1. Abstract

Cybersecurity for data networks is in its infancy while attackers on networks are becoming increasingly sophisticated. The necessary widespread use of wireless networks provides more vulnerabilities. Network routing is key to a functioning network; once compromised, it can be difficult to recover. For many years, network practitioners have worked on methods to protect that routing through authentication of the updates passed in the network. The missing piece has been a usable, protectable key management system. This proposal uses recent advances in the creation of locally controlled and administered hierarchical web-of-trust certificates to provide a managed secure identity for routers that can be used to protect network routes from attack and misconfiguration.

This proposal is the first phase of creating an authenticated routing infrastructure. The work involves adapting advances in evidentiary trust to a link state routing protocol, developing naming for certificate chains and an approach to use the certificates in link state updates. This phase is expected to result in a report and a design for a prototype to be added to open source protocols in a later phase. These results will be made publicly available through open source code and discussions and presentations at standards bodies and with router vendors and network operators. A successful approach should create opportunities to the proposer in contracts with government and commercial organizations. This could result in market opportunities for other organizations such as network management tool and router vendors.

The use of evidentiary trust is expected to have other applications, but its application to the routing infrastructure can have nearer- term impact on securing networks important to the government.

2. Identification and Significance of the Problem

The objective of this proposal is to develop a locally controlled and administered secure identity for data network routing elements and an architecture for its use to authenticate router updates. This is a first step in a larger objective of developing routing algorithms that are resilient to attack and capable of steps to self-recovery. Routing ensures that data is delivered to destinations properly and in a timely fashion. A functioning routing algorithm is key to continued network operation in the presence of attacks due to malfeasance as well as misconfiguration. With a secured unique, locally verifiable identity for each routing process in a network, algorithms can be deployed that respond to attack under a variety of conditions. For instance, it is possible to remove routing data of more questionable provenance when an attack is detected, to identify specific sources of doubtful information, or to operate in modes where all information must be absolutely trusted. A flexible, incrementally deployable approach is preferred; one that can interoperate with other routing elements.

Methods for authenticating the information that routers exchange have been proposed and modified over two decades, but still lack deployment despite general agreement that public key based signatures of routing information is the most secure and general approach. Part of the reason has been computational and memory resources in routers, but these have become less

important with technological advances. More importantly, there has been the lack of a method to provide a secure identity to individual routers and processes that can be used to uniquely identify the provenance of the individual updates. Up till now, approaches have proposed trusted entity key management and this entity creates vulnerabilities. The INFOSEC Research Council Hard Problems list has noted the difficulty of Global-Scale Identity Management (#1) and Information Provenance (#6). By applying new research in this area [SNC] to create secure names to ensure the provenance of routing information, it is possible to address Hard Problems #3 (Availability of Time-Critical Systems) and #4 (Building Scalable Secure Systems) in the context of network routing.

1.1 Background of the Problem

As data networks have become a more ubiquitous and critical part of government and civilian institutional operations, their attractiveness as a target for attack is increasing. Many attacks on have taken place and there are likely more in the wings. At the same time, it is clear that advances in networking are not keeping up, with several areas needing attention. [SCS]

With the widespread use of wireless networking, particularly by the warfighter and first responder communities, the difficulty of uniquely distinguishing routing updates from different network domains is necessary to protect against malfeasance and misconfiguration. When access to a network domain cannot be limited by physical barriers, intradomain routing is just as vulnerable as intradomain routing, and often less well defended.

1.2 Current Approaches are Lacking

There have been many proposals to secure routing protocols. These proposals fall into three categories: securing Border Gateway Protocol (BGP), the de facto inter-domain routing protocol in the Internet, securing link state routing protocols, typically used within an Autonomous System (AS) such as a ISP or an enterprise network, and securing routing protocols for mobile ad hoc networks (MANET) and wireless sensor networks (WSN).

None of these proposals have been widely deployed. Secure BGP (S-BGP) [SBGP], the frontrunner among the BGP proposals, is deemed too heavyweight and requires a trusted public key infrastructure [NC]. Other more lightweight approaches have been proposed, such as Secure Origin (so-BGP) [SOBGP]. Both proposals uses public-private key pairs to authorize prefix ownership and authenticate prefix announcements. Currently, networks use best common practices such as route filtering to protect themselves against routing attacks.

Secure link state routing proposals date to Perlman's seminal work two decades ago [IBR] and have continued through research proposals and both Experimental and Standards Track documents in the Internet Engineering Task Force (IETF), the standards body for the Internet Protocol [RFC2154, RFC5304] for over a decade. The standards are ambiguous on the address ranges allowed on an advertising router. It is up to the network to specify what is allowed but the standards provide little guidance. Enhancements to achieve efficient message authentication have also been proposed [CHE] but have the same problems.

Therefore, although algorithms to authenticate routing are specified, these have not been deployed due to problems of an appropriate key management structure. Simply put, if names must be resolved through a central entity, the system is open to attack by attacking that trusted entity. Consequently, when authentication is used at all, it involves the use of MD5 algorithms to authenticate particular neighbor links of the topology, but cannot be used to authenticate the routing updates themselves and unambiguously determine the signer. This lack of an appropriate key management protocol to secure routing with public-private key pairs has been a stumbling block for sometime. The 1997 approach in [RFC2154] notes that: "Any router can send out a public key and claim to be a given router, so the public key itself provides no assurance of the actual identity of the sender. This assurance must be provided by a Trusted Entity. The Trusted Entity (TE) is a system that generates certificates for routers." However, the issue of how to establish and defend the Trusted Entity was not addressed. That this is still an issue can be seen from the more recent [RFC5304]: "If and when a key management protocol appears that is both widely implemented and readily deployed to secure routing protocols such as IS-IS, a different authentication mechanism that is designed for use with that key management schema could be added if desired."

The proposals for MANET [PH] and WSN [KW] are new routing protocols and do not have backward compatibility with existing, deployed routers in mind. Although most of these protocols are based on distance vector or link state routing algorithms, they are tailored to the specifics characteristics of MANET and WSN, making it difficult to incrementally deploy them in an existing network.

1.3 Important Characteristics of a Viable Solution

Using entirely new routing protocols can have appeal, but these new protocols will also lack the years of deployment, observation, and threat experience of currently deployed protocols. Thus, any new protocols need to be deployed in a limited way, starting from a "sandbox" before becoming protocols of choice in critical networks. For this reason, changes to existing protocols that preserve strengths while plugging up known holes are preferred for the shorter term.

For any protocol update, it will not be feasible in most networks to have a "flag day" and swap out all the routers, so a solution should be incrementally deployable, useful when only partially deployed, reasonably easy to add to existing routing protocols, backwards compatible. Further, a good solution should be like public-private key pair encryption: it is possible to be entirely open about the algorithm but still impervious to attack.

1.4 Emerging Work can be Exploited

There is much active work in network security research on architectures to provide names for data that can be used to uniquely identify that data's provenance. [SNC] gives a review of this work and provides a novel solution of evidentiary trust and the secure binding of names to content that permits the use of user-friendly names and does not require a central authority to disperse names. The focus of that work is on naming and retrieving general data content, but the naming and identity notions appear to be suitable for use in network routing to solve the long-term problems of use of public key authentication in this context. This work is based on [SDSI] approach to create user-friendly namespaces creating transitive trust through a certificate chain

that validates locally controlled and managed keys, rather than requiring a global Public Key Infrastructure (PKI). Certificates, or keys, are created that have a particular context in which they should be utilized and trusted rather than conferring total authority.

The fundamental idea for this proposal is to give each router a secure identity and to add the ability to sign and authenticate route updates to existing routing algorithms. Simply put, the problem this proposal is addressing is to create a reliable way for a network to determine who (that is, which network elements) should be believed in order to ensure that routing data is valid. This proposal does not address privacy of routing data. This approach both solves key distribution and makes it easy to distinguish the authority of any particular routing process to update any particular domain topology. Current approaches are vulnerable to attacks from man-in-the-middle. This is exacerbated by today's mobile, wireless, and ad hoc networks. Secure identity for particular router and routing process makes it possible to disambiguate confusing or intentionally misleading updates that are not relevant.

Innovation is primarily in creating trust models and associated namespaces that are relevant, effective and viable in typical network operational scenarios. Opportunity is expected in follow on contracts with commercial and government entities for architectural support and for enhanced routing algorithms using the secured routing topology.

3. Phase I Technical Objectives

The objective of this proposal is to develop a secure identity for Internet Protocol routers and an architecture for its use to authenticate the provenance of router updates. This is a first step in a larger objective of developing routing algorithms that are resilient to attack and capable of steps to self-recovery. The enumerated technical objectives of Phase I are:

1. Investigate creation of an appropriate namespace for routers.
2. Select routing protocol for prototype (IS-IS or OSPF) and an open source software platform.
3. Design for the backwards compatible addition of signed route updates using this namespace.
4. Analysis framework for evaluation of attack resistant algorithms employing the secure ids.

4. Phase I Work Plan

Phase I will be restricted to design and analysis of the namespace for network routing messages and certificates, design for the addition of secure naming and signing to an open source routing protocol implementation, and preliminary analysis on the use of signed routes for networks under attack. Prototyping and local identification of and response to potential attack conditions will be undertaken in Phase II.

4.1 Investigation and design of namespace approaches

There are two somewhat distinct subtasks here. Starting with the naming and data verification approach of [SNC], a naming approach for network routing elements must be designed with a straightforward naming hierarchy that includes organization, network domain, specific router, and routing process instance. These names *become* the secure identity that is used to validate the

provenance of routing information. In addition the approach to configuration and management of these names needs to be developed.

Note that "key" is often interpreted as something obtained through PKI. In the following, "key" is the locally controlled and managed certificate chain. A certificate is a name plus a key plus a signing with another key that is trusted. [SDSI] specifies how a certificate chain can be hooked up. "Locally" has particular contexts for use and is limited to appropriate actions, unlike PKI keys which have a more global authority, which can let them be used in inappropriate ways. Following [SNC], the name of the data is created so that it can be used to identify which public key should be used to check the signature of the named data, including the certificate chain, so that the receiving router can decide whether or not to trust the signature with respect to the specific data. There can be a number of different "publishers" which can each act as a certification authority. The receiver needs to have a certificate chain that enables it to trust the publisher of any particular data and to check whether the data is from the context within the receiver trusts that publisher. There are two separate questions a router must be able to answer when a routing protocol message is received:

1. Was this message signed by the certificate/key it claims to be signed by?
2. What reason do I have to trust this certificate/key?

Only the receiver can determine if this is a message it wants to accept: is it valid and relevant with an acceptable provenance for its purpose. (e.g., "I received this message that is a routing update, but is it from a peer in my network?"). The publisher/originator of the data has a context (determined from the name and certificate chain) of a particular organization, a particular network domain and a particular routing protocol and and perhaps interface, thus the receiver only trusts the corresponding certificates for that information (e.g., don't trust an adjacent network's updates even if it has the same parent organization).

This can be difficult to visualize so we present an example, using a strawman naming convention. BigCo has a network running separate OSPF protocols in different domains. The relevant ones are its nationwide Core and the San Francisco area SFpop.

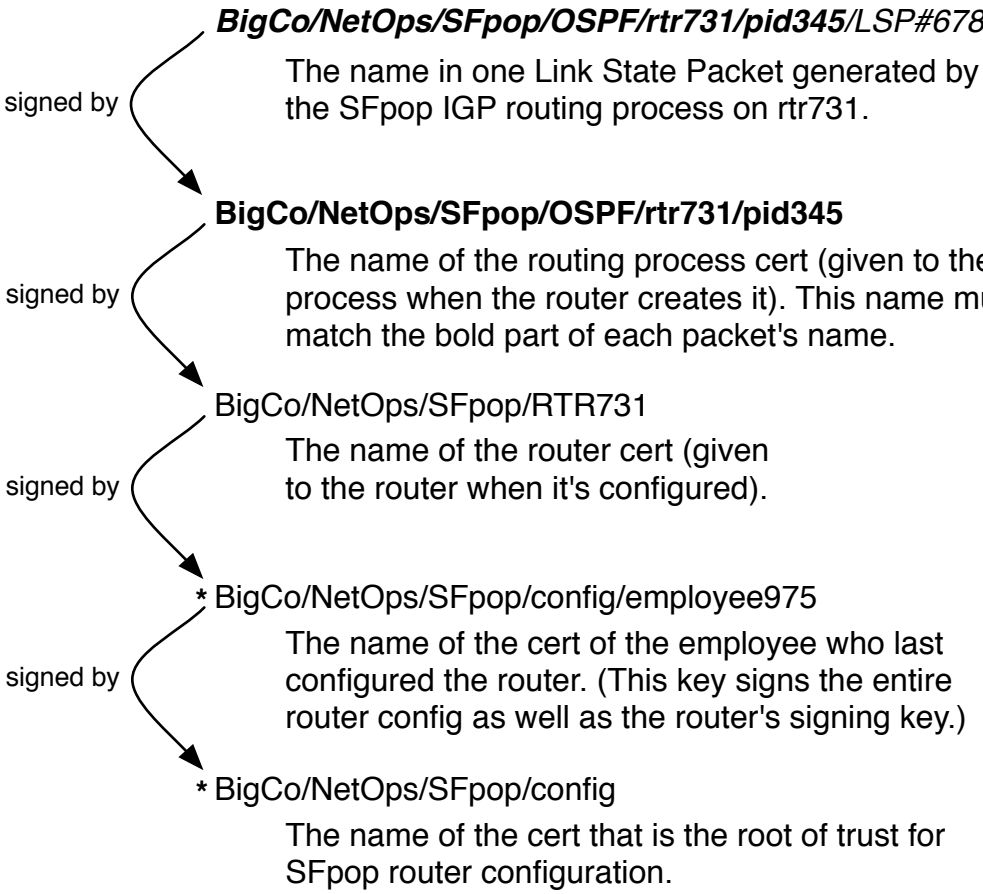

Figure 1: Example Naming and Certification Chain and Configuration

The router rtr 731 sits at the border between Core and SFpop. When rtr731 was initially configured, it was given its own signing certificate BigCo/NetOps/SFpop/RTR731, signed by one of the employees authorized to configure SFpop routers, employee #975. Note that keys in the configuration chain (marked with an * in the figure) must be held very securely. Rtr731 can create certificates for its processes, here for process id 345. The resulting certificate gets added to its LSPs as the key to be used to check its signing. In figure 2, the certification chain for Router 731 and another router in the same domain is shown. When the LSP is received by Router 64, it wants to check if it should trust the signer, rtr731. Since rtr731's certificate is signed by a chain that rtr64 has, trust is established.

Rtr 731 also has processes in Core. A non-peer (to SFpop) router instance signs with the key: BigCo/NetOps/Core/OSPF/rtr731/pid987 which had to be configured by someone with configuration authority in Core (e.g., BigCo/NetOps/Core/config/employee52). When LSPs are received on the same address for different routing instances, the differentiation can be made by the identity resolving process and the LSP handed to the correct process. If an LSP is received for which the receiver has no way to determine trust, then it can be discarded, avoiding routing problems due to both attack and misconfiguration.

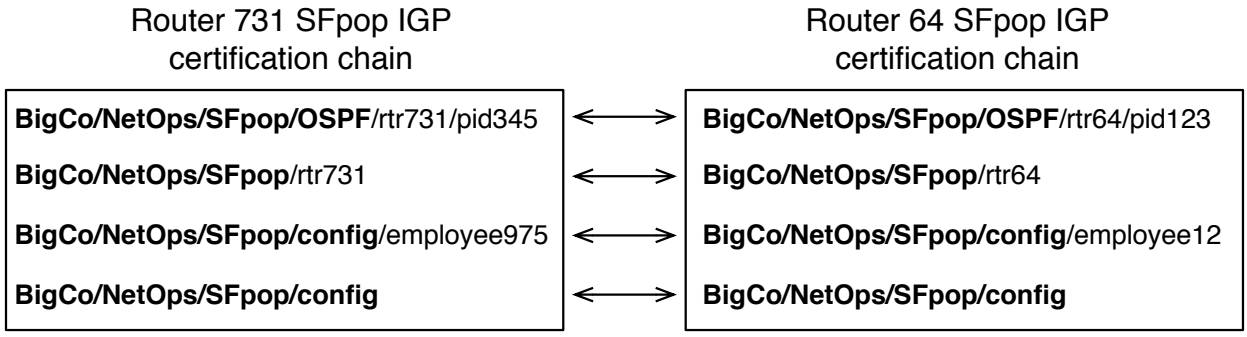

Figure 2: Example Certification Chain for Two Router Peer Processes

To summarize, initial keys/certificates are distributed "out of band" during router set up. These are used to obtain more specific key/certificates and create more specific key/certificates. Trust in keys is established by creating a naming of keys which parallels the naming of the routing data, or link state advertisements. Trust in a particular entity results in trust in the data. When a particular router gets its initial key, that is all it needs to create the richer namespace that gets down to a particular routing instances: e.g., organization/box/protocol/interface/process. This permits use of the same multicast address but lets announcements be specific to particular routing process (solves a problem VLANs often used to address).

4.2 Selection of prototype implementation platform

It is necessary to select the routing protocol to be used for the initial design (and eventual prototype). The prototype routing protocol will be a link state IGP, either IS-IS or OSPF. The choice of IPv4 or IPv6 will be made during the work cycle. Available open source implementations of IS-IS and OSPF are available from several sources, including: Quagga (OSPF, IS-IS), XORP (OSPF), Vyatta (OSPF), and GNU Zebra (OSPF). One of these will will be selected, along with the particular protocol and version.

4.3 Design for introduction into routing protocol

With a preliminary approach and a specific target for a prototype, the design can be detailed for that platform. How to add the use of the certificates and their storage by each router will be addressed. When a message is received on an interface, there's a wrapper with a name plus a signature that can be checked before it is handed upward to appropriate process. When an update is sent, add the wrapper with name plus signature. All routers employing this new option will be expected to sign their routing data; signing not optional though checking is. Preliminary approaches on how to *use* these signatures will be specified (e.g. random checks at some interval after non-deterministic checks on first appearance).

The most obvious use is to ensure management traffic addressed to the router is authenticated and that route updates are authenticated.

4.4 Analysis framework for evaluation of benefits

A framework for evaluating benefits of the signed updates will be developed. Using documented routing protocol threats (e.g., [RFC4593]), we will determine which are candidates for solution

using authenticated updates. For example, Sniffing and Traffic Analysis threats are *not* candidates for solution by authentication. Approaches will be sketched out, a preliminary approach to evaluation of the approach, and a ranking of importance in terms of cost-benefit carried out.

4.5 Schedule

Months 1,2: Investigation and preliminary design of appropriate name space and certificate chain
Months 3: Design generic approach to algorithm for receiving and checking certificate chains
Months 4,5: Select platform for prototype design and complete specific design
Months 6: Complete documentation and develop analysis and evaluation framework

4.6 Deliverables for Phase I

The deliverables for Phase I will include monthly reports, a detailed final report on the design and analysis, and any code base used for data analysis, including documentation.

5. References

[SCS] National Research Council and National Academy of Engineering, *"Toward a Safer and More Secure Cyberspace"*, National Academies Press, 2007.
[NC] O. Nordstrom, and C. Dovrolis, "Beware of BGP Attacks", ACM Computer Communication Review, 2004, VOL 34; NUMB 2, pages 1-8.
[SBGP] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)", IEEE Journal on Selected Areas of Communications, vol 18, no 4, Apr 2000.
[SOBGP] R. White, "Securing BGP: so-BGP", Cisco Systems, Sep 2003.
[MSIP] E. Winjum and T Berg, *"Multilevel Security for IP Routing"*, Proceedings of MILCOM 2008, pp. 1-8.
[RRR] R. Morera, A. McAuley, L. Wong, *"Robust Router Reconfiguration in Large Dynamic Networks"*, Proceedings of MILCOM 2003, pp 1343-1347.
[NSMD5] Naik et. al., *"Improving and Maintaining Network Security Using MD5 Algorithm"*, Global Journal of Computer Science and Technology, Vol 9 Issue 5, January 2010, p 103.
[RFC2154], S. Murphy, M. Badger, and B. Wellington, *"OSPF with Digital Signatures"*, RFC 2154, Experimental, June 1997.
[RFC4593] A. Barbir, S. Murphy, and Y. Yang, *"Generic Threats to Routing Protocols"*, RFC 4593, October 2006.
[RFC5304] T. Li and R. Atkinson, *"IS-IS Cryptographic Authentication"*, RFC 5304, Standards Track, October 2008.
[CHE] S. Cheung, "An Efficient Message Authentication Scheme for Link State Routing", Annual Computer Security Applications Conference, Dec 1997.
[KW] C. Karlof and D. Wagner, "Secure routing in wireless sensor networks: attacks and countermeasures", First International Workshop on Sensor Network Protocols and Applications, 2002.
[PH] P. Papadimitratos and Z. J. Haas, "Secure Routing for Mobile Ad hoc Networks", SCS Communication Networks and Distributed Systems Modeling and Simulation Conference, Jan 2002.
[IBR] R. Perlman, Interconnections: Bridges and Routers, Addisson-Wesley, 1992.

[SDSI] Rivest and Lampson, *"SDSI - A Simple Distributed Security Infrastructure"*, MIT Technical Report, April 30, 1996
[BST] V. Jacobson, C. Alaettinoglu, and K. Poduri, *"BST - BGP Scalable Transport"*, NANOG 17, February, 2003
[CCN] V. Jacobson et. al., *"Networking Named Content"*, ACM CoNEXT, December, 2009
[SNC] D. Smetters and V. Jacobson, *"Securing Network Content"*, Palo Alto Research Center Tech Report, http://www.parc.com/content/attachments/securing-network-content-tr.pdf, October 2009.
[Wang98] F. Wang and S.Wu, *"On the Vulnerabilities and Protection of OSPF Routing Protocol"*, IEEE 7th International Conference on Computer Communication and Networks, 1998.
[SPKI] E. Ellison et al, *"SPKI Certificate Theory"*, RFC 2693, September 1999.