Analysing the Security Threats against Network Convergence Architectures

Patroklos Argyroudis, Robert McAdoo, Stephen Toner,
Linda Doyle and Donal O'Mahony
Centre for Telecommunications Value-chain Research
University of Dublin, Trinity College, Ireland
{Firstname.Lastname}@tcd.ie

ABSTRACT

Current research trends in networks and telecommunications suggest that future architectures will aim for the convergence not only of fixed and mobile infrastructures, but also of different network layer technologies. Proposals like OCALA from UC Berkeley and our own TRANSIT can be used to converge different architectures and network layer overlays while supporting legacy applications. In this paper we analyse and assess the security threats against such systems. Our contribution constitutes the first step in designing security models for convergence architectures as part of their design and deployment phases, rather than as retrofitted mechanisms.

1. INTRODUCTION

The Internet is currently experiencing a huge change in the way that users are communicating over it and using provided services. Unstructured peer-to-peer protocols and overlay network architectures have been proposed, and to a lesser extent deployed, to address limitations of today's Internet. Furthermore, the vision of utilising the Internet for providing Voice-over-IP (VoIP) services and converging mobile networks with fixed telecommunications and data networks is quickly becoming a reality. Even today users are able to access Internet services via their mobile phones, and communicate in the reverse direction as well.

Since it is unlikely that a single network architecture will become prevalent, the research community has developed a number of proposals that aim to allow the interoperation between them and converge them to a unified system. In this paper we analyse the security threats against the OCALA and TRANSIT network convergence architectures. Although we specifically look at these two systems as examples, our results can be generalised to address other current and future similar proposals. Our contribution constitutes the first step in designing security models for convergence architectures as part of their design and deployment phases, rather than as retrofitted mechanisms. Security solutions must be designed concurrently with the basic systems, since retrofitted solutions may leave unpredictable and undetectable vulnerabilities.

The rest of this paper is structured as follows. Section 2 analyses the need for network convergence and presents two proposals that satisfy this need; OCALA and TRANSIT. In section 3 we present our main contribution; a detailed threat modelling of OCALA and TRANSIT using the attack trees methodology that can be generalised to other network convergence systems as well. Section 4 explores the different existing processes for assessing the threats we have identified and gives an example application. Section 5 presents related work on the subject. We conclude in section 6 by summarising the results of our work and our contributions.

2. NETWORK CONVERGENCE

The basic goal of network convergence architectures is to enable the interconnection of different networking technologies. The common element between all such technologies is that they offer end-to-end packet delivery services. The primary examples from the traditional wired Internet are IPv4 and IPv6. Newer approaches follow the structured peer-to-peer, or overlay, communication paradigm in order to enhance several aspects of IP, like routing, mobility and security. Examples of such protocols are the *i*3 [1] and RON [2], among many others. At the same time we are witnessing increasing interest for networking protocols that address the edges of the Internet and are able to operate in a completely ad hoc manner without the need for dedicated routers. In the mobile telecommunications world the current trend is to enable subscribers to use Internet services and to offer services to Internet users as well.

The research literature clearly demonstrates that no one networking solution will become dominant as they all address different problems and have different operational requirements and assumptions. Therefore, it is evident that uniform connectivity will become, if it is not already, the main focus of networks and telecommunications research efforts. Convergence architectures must provide strong security guarantees, such as authentication and confidentiality among others, in addition to their main goal of uniform connectivity if they are to be adopted and used.

2.1 OCALA

Overlay Convergence Architecture for Legacy Applications (OCALA) [3] from UC Berkeley was designed with two primary goals in mind; to enable communications between hosts and users that operate in different overlays, and to allow legacy applications to be used in existing overlays without the need for code changes and recompilation.

In order to implement its stated goals, OCALA defines a new layer in the TCP/IP stack below the transport layer and above the overlay network layer. This is called the Overlay Convergence (OC) layer and is divided into two sublayers, the overlay dependent sublayer (OC-D), which interacts with the underlying overlays, and the overlay independent (OC-I) sublayer, which interacts with legacy applications, see Figure 1. Although Figure 1 illustrates the bridging of two overlays (*i*3 and RON), OCALA also supports the bridging of IPv4/IPv6 networks. The OC-D sublayer follows a modular design supporting a number of different overlay and traditional network routing protocols. The OCALA layer

0-7695-2876-7/07 $25.00 © 2007 IEEE
DOI 10.1109/IAS.2007.14

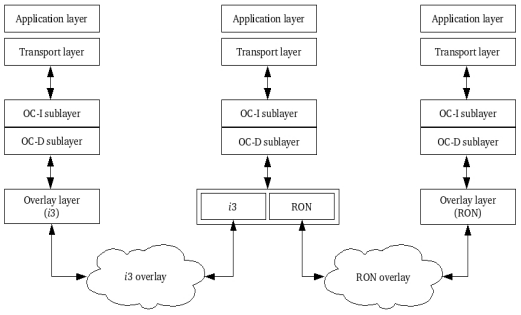

Figure 1: The OC layer in the TCP/IP stack allows connectivity between different overlays.

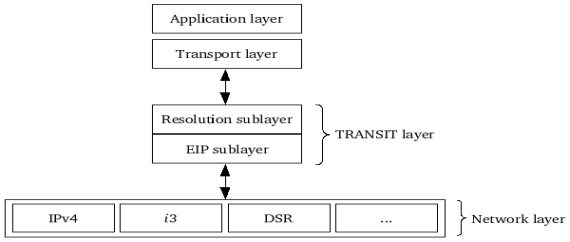

Figure 2: A network layer agnostic, TRANSIT-enabled TCP/IP stack.

operates as a proxy intercepting IP packets coming down the TCP/IP stack from legacy applications and transmitting them over the overlay network that the OC-D sublayer has loaded the required module for.

Identification of end hosts in OCALA is accomplished via the use of DNS-like names [3]. Instead of the traditional DNS hierarchy, OCALA identifiers follow a three-level dot-separated structure. The suffix denotes the overlay type, the middle part identifies the overlay instance and the prefix specifies the overlay-specific name. For example, the name `foo.bar.ron` represents the host `foo` on the `bar` instance of a RON type overlay. The identifier's type part is also used by the OCALA layer in order to select the appropriate OC-D module.

2.2 TRANSIT

Like OCALA, the TRANSIT architecture, developed at the University of Dublin, Trinity College by the Centre for Telecommunications Value-chain Research (CTVR), defines a new layer in the TCP/IP stack between the transport and the network layers (see Figure 2). TRANSIT has been designed as a fully backwards (IPv4) compatible solution to the current address depletion, mobility and decentralisation problems that the Internet faces. In this paper we will only concentrate on its uniform connectivity features; the interested reader is referred to [4] for further details.

TRANSIT's basic component is the Extended IP (EIP) sublayer. EIP uses a triplet of 32-bit addresses. At each level the use of 32-bit addresses allows compatibility with existing network layers, enabling incremental device updates. This three level hierarchy provides a balance between flexibility and compatibility with existing devices. The bottom level enables a local network to locally utilise the entire 32-bit address space, without enforcing structured addressing internally within this network. This allows a network to internally use an addressing scheme that is not dependent on the location or point of attachment of that network to the conventional Internet. The middle level, allows for the arbitrary interconnection of these networks in a mesh topology, even in the absence of access to fixed infrastructure, and further expands the address space to ensure availability for future networking. The top level of the hierarchy enables continued interaction with unmodified IPv4-based devices and applications.

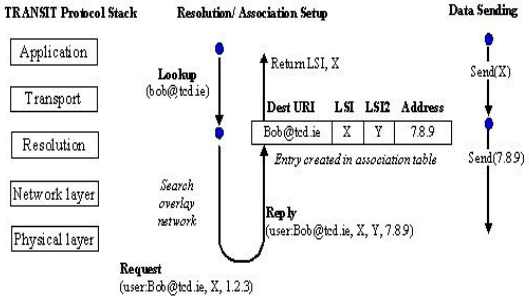

Figure 3: TRANSIT connection establishment and data transfer example.

TRANSIT also consists of a name resolution layer. This is responsible for performing the resolution between a persistent identifier and a transient locator. If identifiers are not to be overloaded with location-dependent information then a scalable resolution mechanism is required which does not rely on hierarchy. As in the case of OCALA, this can be a data storing overlay algorithm, like for example the ones presented in [5]. In aiming for a unified form of identification, TRANSIT proposes the use of Universal Resource Identifiers (URIs) for endpoint identification. A URI is hierarchically structured and assigned, at least in the URI-scheme part. This enables scalable delegation of assignment to different authorities. Although assigned in a hierarchical manner, TRANSIT's directory service treats URIs as flat identifiers.

2.2.1 *Connection Establishment and Data Transfer*

TRANSIT hides location-dependent addresses from applications and upper layers within the protocol stack. However, to avoid rewriting existing applications to create sockets based on URI strings it introduces a 32-bit representation of this identity, termed the Locally Significant Identifier (LSI). This acts as a "placeholder" in existing protocols and APIs. The LSI is generated locally within each host, during the resolution process. This is also exchanged with the peer during *association setup* (Figure 3), which is required to create the necessary state information within devices. By exchanging LSI information, an entity may discover its representation used within the peer device, which may be important for example for security reasons. Resolution is now required to create the necessary association state between communicating entities. An application uses the LSI in place of an IP address. The socket API then uses this LSI rather than the IP address. Thus the semantics of higher level connections are changed without having to modify applications.

3. THREAT ANALYSIS

Threat modelling (or analysis) is essential in order to help us develop a security model than can focus on protecting

Figure 4: General threat categories for network convergence architectures.

against certain threats and manage the related assumptions. One methodology to discover and list all possible security attacks against a system is known as *attack trees*. To create an attack tree we represent attacks against a system in a tree structure; the attack goals as root nodes and the different subgoals necessary to achieve them as their leaf nodes [6].

Figure 4 presents the general threat categories we have identified against network convergence architectures, namely attacks on the network processes responsible for packet routing, intra-realm routing threats and name resolution threats. These categories are divided further. Name resolution threats are different depending on the employed resolution mechanism. As we have discussed, convergence architectures rely either on traditional DNS or overlay protocols for resolving names to routable identifiers. The security problems of DNS have been analysed extensively in the past and threat models have been proposed [7]; therefore we do not analyse them again in this paper. Instead, we focus on attacks against overlay resolution mechanisms. Intra-realm routing security has also been extensively studied. We refer the interested reader to [8] for a threat analysis against traditional routing protocols, and to [9] for a survey of ad hoc routing security. In this paper we analyse the threats introduced by the convergence of different routing realms (inter-realm routing threats).

During the development of the model we have identified that several attacks lead to other attacks which we have previously included and analysed. These are represented in the tree as identical nodes in different locations. A node that appears in more than one location of the tree has the same sub-tree everywhere. In order to avoid including the same sub-tree multiple times in the model we have used arrow icons to denote duplicate nodes. Therefore, a node marked with an arrow means that it exists somewhere else in the tree and if that node has a sub-tree this sub-tree is included only once.

Although we have not explicitly included privacy threats in our model, a lot of the enumerated attacks may result in reducing the privacy of participating users. We consider these to be outside the scope of our study.

3.1 Inter-realm Routing Threats

Since network convergence systems aim to enable uniform connectivity between distinct realms that utilise different routing protocols, attackers can focus on disrupting this process in order to achieve their goals. Figure 5 illustrates the various threat subcategories we have identified for inter-realm routing.

3.1.1 Eavesdropping

Eavesdropping on the communication channel is a common threat for networking systems and protocols that do not employ cryptographic mechanisms to protect the confidentiality of the exchanged messages. An attacker exploits the assumptions of the underlying networking technology, like for example the broadcasting nature of Ethernet, and receives on the local interface the entire traffic of the current subnet. If the confidentiality of the protocol messages is not protected, the attacker has access to all the information included in them. Inter-realm routing data can be captured and analysed in order to discover the structure of realms, important nodes such as gateways, or even the number of nodes that participate in a given realm [10].

3.1.2 Identity Impersonation

In identity impersonation attacks a malicious entity takes advantage of the absence of end-to-end authentication mechanisms in a network protocol and assumes the identity of another entity. In the context of inter-realm routing an attacker can assume the identity of a particular important network node like a realm gateway and send false routing information to other gateway nodes. We analyse such routing table poisoning threats in the next paragraph. Furthermore, an identity impersonator can force a peer node to disclose its routing table by sending specially constructed signalling requests to it.

3.1.3 Routing Table Poisoning

Routing table poisoning attacks allow a malicious entity to insert false data into the routing tables of legitimate participating nodes. When the target of such attacks is a realm relay node then the attacker can influence inter-realm routing to implement a series of other attacks. Non-optimal routes can be presented as optimal and replace legitimate entries with the goal of redirecting traffic through a specific path for eavesdropping or other purposes. Also, by spoofing routes that use the same paths an attacker can implement a difficult to detect denial of service attack, forcing all traffic to go through specific nodes, creating artificial bottlenecks. More traditional denial of service attacks will be analysed in the next paragraph. Another possible attack is the isolation of network realms by inserting false data that portray broken connections, or non-existent link failures. Realm isolation can also be accomplished by creating loops between two or more different routing realms, separating them from the rest of the network.

3.1.4 Denial of Service

Denial of service threats aim at disrupting the requirement of availability, which can be defined as the problem of enabling systems to perform their advertised services in a timely manner. An attacker can generate false routing data and forward them to a relay node which will try to save all of them to its routing table, possibly deleting legitimate data to do so. Although the results of such attacks are devastating since they render all inter-realm routing impossible, they are easily detected. A much more subtle way to perform denial of service attacks is to carefully create excessive amounts of signalling traffic. A motivated attacker can cause the performance of a relay node to degrade to unusable levels by mimicking heavy routing traffic. By periodically advertising routes from falsely generated sources the attacker can hide from the relay node the fact that an attack is taking place.

3.1.5 Replay

Replay attacks involve the transmission of previously captured legitimate information. Such information can be whole routing packets, or just the authentication information included in them appended to false payloads. The latter can be used by an attacker to undermine poorly designed security solutions and furthermore as the first step of identity impersonation attacks. By replaying previously captured packets verbatim an attacker can break the synchronisation between

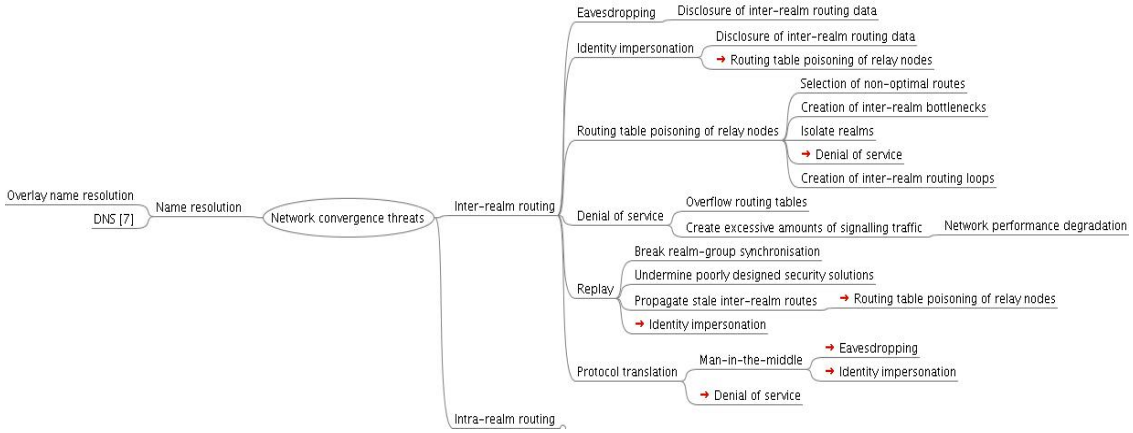

Figure 5: Inter-realm routing threats.

different routing realms. Stale routes can be presented as fresh and routing tables of relay nodes can be polluted with them.

3.1.6 Protocol Translation

When gateway nodes relay traffic between two realms that use different network layer protocols they have to perform the necessary translation between them. This is accomplished by removing the network layer header of the source realm and replacing it with the header of the destination realm. However, in the case that a transport layer security mechanism like TLS/SSL is used, the process of translation becomes problematic and opens new attack avenues. TLS channels are established in an end-to-end manner based on the identifiers of the two communicating hosts. These identifiers are the ones employed at the convergence layer. When protocol translation is required, the two end hosts are identified within their respective network realms with different identifiers than the ones used at the convergence layer. Therefore, the TLS channel cannot be established in an end-to-end manner; the host that performs the protocol translation, i.e. the relay node, effectively behaves as a man-in-the-middle. If it is not trusted by both end hosts to perform the translation and participate in the TLS channel then it can eavesdrop on the connection and impersonate any one of them. Furthermore, untrusted nodes that translate between network protocols are in a position to cause denial of service on every routing path that they participate in.

3.2 Overlay Name Resolution Threats

Overlay name resolution algorithms are often proposed to address the need for uniform naming management in network convergence architectures. Since the first step of establishing a connection between two parties that need to communicate is the resolution of an identifier to the corresponding routing token, an attacker can take advantage of the assumptions made by the resolution algorithm and implement a number of threats. Moreover, when an entity wants to join a network usually the first step is to make an initial contact with an existing network participant and through them to request a particular name. If this process is left unprotected an attacker is able to prevent new entities to join the network. Figure 6 illustrates the threats we have identified in this area and the following paragraphs analyse them in detail.

3.2.1 Eavesdropping

If the confidentiality of the exchanged messages is not protected by a security mechanism employed at the overlay resolution layer, then malicious entities that participate in the network are free to gain access to their payloads. This usually constitutes the first step of further attacks, since it allows an attacker to have a clear view of the network, its topology and details regarding its participants. An attacker may wish to obtain information about the queries made by a host, or group of hosts. In an overlay network resolution system, they are afforded this opportunity since each host in the network participates in routing. If the contents of messages are not protected as they are routed then attackers may passively monitor queries that are routed through them.

3.2.2 Invalid Messages

Invalid message attacks compromise the integrity of the name resolution system. An attacker may forge results for queries that it has received, or eavesdropped on. In a situation where they do not have the ability to eavesdrop on queries they may attempt to predict queries and reply to them optimistically in the hope that they successfully cause an invalid result to be returned. In the absence of mechanisms to protect message integrity a malicious node may alter packets that they are responsible for routing, causing false results to be returned, or denying nodes the ability to successfully perform queries. Attacks of this nature may be mitigated by the inclusion of a mechanism to protect data integrity. However, the semantics of the name resolution system itself could be taken advantage of to provide invalid results. These systems may offer no guarantee that a result will be returned, and often there is no differentiation between a query that failed and one for which there is no valid result. A malicious node may exploit this, and forward packets for queries that it wishes to let through, giving the impression that it is functional. However, when it receives a query that it wishes not to return a result for, it may choose not to forward that packet. The node performing the query may then interpret this silence as a lack of results. Finally, the lack of a node authentication architecture means there is no mechanism in place to ensure that those who insert data into the network have the authority to do so. This may allow nodes to overwrite data that belong to other nodes.

Figure 6: Overlay name resolution threats.

3.2.3 Denial of Service

Denial of service occurs when an attacker or a group of attackers prevent one or more legitimate nodes from availing of the services offered by the network. In an overlay network this may occur when a node is either prevented from inserting itself in the network, or when it is prevented from participating in the network. Malicious nodes wishing to prevent others from entering the network may attack either the joining algorithm, or the identifier allocation mechanism. Nodes joining the network rely on the existence of at least one bootstrap node. As this node is already a part of the network, it may act as a bridge between the node wishing to join and the rest of the network. This will then allow the joining node to query the network for the information it needs to join. If, however, the bootstrap node acts maliciously then it may prevent a node from inserting its information into the network. The bootstrap node must therefore be a node that is trusted, and should provide a method for joining nodes to authenticate it.

As in any naming system, a mechanism for secure allocation of identifiers is a requirement to prevent malicious nodes from stealing identities. In an overlay network node insertion may fail if its identifier already exists in the network. This may be targeted against a single node or a general attack, known as a Sybil attack [11], against a group of nodes. Service may be disrupted for nodes already inserted into the network if an attacker prevents them from querying, their messages from being routed, or results of queries from being returned. Nodes may be prevented from querying, or have replies from legitimate nodes go unnoticed by deploying attacks discussed in the previous paragraph. Alternatively a node may spoof a *node leave* message, which would subsequently result in that node being deleted from the routing tables of the other nodes in the network, rendering it unable to perform queries. Also, if nodes may distort the routing tables of other nodes then routing may be prevented. Finally, a flooding of the network by a malicious node would result in nodes being overloaded with queries, and unable to differentiate between legitimate and malicious queries.

4. THREAT ASSESSMENT

Threat assessment methodologies can be used to indicate the severity of an identified attack. Tregear has identified two main approaches for analysing threats and assessing their impact, namely quantitative and qualitative ones [12]. To assess the threats we have previously identified in the area of network convergence we rely on the qualitative approach and specifically we use the metrics proposed in [13], initially used to assess X.509 Public Key Infrastructure (PKI) related compromises.

Our goal is to use these metrics to evaluate the possibility of violations of security properties like confidentiality, integrity, authentication, authorisation, availability and non-repudiation in regard to the threats we have identified in section 3. Table 1 presents our assessment. We acknowledge that the process of assigning impact levels to threats is subjective, primarily depends on the application environment and therefore may be challenged. However, our analysis provides a guideline according to which specific convergence architectures can be evaluated. For example, as can be seen in Table 1, the eavesdropping threat has different impact levels in the inter-realm routing and the overlay name resolution categories. We believe that generally the information exchanged in the latter category is of higher value than in the former, but this may not be the case in particular environments.

5. RELATED WORK

Although there exist in the literature previous threat modelling efforts for overlay name resolution algorithms, to our knowledge we are the first to present a detailed analysis of network convergence threats. In [14] a security analysis of IP-based peer-to-peer distributed hash tables is presented. The authors make the assumption that an attacker can generate packets with arbitrary contents, but she can only examine packets addressed to herself. This means that in their adversarial model an IP address can be used as a weak form of node identity, an assumption that cannot be made in ad hoc or other non IP-based networks. Security problems in peer-to-peer networks have also been examined by Wallach [15]. His focus is on intra-realm overlay routing algorithms and approaches to secure these. A more complete representation of threats was given by Keely [16]. Although he focuses on threats to wireless mobile networking for e-business applications, several of the attacks he identified are applicable to

Table 1: Impact of identified threats on security properties.

Threat	Security property						Impact level			
	Confidentiality	Integrity	Authentication	Authorisation	Availability	Non-repudiation	High	Medium	Basic	Rudimentary
Inter-realm routing										
Eavesdropping	X								X	
Identity impersonation	X	X	X	X	X	X	X			
Routing table poisoning of relay nodes		X			X			X		
Denial of service					X				X	
Replay	X	X	X	X	X	X	X			
Protocol translation	X	X	X	X	X	X	X			
Overlay name resolution										
Eavesdropping	X							X		
Invalid messages		X	X	X	X	X	X			
Denial of service					X			X		

network convergence architectures.

6. CONCLUSION

The unification of different networking technologies is becoming a reality due to the increasing need for interoperation between fixed, mobile and peer-to-peer communication protocols. Although a number of architectures have been proposed that are able to implement the required convergence, the associated security problems and the new threat avenues opened by this trend have not received similar attention. In this paper we have presented an initial attempt to rectify this. Our contribution constitutes the first step towards the design of security solutions that are part of the proposed unification systems, and not as retrofitted mechanisms. By examining two existing proposals, namely OCALA and TRANSIT, we have presented a threat model that can be generalised and applied to other similar systems. We are currently investigating the applicability of existing security models and protocols in providing solutions for the identified threats.

7. ACKNOWLEDGEMENTS

This material is based upon works supported by Science Foundation Ireland under grant number 03/CE3/I405.

8. REFERENCES

[1] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet Indirection Infrastructure," in *Proceedings of ACM SIGCOMM 2001*, 2001, pp. 149–160.

[2] D.G. Andersen, H. Balakrishnan, M.F. Kaashoek, and R. Morris, "Resilient Overlay Networks," *Operating Systems Review*, vol. 35, no. 5, pp. 131–145, 2001.

[3] D. Joseph, J. Kannan, A. Kubota, K. Lakshminarayanan, I. Stoica, and K. Wehrle, "OCALA: an Architecture for Supporting Legacy Applications over Overlays," in *Proceedings of 3rd USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI '06)*, 2006.

[4] S. Toner, *TRANSIT: Adapting the Internet for Mobile and Ad hoc Operation*, Ph.D. thesis, School of Computer Science and Statistics, University of Dublin, Trinity College, 2006.

[5] D. Doval and D. O'Mahony, "Overlay Networks: A Scalable Alternative for P2P," *IEEE Internet Computing*, vol. 7, no. 3, pp. 2–5, 2003.

[6] B. Schneier, "Attack Trees," *Dr. Dobb's Journal*, pp. 21–29, 1999.

[7] D. Atkins and R. Austein, "Threat Analysis of the Domain Name System (DNS)," RFC 3833, 2004.

[8] A. Barbir, S. Murphy, and Y. Yang, "Generic Threats to Routing Protocols," IETF Draft draft-ietf-rpsec-routing-threats-07, 2004.

[9] P. Argyroudis and D. O'Mahony, "Secure Routing for Mobile Ad hoc Networks," *IEEE Communications Surveys and Tutorials*, vol. 7, no. 3, pp. 2–21, 2005.

[10] S.M. Bellovin, "A Technique for Counting NATted Hosts," in *Proceedings of 2nd ACM SIGCOMM Internet Measurement Workshop (IMW'02)*, 2002, pp. 267–272.

[11] J. Douceur, "The Sybil Attack," in *Proceedings of 1st International Workshop on Peer-to-Peer Systems (IPTPS'02)*, 2002, pp. 251–260.

[12] J. Tregear, "Risk Assessment," *Information Security Technical Report*, vol. 6, no. 3, pp. 19–27, 2001.

[13] Federal Bridge Certification Authority, "X.509 Certificate Policy for the Federal Bridge Certification Authority (FBCA)," `http://www.cio.gov/fpkipa/documents/fbca_cp_09-10-02.pdf`, 2002.

[14] E. Sit and R. Morris, "Security Considerations for Peer-to-Peer Distributed Hash Tables," in *Proceedings of 1st International Workshop on Peer-to-Peer Systems (IPTPS'02)*, 2002, pp. 261–269.

[15] D.S. Wallach, "A Survey of Peer-to-Peer Security Issues," in *Proceedings of 2003 International Symposium on Software Security (ISSS'03)*, 2003, pp. 42–57.

[16] D. Keely, "A Security Strategy for Mobile E-business," Tech. Rep. GSOEE213, IBM Global Services, 2001.