-layer security using IP Security for ILNP, was given in our previous work [10]. In summary, IP Security for ILNP provides the same security capabilities for ILNP that IP Security provides for IP, using the same basic mechanisms. Since the *High Assurance IP Encryptor (HAIPE)* used to protect existing military IP networks is a US DoD profile of the IETF IPsec standards, ILNP can support military-grade network security requirements.

IV. DATA CENTRE USE CASES

There are several different data centre use cases that could benefit from the approach to VM mobility and migration described here, as compared with existing approaches that rely upon very large flat bridged layer-2 networks (both within and between data centres).

A common use for VM migration is load-balancing a set of VMs across a set of physical devices, either within the same data centre or between different data centres. Unlike current approaches, our proposed approach allows a data centre to be broken down into a number of separate Virtual LANs (VLANs), where each VLAN typically is a separately routed IP subnetwork. This separation provides better scalability of the network and for the associated server devices. For example, broadcast traffic (e.g. ARP) is scoped to a single VLAN/IP subnet, and does not affect VMs or devices on separate VLANs/IP subnets. Indeed, our approach could use existing network configurations, whereas the ‘extended LAN’ approach being used widely requires additional network configuration and management overhead.

VM migration also is used to optimise service resilience & fault tolerance. For example, if a data centre is impaired but still available (e.g. primary power has failed and the facility is operating on backup generators), particular VMs might be relocated to a different data centre to reduce the load on the impaired data centre, while keeping network-enabled services operational continuously, including maintaining existing communication sessions.

A third use is to improve service performance by moving VMs, providing selected services closer topologically to a set (or subset) of their client nodes. Since TCP uses a sliding-window algorithm, reducing the delay along the path between two nodes can improve the effective throughput for a TCP session between those nodes. However, for many services, offered load from a particular geographic location varies with the time of day. Scalable wide-area VM mobility can enable user-facing, performance-sensitive services to move to the data centres near the current cluster of users, and also enable delay-tolerant background batch-processing tasks to move to data centres away from the current cluster of users.

These various use cases can be broken down into three separate functional scenarios:

- *Same data centre, internal VM mobility:* This could be beneficial in load-balancing, dynamically, where load changes are taking place. The approach described here hides the VM migration from the remote correspondent of the VM that has moved, while still allowing internal networks to be broken into a number of different IP subnetworks.
- *Different data centres, transparent VM mobility:* This is where the data centre resources may be geographically distributed, but the geographical movement is transparent to the remote user.
- *Different data centres, VM mobility is visible:* This is where the data centre resources may be geographically

distributed, but the geographical movement is visible to the remote user.

These are three operational scenarios which can be supported by ILNP, but they are not the only ones. We provide these as examples, but they are not intended to be prescriptive. The intention is to show the flexibility that is possible through the use of ILNPv6.

The following sections describe some VM mobility capabilities that are possible with ILNP. Depending on the internal details and virtualisation model provided by a VM platform, it might be sufficient for the guest operating system to support ILNP. In a few cases, again depending on the internal details and virtualisation model provided by a VM platform, the VM platform itself also might need to include support for ILNPv6.

We present here the conceptual realisation of these three deployment scenarios, to demonstrate that the ILNP architecture is able to support such VM mobility. Details beyond the scope of this paper include:

- Which virtualisation model(s) a VM platform supports.
- Implementation details of a particular VM platform.
- Implementation details of ILNP in a VM platform.
- Implementation details for other ILNP systems.

V. VM MOBILITY – SINGLE DATA CENTRE

First, consider the scenario of Figure 3, noting its similarity to Figure 2 for use of localised numbering. L_L is a Locator value used for the ILNP hosts H1 and H2 and V has Identifier I_V . Here, the ‘:’ in the diagram signifies that the virtual machine V is currently resident on H1.

A. Intra-Site Mobility

As H1 and H2 have the same L_{64} value, if V is resident either on H1 or H2, all transport packets between V and CN will have the same state as far as CN is concerned, e.g. for a TCP flow we have expression (3) as the TCP state at CN, and expression (4) as the TCP state at V.

$$\langle TCP : P_{CN}, P_V, I_{CN}, I_V \rangle \langle ILNP : L_{CN}, L_L \rangle \quad (3)$$

$$\langle TCP : P_V, P_{CN}, I_V, I_{CN} \rangle \langle ILNP : L_L, L_{CN} \rangle \quad (4)$$




Fig. 3. VM mobility within the same site. H1 and H2 have the same locator value, L_L . So, the CN would be unaware of the VM moving from H1 – position (1) – to H2 – position (2). The ‘:’ between V and H1 signifies that V is bound to H1.

Now, if V were to migrate to H2, the migration would be an issue purely local to the site-network, and the end-to-end integrity of the transport flow would be maintained.

Of course, there are practical operating systems issues in enabling such a migration locally, but products exist today that could be enhanced to be ILNP-aware in order to enable such VM mobility.

For convenience, above, we have used localised numbering for ILNP, but if local Locator values were not used and the whole site simply used L_1 (a globally unique prefix/Locator value), the principle would be the same, without the Site Border Router (SBR) needing to rewrite any Locator values.

Note that such capability is available today in existing systems: this section is to show how ILNP supports the capability that is available today.

VI. VM MOBILITY – BETWEEN DATA CENTRES, INVISIBLE

Now, consider an extended version of the scenario above, in Figure 4, where we see that there is a second site network, geographically distant to the first site network, and the two site networks are interconnected via their respective SBRs.

A. Inter-Site Mobility, Invisible to CN

Unlike the case above for Figure 3, now H1 and H2 have different Locator values, L_{L1} and L_{L2} respectively. However, L_{L1} and L_{L2} could be considered as two local subnets reachable via SBR1. So, any communication from CN still routes through SBR1, but SBR1 would select the correct current Locator value for V, for example, as a table look-up keyed by I_V . While we show two sites for simplicity, more than two sites also could be used. The end-system state expressions for CN, V:H1, and V:H2 are shown, respectively in expressions (5), (6), and (7)). Note that in the end-system state for CN, there is still no difference between V being at H1 or H2.

The logical inter-router link between SBR1 and SBR2 could be realised physically in many different ways that are available today and are not ILNP-specific, e.g. fibre or copper wire, leased circuit, secure IP-layer (IPsec) tunnel, MPLS tunnel, etc. This link also allows coordination between the two SBRs: some existing commercial security gateways already allow such coordination functions and these could be adapted as required. For now, we ignore external link L_2 on SBR2 with respect to CN, and assume that the remote node, CN, is in communication with V through SBR1. We can see that L_{L1} and L_{L2} are names for, effectively, two internal (private) subnetworks, and are not visible to CN.

However, SBR2 and SBR1 must coordinate so that any further communication to V via SBR1 is routed across the inter-router link. Again, existing commercial security gateways could be adapted to manage such shared state.

$$\langle TCP : P_{CN}, P_V, I_{CN}, I_V \rangle \langle ILNP : L_{CN}, L_1 \rangle \quad (5)$$

$$\langle TCP : P_V, P_{CN}, I_V, I_{CN} \rangle \langle ILNP : L_{L1}, L_{CN} \rangle \quad (6)$$

$$\langle TCP : P_V, P_{CN}, I_V, I_{CN} \rangle \langle ILNP : L_{L2}, L_{CN} \rangle \quad (7)$$

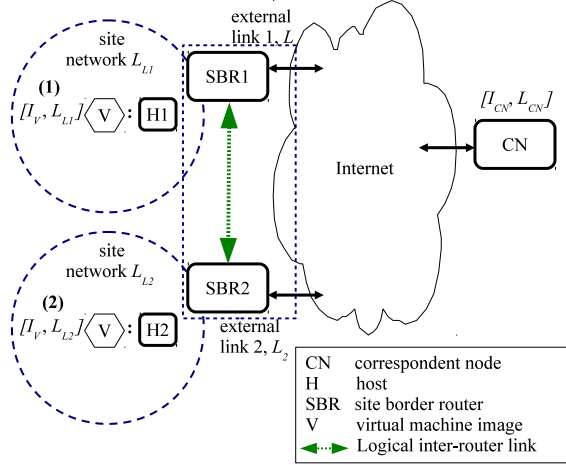


Fig. 4. VM mobility across sites. H1 – position (1) – and H2 – position (2) – now have the different locator values, L_{L1} and L_{L2} . However, these can be realised as subnets behind SBR1 (or SBR2), so the CN would still be unaware of the VM moving from H1 to H2. The logical link between SBR1 and SBR2 would carry a coordination protocol between SBR1 and SBR2, and would also carry the data for the VM migration, so it might be encrypted. The dotted box around SBR1 and SBR2 depicts the logical boundary to the data centre that is created via SBR1 and SBR2 using ILNP. The ‘:’ between V and H1 signifies that V is resident on H1.

Such capability is also available in systems today, but requires that the hosts involved in the VM migration must belong to the same IP subnetwork, which adds network configuration and management overhead. This constraint is not applicable to ILNP.

VII. VM MOBILITY – BETWEEN DATA CENTRES, VISIBLE

In the capability scenario of the Section above, once V has moved to the second site-network, it may be beneficial, for a number of reasons, for communication to V to be routed via SBR2 rather than SBR1. Reasons to route via SBR2 rather than SBR1 include:

- Moving network traffic away from site-network 1 and SBR1 for management purposes.
- Moving VMs away from site-network 1 as it has a fault, e.g. a systems fault, a power problem etc.
- Moving VMs away from site-network 1, if, for example, it is threatened by an attack or is being attacked, either by a physical offensive against the site, or it is subject to a cyberattack (such as a traffic-based DDoS attack).
- Having traffic always ingress through a single point and then re-routing to different, geographically-diverse data centres creates a single-point of failure, a performance bottleneck, and leads to sub-optimal traffic flows between the data centre and CN.

These are examples, but there may be other reasons. We do not make a comprehensive analyses of the reasons for inter-site migration being visible to remote clients. However, we take the position that such visibility may be beneficial to remote systems in the interests of both performance and security.

A. Inter-site Mobility, Visible to CN

When V moves from site network 1 to site network 2, the visibility of the mobility would be enabled by V sending ILNP *Locator Update (LU)* messages to the CN during the mobility process, to inform CN that it can now be reached via L_2 rather than L_1 . The use-case is depicted in Figure 5. LU messages are analogous to Mobile IPv6 Binding Update (BU) Messages [27]. Also, V would update any relevant ILNP DNS records, such as *L64* records, with the value L_2 , so that new session requests to V would be routed via SBR2. Indeed, V need not undertake such management functions: as SBR1 and SBR2 are aware of the mobility, either could perform these functions, e.g. as described in previous work [8], [9], [11].

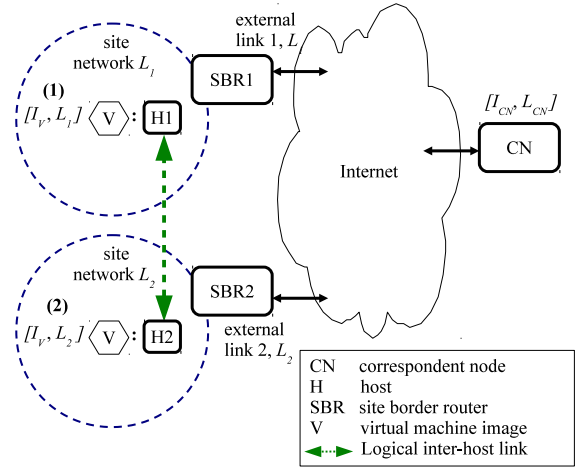


Fig. 5. VM mobility across sites. H1 – position (1) – and H2 – position (2) – now have the different global locator values, L_1 and L_2 . The inter-host link is an IP session that allows the VM images to move between the two hosts. The ‘:’ between V and H1 signifies that V is resident on H1.

A great advantage of having ILNP capability in both the CN and V is that then SBRs need not be ILNP-aware. That is, the flexible VM mobility is enabled purely by updating end-systems and does not require any infrastructure to be upgraded. This allows for incremental deployment, as well as trialling new applications or services using ILNP in parallel with existing deployments. Indeed, this is the primary model for ILNP mobility [18].

The logical inter-host link in Figure 5 allows a VM image to move from H1 to H2. For this, only an ordinary IP session between the two hosts is required. In situations where security policy requires, the IP session used to migrate the VM may be protected using Transport Layer Security (TLS), transport-mode or tunnel-mode IPsec, or other standard confidentiality and authentication mechanisms. A difference compared to today’s systems is that for ILNP, this is an IP session and

can cross routed network boundaries, whereas in many cases, existing VM migration tools require that H1 and H2 always reside on the same IP subnetwork.

B. Integrated Multi-homing

Let us now consider again Figure 4, but assume that Local Locators L_{L1} and L_{L2} are not in use on either site network, neither SBR1 or SBR2 perform Locator re-writing, and each site network uses its own global Locator values, L_1 and L_2 , respectively, and these vales are used directly by nodes in each site network respectively. In that case, the packet flow state for V when it is in site network 1 as viewed from CN is, again as given in expression (5). However, when V moves to site network 2, it would simply use L_2 as its new Locator, send Locator Update messages to CN (as would a normal mobile node for ILNP), and complete its migration to H2. Then, CN would see the flow state as in expression (8). Thus, the entire data centre becomes multi-homed, providing more resilience.

$$\langle TCP : P_{CN}, P_V, I_{CN}, I_V \rangle \langle ILNP : L_{CN}, L_2 \rangle \quad (8)$$

In this case, no special inter-router link is required for multi-homing – normal Internet connectivity between SBR1 and SBR2 suffices. However, it is quite likely that some sort of secure link, between SBR1 and SBR2 or between H1 and H2, would be desirable to protect the bytes that constitute the VM instance V, as they migrate from H1 to H2 across the sites.

Note that this capability scenario also would support network (site) mobility, as previously described for ILNP [8], [9].

VIII. DISTRIBUTED APPLICATIONS

Extending the capability scenario in the section above, consider now that VMs are all ILNP-capable and distributed across data centres. This may be, for example, a distributed application running across several sites. An example is as depicted in Figure 6. Here VMs V_X to V_Z are each running on separate hosts in separate data centres. If all instances of V are ILNP-capable, then each of the VMs can behave as a mobile server, using the standard mobility-model for ILNP [18]. Each VM can migrate to any other host at another data centre, as required, while maintaining existing communication sessions, either with other VM images that form part of the distributed applications, or with CNs (not shown in the figure).

In this case, mobility can be across networks, could be global, and special inter-router links between SBRs are not needed. Indeed, the SBRs need not be ILNP-capable. If the SBRs are not ILNP-capable, but $V_X - V_Z$ are, then this allows incremental and autonomous deployment of the ILNP-capable VMs. Subject to policy, hosts H1 to H4 would configure inter-host links to carry the moving VM image as required. For example, secure sessions might be set up directly between the hosts as described in the section above.

We believe that this capability scenario gives great flexibility and allows the site-networks and application deployments to

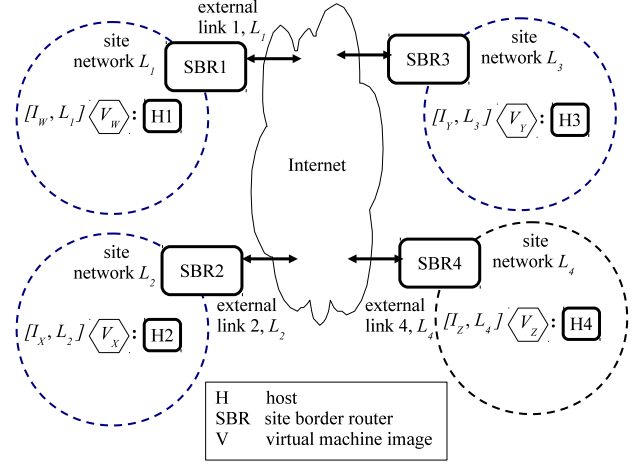


Fig. 6. Distributed applications: multiple instances of VM images form part of a distributed application. If $V_X - V_Z$ are ILNP-enabled systems, they can move to any other site as required while maintaining communication sessions with any other instances of $V_X - V_Z$. Each VM image will maintain its respective NID value, $I_X - I_Z$, but change its respective $L64$ value as it moves from site to site. In this use case, the SBRs need not be ILNP-enabled.

have a very high-degree of independence from the network service provisioned at each site.

If multi-homing is required, then ILNP supports multi-homing for individual hosts, with suitably modified transport protocols. For example, the case for multi-path TCP in data centres [28] becomes extremely attractive when ILNPv6 handles the ‘address management’ aspect of multi-path TCP – a single NID value is bound to multiple $L64$ values so an instance of V could use multi-path TCP.

IX. ILNP-CAPABLE CORRESPONDENT NODE (CN)

The scenarios presented above would require ILNP-capable end-systems in the data centre. However, it is possible to have either ILNP-capable hosts or IP-only hosts as the remote CN and these options are discussed in this section.

A. IP-only CN

For the remote host – the correspondent node (CN) – the availability of ILNP would be beneficial. However, for the first two scenarios listed above, the state of the transport flows remains fixed from the viewpoint of the CN. So, it the benefits of ILNP VM mobility could be employed in datacentres even while CNs remain normal IP hosts

Consider expression (3) and (5), which represent, respectively, the state for the transport connection as the VM moves hosts. We can see that this state remains constant as the VM moves, i.e. NID and $L64$ values do not change. So, in principle, an IP-only CN could be supported *that has not been modified to support ILNP*, achieving the ILNP benefits for mobility in the data centre by upgrading only the data centre systems, and not requiring client systems to upgrade. This is made possible, essentially, by the ILNP Locator re-writing capability used in the conjunction with the simple mapping between a Local

locator value (e.g. L_L , or L_{L1} and L_{L2}) and a global Locator value at the site-network.

B. ILNP-capable CN

If the remote host is ILNP-capable, then extensive additional functionality is possible, compared with an IP-only host. For example, above in Section VII-B, flexible multi-homing is possible for the data centre if the remote CN is ILNP capable. Whilst multi-homing is possible for IP today, it relies on extra state being inserted in the *global* IP routing tables, creating additional technical and administrative overhead, advertising more details about one's network deployment to potential adversaries, as well as creating an additional trust dependency. If V and CN wished to use end-to-end packet-level security, then ILNP-enhanced IPsec between V and CN still could be used.

If the whole site were mobile, i.e. site-network 2 of Figure 4 represents the site's new 'position' (e.g. a Humvee or aircraft carrier with a new uplink), then site-mobility is also possible for the data centre [9].

Just as V can send Locator Update (LU) messages to the CN when V moves, if the CN were mobile, then the CN would send LU messages to V. Thus, the data centre also could support mobile remote systems.

More details of the possible features are highlighted in previous work. For functionality and security issues for the site-network, see [9], [10], [20]. For functionality and security issues for the remote system (e.g. CN), see [8], [11].

X. CONCLUSION

We have described several use-cases to demonstrate how to enable secure and agile virtual machine (VM) image mobility based on the use of the Identifier Locator Network Protocol (ILNP). There are use cases to support when the the correspondent node (e.g. client system) is not ILNP-enabled. When the CN is ILNP enabled, the mobility can operate with enhanced functionality, e.g. with IPsec and multi-homing.

There are additional security benefits for the site-network, mobile VM images, and for the CNs when ILNPv6 is used throughout, e.g. we show a use-case with distributed, mobile VM images running across multiple datacentres. Our proposal includes mechanisms that:

- Support continuity of communication sessions for VM images across mobility and migration scenarios.
- Support multi-homing and other traffic management options from the site-network.
- Eliminate restrictions on network design, i.e. network connectivity between data centre sites can be bridged or routed and reside on completely different IP subnetworks.
- Support wide-area network mobility, as well as local-area mobility.

We have described use-cases and capability scenarios which would use ILNPv6 deployment to give extremely flexible VM mobility in the data centre. The capability scenarios cover situations where only the data centre hosts are ILNP enabled, when data centre hosts and correspondent nodes are ILNP

enabled, and when the the data centre hosts and applications are ILNP enabled.

REFERENCES

- [1] W. Kash. DISA's new focus: supporting mobility. *Government Computer News*, February 2011.
- [2] C. Clark et al. Live Migration of Virtual Machines. In *Proc. NSDI'05 - 2nd Conf. on Networked Systems Design & Implementation*, pages 273–286, Berkeley, CA, USA, 2005. USENIX Association.
- [3] VMware Inc. VMware vsphere. Product Brochure VMW-BRO-VSPHR5-USLET-101, VMware Inc., Palo Alto, CA, USA, 2011.
- [4] Aileen Black. The Virtual Armed Forces: US Military Turns to Virtualization, August 2011.
- [5] Juniper Networks. Solutions For Deploying Server Virtualization In Data Center Networks. White paper 2000349-003-EN, Juniper Networks, Mountain View, CA, USA, Sep 2011.
- [6] Cisco Systems. Scalable Cloud Networking with Cisco Nexus 1000V Series Switches and VXLAN. White paper C11-685115-01, Cisco Systems, San Jose, CA, USA, Mar 2012.
- [7] T. M. Takai. Cloud Computing Strategy. Technical report, Chief Information Officer, Dept. of Defense, USA, Jul 2012.
- [8] R. Atkinson, S. Bhatti, and S. Hailes. Harmonised Resilience, Security and Mobility Capability for IP. In *Proc. IEEE MILCOM 2008*, Nov 2008.
- [9] D. Rehunathan, R. Atkinson, and S. Bhatti. Enabling Mobile Networks Through Secure Naming. In *Proc. IEEE MILCOM 2009*, Oct 2009.
- [10] R. Atkinson and S. Bhatti. Site-Controlled Secure Multi-homing and Traffic Engineering for IP. In *Proc. IEEE MILCOM 2009*, Oct 2009.
- [11] S. Bhatti and R. Atkinson. Integrating Challenged Networks. In *Proc. IEEE MILCOM 2011*, Nov 2011.
- [12] A. Fischer et al. Wide-Area Virtual Machine Migration as Resilience Mechanism. In *Proc. SRDSW 2011 - 30th IEEE Sym. on Reliable Dist. Systems Workshops*, pages 72–77, Oct 2011.
- [13] R. Bradford et al. Live Wide-Area Migration of Virtual Machines including Local Persistent State. In *Proc. VEE'07 - 3rd Intl. Conf. on Virtual Execution Environments*, pages 169–179, New York, NY, USA, 2007. ACM.
- [14] T. Li. Recommendation for a Routing Architecture. RFC 6115, IRTF, Feb 2011.
- [15] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson. Host Identity Protocol. RFC 5201, IETF, Apr 2008.
- [16] D. Meyer. The Locator Identifier Separation Protocol (LISP). *Internet Protocol Journal*, 11(1), Mar 2008.
- [17] R. Atkinson, S. Bhatti, and S. Hailes. Evolving the Internet Architecture Through Naming. *IEEE JSAC*, 28(8):1319–1325, Oct 2010.
- [18] R. Atkinson, S. Bhatti, and S. Hailes. ILNP: Mobility, Multi-homing, Localised Addressing and Security Through Naming. *Telecommunication Systems*, 42(3):273–291, Dec 2009.
- [19] R. Atkinson, S. Bhatti, and S. Hailes. A Proposal for Unifying Mobility with Multi-Homing, NAT, and Security. In *Proc. MobiWAC'07 - 5th ACM Intl. Wkshp. on Mobility Mgmt. and W'less Access*, Oct 2007.
- [20] R. Atkinson, M. Lad, S. Bhatti, and S. Hailes. A Proposal for Coalition Networking in Dynamic Operational Environments. In *Proc. IEEE MILCOM 2006*, Oct 2006.
- [21] B. Carpenter. Architectural Principles of the Internet. RFC 1958, Internet Architecture Board, Jun 1996.
- [22] R. Hinden and S. Deering. IP Version 6 Addressing Architecture. RFC 4291, IETF, Feb 2006.
- [23] T. Aura. Cryptographically Generated Addresses (CGA). RFC 3972, IETF, Mar 2005.
- [24] T. Narten, R. Draves, and S. Krishnan. Privacy Extensions for Stateless Address Autoconfiguration in IPv6. RFC 4941, IETF, Sep 2007.
- [25] R. Hinden, S. Deering, and E. Nordmark. IPv6 Global Unicast Address Format. RFC 3587, IETF, Aug 2003.
- [26] W. Wasserman and F. Baker. IPv6-to-IPv6 Network Prefix Translation. RFC 6296, IETF, Jun 2011.
- [27] C. Perkins, D. Johnson, and J. Arkko. Mobility Support in IPv6. RFC 6275, IETF, Jul 2011.
- [28] C. Raiciu et al. Improving Datacenter Performance and Robustness with Multipath TCP. In *Proc. ACM SIGCOMM 2011*, pages 266–277, 2011.