

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
16 March 2006 (16.03.2006)

PCT

(10) International Publication Number
WO 2006/028651 A2

- (51) International Patent Classification: Not classified
- (21) International Application Number: PCT/US2005/028698
- (22) International Filing Date: 12 August 2005 (12.08.2005)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/933,696 3 September 2004 (03.09.2004) US
- (71) Applicant (for all designated States except US): INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): LORTZ, Victor [US/US]; 7716 S.W. OVIATT DRIVE, BEAVERTON, OR 97007 (US). WALKER, Jesse [US/US]; 10992 N.W. APPELLATE WAY, PORTLAND, OR 97229 (US). HEGDE, Shriharsha [IN/US]; 21588 SW RANKIN COURT, BEAVERTON, OR 97006 (US). KULKARNI, Amol [IN/US]; 6476 NE ALDER ST, HILLSBORO, OR 97124 (US). TAI, Tsung-Yuan [KR/US]; 2496 NW 141ST PLACE, PORTLAND, OR 97229 (US).
- (74) Agents: MCCRACKIN, Ann M. et al.; SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A., P.O. Box 2938, Minneapolis, MN 55402 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.



WO 2006/028651 A2

(54) Title: A DEVICE INTRODUCTION AND ACCESS CONTROL FRAMEWORK

(57) Abstract: In an embodiment, a method includes registering applications and network services for notification of an out-of-band introduction, and using the out-of-band introduction to bootstrap secure in-band provisioning of credentials and policies that are used to control subsequent access and resource sharing on an in-band channel. In another embodiment, an apparatus implements the method.

A DEVICE INTRODUCTION AND ACCESS CONTROL FRAMEWORK

5 Technical Field

Various embodiments described herein relate to the field of wireless networks generally, including communications and network devices generally, including apparatus, systems, and methods for allowing network access from one device to another.

10 Background

Various devices may operate wirelessly, communicating over radio frequency links or infra-red links in a network. Network security requires either physical isolation or cryptographic mechanisms or some combination of the two to assure that devices accessing the network are authorized to do so. Utility and/or security of the network may suffer as the techniques require increasing user attention and intervention to assure device verification.

Present network security either requires physical isolation or cryptographic mechanisms bootstrapped via some out-of-band (OOB) mechanism (or some combination of the two). Some out-of-band techniques require substantial user attention and intervention (such as entering passwords or examining digital certificates, names, or key identifiers). The user experience of some of these techniques is poor. To make matters worse, devices are increasingly multifunctional, with a variety of embedded applications and sharable resources that the user would like to selectively make available to other devices. These applications and services may operate at different layers of the software stack (communication stack), and credentials applicable at one level are not necessarily usable at another layer for a variety of reasons. Since each credential type may (and typically does) use a different method for trust bootstrapping, the user today is faced with the prospect of performing multiple separate trust bootstrapping operations to configure mutual trust between a given pair of devices. For devices with limited user interfaces, the complexity of current trust bootstrapping methods is such a serious obstacle that many devices simply omit security altogether or include only simple security built

around extremely weak measures such as a 4-digit PIN (personal identification number).

Brief Description of the Drawings

5 Embodiments of the present method and apparatus may be best understood by referring to the following description and accompanying drawings, which illustrate such embodiments. In the drawings:

FIG. 1 depicts an embodiment of the present method and apparatus.

10 FIG. 2 depicts in very general terms an embodiment of the present method for bootstrapping applications and services.

FIG. 3 depicts in general terms an embodiment of the present method.

FIG. 4 illustrates an embodiment of a concept of layering in a layered framework.

FIG. 5 illustrates an example embodiment of an introduction ceremony.

FIG. 6 depicts an embodiment of an alternative secure introduction embodiment.

15 FIG. 7 depicts a scenario where the server initiates the connection.

FIG. 8 depicts a scenario where the client initiates the connection.

FIG. 9 depicts two devices that want to conduct secure communications using a network that is also shared by many other devices/users.

20 FIG. 10 shows an embodiment of a detailed message exchange between the two devices.

Detailed Description

In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific
25 details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the understanding of this description.

Embodiments of the present method and apparatus are directed to initial trust bootstrapping and provisioning of domain and/or role-based credentials.

30 FIG. 1 depicts an embodiment of the present method and apparatus that addresses the problems in the prior art with a device introduction framework 100 that is extensible to bootstrap multiple applications and services at different layers in the communications stack.

The FIG. 1 embodiment of the present method and apparatus may have the following architectural elements: an out-of-band introduction mechanism (e.g., wire, USB token, RFID, infrared, or short-range radio) and data structure 102 for conveying keying material, names for user display, and in-band addressing information for the next phase of the credential and policy bootstrapping process; an extensible data structure 104 for conveying across the in-band channel additional credentials, policy, and configuration data associated with a domain to which one of the newly-introduced devices belongs; a protocol 106 that confirms knowledge of the out-of-band keying material and derives a secure in-band channel for exchanging the data structures used to provision long-lived credentials and domain policy; and a processing framework 108 that maintains credentials and associated domain policy and makes them available to the appropriate trusted application and driver components on the device platform.

FIG. 2 depicts in very general terms an embodiment of the present method for bootstrapping applications and services at different layers in a communications stack. This embodiment may include: registering applications and network services for notification of an out-of-band introduction (201); and using the out-of-band introduction to bootstrap secure in-band provisioning of credentials and policies that are used to control subsequent access and resource sharing on an in-band channel (202). These credentials may be domain credentials. The domain credentials issued to a newly added device may include implicit role-based authorization so that no additional access control configuration is needed after the introduction is complete.

In an embodiment of the present method and apparatus the domain credentials issued to a newly added device may include implicit role-based authorization so that in most cases no additional access control configuration is needed after the introduction process is complete.

FIG. 3 depicts in general terms an embodiment of the present method for an introduction process that provides a device introduction framework that is extensible to bootstrap multiple applications and services at different layers in a communications stack. This embodiment may have: registering applications and network services for notification of out-of-band (OOB) introductions (301); OOB introducing at least one of shared secret, in-band addressing info, and/or public key(s) (302); receiving, by the applications, callback with OOB channel characteristics, domain identifier, and

peer device name, and providing an opportunity to pass policy data to peer over an in-band channel (303); conducting at least one of in-band credential establishment and in-band policy exchange (304); establishing subsequent secure interaction based on identities and keys (305); and consulting, by multiple software components, with introduction
5 framework to determine domain and use appropriate credentials and policies (306).

In some embodiments of the present method and apparatus, applications running on a platform during subsequent operations may consult the processing framework for credentials and policy associated with a currently active domain. Since the framework may maintain a centralized policy store using an extensible data structure,
10 service providers that own or manage devices can easily provision policy for a variety of different applications and services by updating the policy data for their domain. Globally unique domain identifiers may be used to prevent policy updates from one domain owner from overwriting policy associated with a different domain. At runtime, the processing framework may provide the appropriate policy to
15 applications and services according to a currently active domain.

When digital devices need to securely share resources or communicate across networks (especially wireless networks), they must first establish a virtual security perimeter based on cryptographic keys. Of course, a device has no intrinsic understanding of its owner's goals or the trustworthiness of another device on the network. The security
20 mechanisms and trust relationships used by the device must first be configured to reflect the desires of the device owner. The device owner in turn must have some means of conveying his intent that the devices work together. Translating the owner's intent into device configuration is variously known as provisioning or trust bootstrapping. Historically, many computing systems and applications have neglected the problem of
25 trust bootstrapping and have simply assumed that initial shared secrets will be configured through some unspecified method. However, secure bootstrapping turns out to be a very challenging. Security vulnerabilities in trust bootstrapping processes and unwillingness of many users to follow good practices in managing credentials have done much to undermine the value and robustness of network security implementations.

30 Some embodiments of the present method and apparatus may include a framework for securely establishing initial trust relationships between wireless devices and for communicating those trust relationships to various operating system, device driver, and application software components. The intent is not to subsume or replace other security or

access control mechanisms but rather to augment them by improving their trust bootstrapping.

Part of the framework is a user experience that is tangible and intuitive, based on physical operations such as pressing designated “introduction” buttons, moving devices close together, and/or physically connecting them to transfer introduction data from one device to another. The level of security achieved via this introduction process will depend upon the physical properties of the so-called Out-of-Band (OOB) channels and the cryptographic capabilities of the two devices.

It is well known that security has the following properties: the more complex a system the less likely it is to be secure, and the security strength of a system is limited by its least-secure element. System security is like a chain – only as strong as its weakest link.

Because of these properties, embodiments of the present framework are structured in well-defined layers, with simple mechanisms and protocols at the bottom and more complex mechanisms and applications higher up. This way, the lower layers can be made simple and sufficiently well understood to have a high degree of confidence in their integrity. Higher layers built on top of this foundation can stand or fall according to their individual security characteristics without compromising each others integrity.

FIG. 4 illustrates an embodiment of this concept of layering in a layered framework 401. The core elements may be represented by boxes 400 (OOB invitation), 402 (in-band identification), and 404 (in-band credential registration). Given these foundational elements, other applications and network security layers (App1, App2, Radios, etc.) may be constructed on top (the AppN boxes correspond to the trust bootstrapping and configuration aspects of those applications).

Of course, it is not necessary that all software layers base their trust bootstrapping on this foundation. It is always possible for applications or network services to take care of all aspects of their own security and configuration on their own. However, embodiments of the present method and apparatus provide an opportunity for applications or service layers to take advantage of a common introduction process and thereby improve the user experience.

The following define various core elements:

Device. Devices are independent physical or logical entities that communicate with other devices across a network. Devices consume and/or provide resources to other

devices on a peer-to-peer or master/slave basis. Devices may have heterogeneous capabilities in terms of computational power and I/O channels.

Domain. A Domain is a set of one or more devices that recognize a common authority over membership and access to network resources.

5 Enrollee. An Enrollee is a device seeking credentials to access a Domain resource. Once an Enrollee obtains a credential, it becomes a Member.

Member. A Member is a device that holds credentials of one or more Domains.

Resource Monitor. A Resource Monitor is a device mediating access to some resource. It requires a Credential to decide whether to permit a Member to access the
10 resource.

Credential. A Credential is a data structure issued by a Registrar to an Enrollee, allowing the latter to authenticate itself as a Domain Member. The Credential in turn allows the Member to authenticate a Resource Monitor and a Resource Monitor to decide whether to permit the Member to access its resource.

15 Registrar. A Registrar is a device trusted by one or more Resource Monitors to administer a Domain. It issues an invitation to an Enrollee, establishes an Enrollment Security Association with it, verifies that the security association is with the correct Enrollee, and finally issues Member Credentials to the Enrollee. The Registrar may be a distinct physical entity from the Resource Monitor, or it may be integrated into the
20 Resource Monitor but logically distinct. The role of Registrar may be played by any or all of the Devices in a Domain, depending upon the Domain policy.

Invitation Protocol. An Invitation Protocol is an optional out-of-band protocol between a Registrar and an Enrollee. This protocol transfers parameters from a Registrar to the Enrollee to enable the Enrollee to locate and mutually authenticate with the
25 Registrar using the Identification Protocol.

Identification Protocol. An Identification Protocol is a two-way in-band protocol between the Enrollee and the Registrar that establishes an Enrollment Security Association.

Enrollment Security Association. An Enrollment Security Association is state
30 shared between the Enrollee and Registrar to protect the provisioning of a Credential.

Verification Protocol. A Verification protocol is an optional out-of-band protocol between the Enrollee and the Registrar that is used after the Identification Protocol and prior to the Registration Protocol.

Registration Protocol. A Registration Protocol is a (logically) three party protocol to assign a Credential to the Enrollee. The protocol operates between the Enrollee and the Resource Monitor, via the auspices of the Registrar.

Embodiments of the present method and apparatus may also involve an
5 introduction ceremony. In general, a ceremony is a collection of coordinated management activities, out-of-band protocols, and in-band protocols, which work together in a prescribed fashion to accomplish some end. An introduction ceremony “introduces” two computing devices that have never before met and have no common root of trust. The ceremony goal is to securely establish credentials, so the pair of devices (and by extension
10 others belonging to the same domain) need not be introduced again.

FIG. 5 illustrates an example embodiment of an introduction ceremony. In this embodiment of the present method and apparatus, the verification protocol is depicted in brackets because it may be omitted if the invitation protocol is used. Transport for the identification protocol between an enrollee and a registrar may use a temporary guest
15 credential, a separate logical or physical network used only for enrollment, or a framework such as 802.1X that permits only limited access to the network prior to domain enrollment. A given domain may include multiple registrars. Devices acting as registrars may also be resource monitors.

Some devices may choose to respond differently to subsequent introduction
20 attempts once they have been successfully enrolled in a domain. For example, after initial enrollment, a device may reject attempts to enroll it in other domains unless this is explicitly authorized by the original domain. Typically, a “factory reset” capability is included in such devices to enable recovery from failure of the domain authority.

The initiation of an introduction may ordinarily be accomplished by some user
25 action such as pressing designated “connect” or “introduction” buttons on the two devices. The introduction needs to reflect the user’s intent, and introductions that happen entirely automatically and without the user’s involvement are likely to violate this principle. Once the user has indicated to the registrar that an introduction is needed, the registrar may prompt the user for the desired role to assign the enrollee. This information will be used
30 in the registration protocol when deciding which credential to issue to the new member. Note that the role assignment is optional.

Execution of an invitation protocol initiates enrollment. The invitation protocol transfers out-of-band parameters between the registrar and the enrollee. These parameters

allow the enrollee to locate and subsequently authenticate the correct registrar using the identification protocol. Although it is possible for an invitation protocol to transfer parameters in either or both directions, the invitation protocol may transfer parameters one-way from the registrar to the enrollee. By limiting the exchange to one way, it becomes possible to implement the protocol with one-way out-of-band channels and simplifies the user experience when using introducer devices that are carried from the registrar to the enrollee.

The specific parameters transferred by the invitation protocol may vary according to the characteristics of the out-of-band channel and the capabilities of the registrar and enrollee. Typical values may include:

- Message type and version – makes the invitation message format extensible;
- In-band network descriptor for the identification protocol (type and/or network identification);
- Registrar's public key identifier;
- Registrar's in-band network address for the identification protocol;
- Temporary in-band network access credential; and
- Identification protocol type and options (may be implicit in message type).

To improve resistance to attacks against the out-of-band channel, especially if an introducer device is used, the registrar may assign an expiration time to an introduction secret and reject it after that time, or to expire the introduction secret after its first use.

Once it has received out-of-band parameters via the invitation protocol, the enrollee may use these parameters to initiate the in-band identification protocol with the registrar. This protocol may mutually prove possession of cryptographic keys by the enrollee and the registrar and establishes an enrollment security association that may be used to protect the registration protocol.

The identification protocol is typically not run across the same out-of-band channel as the invitation protocol because some out-of-band channels have very limited bandwidth, and some are unidirectional. For example, RFID (radio frequency identification) tags have both of these limitations. USB flash drives have ample bandwidth but are unidirectional (unless the user shuttles the drive back and forth multiple times between the devices, which is unworkable). In an embodiment of the present method and apparatus the invitation protocol may only require a one-way distribution of initial network addresses, and public key identifiers, and may be readily adapted to

channels with these limitations. The identification and registration protocols, which typically involve mutual authentication, private key proof of possession, and certificate issuance, may require a bi-directional channel with moderate bandwidth.

In some embodiments, a single round-trip may be performed using the out-of-band
5 channel to establish an enrollment security association for use by the registration protocol or for direct use as a credential for network access. In this case, the identification protocol may be omitted. FIG. 6 illustrates an example of this type of introduction ceremony.

When the verification protocol is used, it transfers information that allows the enrollee and the registrar to verify that the enrollment security association has been
10 established with a designated peer. The verification protocol may typically only be used if the invitation protocol is omitted or if the invitation protocol's out-of-band channel properties or parameters provide insufficient authentication of the peer.

The registration protocol is used to issue credentials to the enrollee. It is protected by the enrollment security association that is established by the identification protocol or
15 by the invitation protocol.

A single introduction ceremony may result in the establishment of multiple credentials of different types. This is because devices often contain multiple applications and networking layers, each of which may require an independent credential. Indeed, it is often beneficial to keep such credentials separate, because shared credentials used
20 concurrently in different contexts are more susceptible to attack than independent credentials.

For security reasons, it is desirable that the introduction ceremony maintains perfect forward secrecy between the invitation protocol and the identification protocol. In other words, if an eavesdropper records the encrypted message exchanges of the
25 identification and registration protocols and subsequently obtains the out-of-band data passed in the invitation protocol, this information should not enable it to compromise the credentials established in the registration protocol. Likewise, later disclosure of data exchanged in the verification protocol should not compromise the credentials. An alternative approach for protecting against disclosure of the invitation secret on the out-of-
30 band channel is to encrypt this data using public keys so that an attacker that eavesdrops on the out-of-band channel cannot thereby discover the secret and compromise the introduction. FIG. 6 illustrates this latter approach.

Embodiments of the present method and apparatus are agnostic to the specific out-of-band channel that may be for the invitation protocol. Nevertheless, some guidelines may be provided to ensure that the out-of-band channel has appropriate characteristics for secure device introduction.

5 Some out-of-band channels are more resistant to eavesdropping than others. For example, physical wires and USB flash drives are more resistant to eavesdropping attacks than are infrared or RFID. Public keys may be utilized for introductions using channels susceptible to eavesdropping.

10 Out-of-band channels such as USB flash drives typically have ample capacity to transmit large quantities of data. Others, such as RFID, may have very limited capacity. In the latter case, a compact binary format of the minimal data may be utilized to establish a connection and bootstrap initial trust on the in-band channel.

15 One characteristics of a good out-of-band channel for introduction is that it allows the user to unambiguously express which two devices are being introduced. For this purpose, channels that ensure physical proximity of either the devices or an introducer device may be utilized. Out-of-band channels may also include a physical point-to-point wire. This has the advantage of resistant to eavesdroppers, high bandwidth, and simple and intuitive to implement and to use. However, this has the disadvantage of being impractical or unusable if devices are far apart, user may be confused about why a wire is needed to
20 set up a wireless device, and may require an extra connector.

 Out-of-band channels may also include an USB flash key or equivalent device. This has the advantage of resistant to eavesdroppers, high bandwidth, and is a fit metaphor of physical "security key". However, this has the disadvantage of an extra device that a user may lose or misplace, extra expense for the key, extra expense for USB host interface
25 and connector on devices.

 Out-of-band channels may also include the use of infrared. This has the advantage of being very inexpensive, and is built-in already too many devices, especially consumer electronics. However, this has the disadvantage of being susceptible to eavesdroppers in the same room, physical proximity is not as precise as wire or USB key, and monitoring
30 infrared channel consumes power. This may be an issue for small battery-powered devices. Aligning directional infrared ports is not always easy or even possible (if devices cannot be moved or are far apart). This disadvantage may be overcome using a portable IR-based introducer device.

Out-of-band channels may also include near field communications and RFID (radio frequency identifier). This has the advantage of being fully wireless, no external connector needed, no need to precisely orient devices being introduced. However, this has the disadvantage of being susceptible to eavesdroppers, has limited channel data capacity
5 for RFID, and currently less widely available than other channels mentioned previously.

Access control is the process by which a device accepts or rejects a request by a peer device to perform some action or obtain some data. Access control may be rooted in an access control list that maps credentials of other members or domain authorities onto corresponding privileges. If domain authorities are used, then credentials issued by that
10 domain may convey a subset of the domain privileges to the credential holder. When a device receives a request, it may first obtain the credential(s) of the requester. It may then determine based on the credential and the access control list whether or not to grant the request.

If role-based access control is used, the following sub-categories of members are
15 often useful. These roles may be assigned directly in the issued credentials or may be associated with the credentials using a data structure in the resource monitor. An administrator is a member with access permissions that are a superset of ordinary members. A guest is a member with access permissions that are a subset of ordinary members.

There are two main alternatives to enable cross-domain access control, where a
20 Resource Manager requires Members to present Credentials issued by more than one Domain to obtain access. One possibility is that cross-certification may be performed (i.e., a Domain authority issues a certificate for another Domain authority). Such cross-Domain credentials must accurately reflect the desired semantics to avoid granting unintended
25 permissions. Another possibility is that a Resource Manager may recognize multiple Domains and require multiple independent Credentials be provided by a Member to gain access.

In some circumstances, a user may wish to remove security associations and/or revoke Member access granted through a prior introduction. A straightforward method of
30 accomplishing this is to unilaterally delete data associated with that introduction using a device's user interface. For devices such as laptop computers, this may be readily accomplished due to the large display and rich I/O capabilities. For headless devices, the problem is more difficult. Therefore, a special mode of the Registration Protocol may be

used to revoke Security Associations and/or Credentials associated with a prior introduction. This function may be initiated by either the Registrar or the Member. Note that one Device cannot ultimately force another Device to forget a key. All it can do is signal its own intention to revoke a key and request that the peer do likewise.

5 The efficacy of the revoke operation will depend upon the type of credential and access control mechanism. For example, if access control is solely based on checking the root signature of certificates signed by a Domain, then if the device holding the Credential refuses the revocation request, it may be necessary to wait for certificate expiration to implicitly revoke access. Alternatively, a certificate revocation list may be maintained
10 (although maintenance of such lists has proven to be a substantial burden in PKI deployments). An additional option is to maintain separate access control lists that may be edited to revoke privileges previously granted to specific Domain Members. In this case, the Credentials would continue to be valid, but the privileges conveyed by them would be removed.

15 As described in earlier sections, the process of introducing a device to a domain is accomplished via a combination of protocols, viz. Invitation Protocol, Identification Protocol, Verification Protocol and Registration protocol.

 Some devices may be able to support relatively heavyweight Public Key encryption, whereas others may only have the capability of supporting lightweight
20 symmetric key encryption. Accordingly, the Introduction Ceremony may have the following flavors: Symmetric Key only Introduction; Symmetric Key + Public Key Introduction; and Public Key only Introduction.

 The following describes how the Introduction Ceremony may be performed using Symmetric Key and/or Public Key methods.

25 The Invitation Protocol may use a one-way OOB channel from the Registrar to the Enrollee. Ceremonies using the Symmetric Key and the Symmetric Key + Public Key use a common Invitation Protocol.

***Registrar* → *Enrollee*: *RegistrarAddress* | *Key* | *KeyId* [|
hash(PublicEncryptionKey)]**

30 Here, *RegistrarAddress* is the address of the registrar on the InBand channel. The address may vary based on the transport used to convey the Introduction Protocol. It may be a MAC Address or an IP Address with an optional TCP port number.

Key is a randomly generated Symmetric Key used by both parties to secure their session. It is composed of two separate keys, the Key Confirmation Key and the Key Encryption Key.

Key: KCK | KEK

5 *KeyId* is used by the Registrar to identify the Key from other Keys it generates for other Enrollees. The *KeyId* is randomly selected, making it infeasible to compute the Key from the *KeyId* i.e. the *Key* → *KeyId* mapping is One Way.

10 *Hash(PublicEncryptionKey)* is the hash of the Public Encryption Key used by the Registrar. This fingerprint is used by the Enrollee to verify the Registrar's public key in the Identification Protocol.

Capability permitting, the Registrar should always include the hash of the Public Encryption Key in the Invitation message. This provides the Enrollee with the option of using either the Symmetric Key Only or Symmetric Key + Public Key method for the Identification protocol.

15 If a device does not support an independent hardware channel with the appropriate characteristics for secure transmittal of Invitation data, then verification of the data must be performed through the Verification Protocol (for instance, including visual comparison of key identifiers or device serial numbers)

20 The Invitation Protocol may provide the Enrollee with the Registrar's In-Band address and other parameters needed to finish the Introduction Ceremony. The Enrollee may now connect to the Registrar In-band and begin the Identification Protocol. Typically, the Registrar provides the Enrollee with the option of using either Symmetric Only or Symmetric + Public Keys to secure their session. Given sufficient capability, the Enrollee should always choose the latter option since it has better security features than using only
25 Symmetric Keys.

The Identification Protocol therefore has two possible flavors as follows.

The Symmetric Key Only Introduction Protocol is the case where the Enrollee decides to use only symmetric keys to secure the session with the Registrar. The Introduction Protocol is as follows:

30

Enrollee → Registrar: M1 = Random1 | KeyId

Registrar generates a random key K1 and computes $\alpha = \text{KeyWrap}(\text{KEK}, K1)$

Registrar → Enrollee: $M2 = \text{Random1} \mid \text{Random2} \mid \text{KeyId} \mid \alpha \mid \text{MAC}(\text{KCK}, \text{Random1} \mid \text{Random2} \mid \text{KeyId} \mid \alpha)$

Enrollee generates a random key K2 and computes $\beta = \text{KeyWrap}(\text{KEK}, \text{K2})$

5 **Enrollee → Registrar:** $M3 = \text{Random2} \mid \text{KeyId} \mid \beta \mid \text{MAC}(\text{KCK}, \text{Random2} \mid \text{KeyId} \mid \beta)$

Enrollee and Registrar compute key $K = \text{kdf}(\text{K1}, \text{K2}, \text{KeyId}, \text{K-Length})$

10 In Message M1, the Enrollee generates a random string Random1 and sends it to the Registrar along with the KeyId. The KeyId assures the Registrar that the Enrollee is indeed the recipient of the message sent in the Invitation Protocol.

 Message M2 contains both, Random1 and a MAC computed using the KCK. This assures the Enrollee that it is indeed communicating with the Registrar. Message M2 also
15 contains a second random string Random2 and a random key K1, which forms the first part of the session key K.

 Random2 and the use of the KCK in message M3 serve as further verification of the Enrollee. Also, Random2 assures the Registrar that the Enrollee received message M2 correctly. The key K2 serves as the other half of the session key K.

20 Following this message exchange, the Registrar and the Enrollee compute the session key K using K1 and K2.

 The Symmetric Key + Public Key Introduction Protocol is the case where the Enrollee decides to use both the Symmetric Key and the Public Key to secure the session.

25 **Enrollee → Registrar:** $M1 = \text{Random1} \mid \text{KeyId} \mid \text{hash}(\text{PublicEncryptionKey})$

Registrar generates a random seed S1 and computes $\alpha = \text{KeyWrap}(\text{KEK}, \text{S1})$

Registrar → Enrollee: $M2 = \text{Random1} \mid \text{Random2} \mid \text{KeyId} \mid \alpha \mid \text{PublicEncryptionKey} \mid \text{MAC}(\text{KCK}, \text{Random1} \mid \text{Random2} \mid \text{KeyId} \mid \alpha \mid \text{PublicEncryptionKey})$

30

Enrollee generates a random seed S2 and computes $\beta =$

Encrypt(PublicEncryptionKey, S2 | KeyId)

Enrollee → Registrar: M3 = Random2 | KeyId | β | MAC(KCK, Random2 |

KeyId | β)

5

Enrollee and Registrar compute $K = \text{kdf}(S1, S2, KDK | KeyId |$

hash(PublicEncryptionKey), K-Length)

In Message M1, the Enrollee generates a random string Random1 and sends it to
 10 the Registrar along with the KeyId and the hash of the Public Encryption Key (PEK). The
 KeyId and hash assure the Registrar that the Enrollee is indeed the recipient of the
 message sent in the Invitation Protocol. Also, the hash serves as an indication that the
 Enrollee wishes to use the PEK along with the Symmetric Key to secure the session.

Again, the Enrollee uses Random1 and KCK to verify the Registrar, while the
 15 Registrar uses Random2 and KCK to verify the Enrollee.

Seeds S1 and S2 are used to compute the session key K.

Since the out-of-band channel may be one-way, and since out-of-band channels
 may use a wide range of possible physical transport mechanisms, one aspect of the out-of-
 band introduction stage is the schema, or format of the information. Ideally, the same
 20 information may be conveyed across a wide variety of out-of-band channels, and the code
 in both the sender and the receiver may be largely agnostic to the channel itself. However,
 since some out-of-band channels have very limited bandwidth, it is desirable to define a
 compact format for this data. An extensible (XML) format as well as a compact (binary)
 format may be used for the out-of-band data.

25 FIG. 6 depicts an alternative secure introduction embodiment. In FIG. 6 E
 designates a public encryption key. SSID is the 802.11 network identifier. WPA-PSK is
 the new link-layer security solution for 802.11 home networks defined by Wi-Fi Alliance.

In embodiments of the present method and apparatus a framework for multi-layer
 bootstrapping based on a single introduction may include the following:

30 A device introduction service is created with a user interface and an
 application program interface (API) that applications aware of this service may
 utilize to help bootstrap their trust and policy.

Using the API, an application may register interest in device introduction events.

When the user wishes to perform an introduction, the user may use the introduction service's API to indicate this. The introduction service may perform platform-specific operations to further secure the OOB channel(s) at this point. For example, it may configure the device drivers to authenticate data received across the I/O channel to prevent undetected modification by viruses or other hostile software running on the platform.

When an introduction event occurs, the introduction service may notify registered applications and provides them an opportunity to reserve keying material and to provide application-specific configuration data to corresponding applications that may be running on the peer device. The introduction service may also provide registered applications with an opportunity to display one or more user interface pages integrated with that of the introduction service (possibly as a separate tab on a tabbed dialog or a separate page in a wizard-driven interaction).

Applications that provide configuration data to peers may also supply an application identifier that may be used to extract the application-specific configuration from the introduction exchange. The API, which may be provided by the introduction service running on each device, enables the delivery and lookup of data based on application identifiers. The introduction service may use the secure channel established via the invitation or identification protocols to authenticate and/or encrypt the configuration and policy data sent across the in-band channel in the registration protocol.

If applications or higher-level network services request key material from the introduction service, a description of the complete set of requests may be constructed and delivered as part of the registration protocol. For example, the registration protocol may derive a master session key to protect its messages. This key may be expanded into additional keying material using well-known key derivation functions (e.g., by repeatedly applying a pseudo-random function to the master session key). A mapping between application identifiers and regions in the derived keying material may be exchanged as part of the registration protocol so that in addition to retrieving policy data, application-specific keying material may be obtained from the introduction service. Note that the registration API may allow applications to designate key identifiers to use to encrypt their

session keys and configuration data so that applications can protect their configuration and keys from each other. When the introduction service receives configuration data and derives keys for a given application, it first encrypts this data using the keys of those applications prior to storing it.

5 The Introduction Service's registration API may be used by multiple applications. Each application may indicate an application ID, data to exchange, a local encryption key identifier to protect the data, and zero or more UI plug-ins. When an introduction occurs, the master secret (enrollment security association) used by the registration protocol may be expanded using a pseudo-random function and the application identifier into
10 application-specific session keys. These keys may then be protected from other applications on the platform by encrypting them using the local encryption keys designated during application registration. Likewise, when data corresponding to each application is received from the peer, it may first be encrypted using the application-specific session key before being made available to applications on the platform. In this
15 way, when applications query the Introduction Service to retrieve keys and data associated with their application IDs, they may have at least a degree of assurance that the data has not been tampered with and cannot be read by other unauthorized applications. The actual level of assurance and application isolation may depend upon the inherent security properties of the host hardware and operating system.

20

RADIO-BASED IN-BAND CHANNEL FOR OOB CHANNEL

Network security either requires physical isolation from untrusted machines or cryptographic mechanisms bootstrapped via some out-of-band (OOB) introduction mechanism (or some combination of the two). There are two forms of OOB channels:
25 physical and logical. A physical OOB channel typically employs independent I/O hardware from the main in-band channel(s). A logical OOB channel uses some other mechanism such as human memory or visual comparison of representations of cryptographic keys. Physical OOB channels have significant usability advantages over logical OOB channels, but they do suffer from some drawbacks. In particular, if a device
30 does not already support multiple independent I/O channels, then adding this capability will make the device more expensive to produce. It can also be problematic to ensure that the same physical OOB channel type is supported for each pair of devices that must be introduced to each other.

For these reasons, there is substantial incentive to use the same in-band channel for both trust bootstrapping as well as subsequent secure communication. Existing approaches to this problem basically fall into two categories: 1) use a logical OOB channel. This approach typically requires either trusted third party
5 certification or requires the user to input a shared password or do manual comparison of cryptographic keys displayed by the devices. 2) apply weak measures such as pressing a button that puts a device in a promiscuous mode to receive an unprotected key over the in-band channel. Method 1 either requires expensive PKI infrastructure or requires the user to correctly perform relatively complex tasks such as typing an
10 identical and secure password into each device. Method 2 is only pseudo-security. It offers no meaningful protection against an attacker with access to the in-band channel at the time of introduction. However, embodiments of the present method and apparatus provide an improved and novel third alternative: modifying the operation of the in-band hardware so that it approximates an OOB channel. This
15 approach obtains the benefits of hardware-based OOB channels without incurring the expense of an independent I/O subsystem.

A good physical OOB channel has two properties: 1) it is resistant to impersonation and man-in-the-middle attacks by virtue of the physical properties of the channel, and 2) it is widely available and easy to use. Examples of good OOB
20 channels include infrared, audio, visual (e.g., printed barcodes or equivalent), RFID, point-to-point physical wires, and memory tokens such as portable USB flash drives. A common characteristic of these OOB channels is that they allow a digest of a key and/or a key itself to be imported into a device based on physical proximity of some sort. The physical proximity property is crucial, because this makes remote
25 impersonation and man-in-the-middle attacks much more difficult. Physical proximity helps a user to know that their device is actually communicating with the intended peer device as opposed to that of an attacker. Introduction methods based on physical OOB channels with proximity are typically quite good at reflecting the intent of the user. Logical OOB channels are more easily subverted, since they often
30 rely on a technically sophisticated and security-conscious user.

Embodiments of the present method and apparatus provide a radio-based in-band channel that approximates the properties of an OOB channel. These embodiments attenuate the signal strength of the radio significantly so that the range

of the radio is diminished to a few inches. Thus the in-band radio function has similar properties to an infrared channel or RFID transponder. The problem with a naïve implementation of this approach, however, is that attackers with more powerful antennas could fool a device into thinking that it is communicating with another
5 device in close physical proximity.

In some embodiments of the present method and apparatus resistance is added to remote amplified antennas by taking advantage of physical properties of the radio signal propagation. Electromagnetic energy emitted by a non-directional antenna diminishes in intensity as the square of the distance from the antenna. This means
10 that if the signal strength is X at 5 inches from the source, it will be $X/9$ at 15 inches. Thus, weak radio signals rapidly diminish in intensity over a short distance. If the devices being introduced can not only reduce their own in-band signal strengths but can also measure the signal strength gradient of signals received from their peers in real-time, then physical proximity can be reliably detected.

Some embodiments of the present method and apparatus rely on the steep
15 signal gradient of very short-range radios. Embodiments of the present method and apparatus may embed a signal gradient detector or a secondary antenna that is physically separated on the device far enough to detect the slope of the signal strength gradient during introduction, when the radio signal strengths are greatly
20 reduced. The physical separation and the required sensitivity of the antennas will depend upon the specific device geometries. If the effective range of the introduction signal is five inches or less (which is desirable to increase assurance of close proximity), then the needed antenna separation would probably be on the order of an inch or two. If the precision of the gradient sensitivity is increased, then even
25 less separation is needed. For example, a signal gradient sensor may be integrated into the wireless chip itself.

Buttons may be pressed on each device at the start of an introduction. The devices each may respond by reducing their signal strengths to an effective range of five inches or less. The user may then move the devices close together. They detect
30 each other's signal and at the same time measure the signal gradient by simultaneously measuring the signals received on both antennas. Any signals that do not diminish according to an acceptable profile are rejected as potential attacks. If a signal does have the correct strength gradient, cryptographic material (such as public

keys or key identifiers, but not precluding symmetric keys) is received across that channel and subsequently used to establish secure connections on the in-band channel after introduction. At least part of the signal that delivers the keying material is checked for the appropriate signal strength gradient. This effectively
5 binds the proximity information with the keying material.

Even if only one of the radios is capable of signal strength measurement, embodiments of the present method and apparatus can still help make the channel more resistant to attack since the user interface on that device will indicate that it did not successfully establish a connection. The user attempting an introduction can at
10 that point abort the introduction operation on the other device and try again. If neither device can measure signal strength, then the introduction method will still work, but it will have some susceptibility to remote attackers. Therefore, embodiments of the present method and apparatus scale well in that they do not require all devices to implement the advanced capabilities. Some security benefit
15 may be obtained with simple signal attenuation; greater benefit may be obtained if one device has a gradient detector; maximum benefit may be obtained if both devices have gradient detectors.

Although remote attackers can manipulate their signal strength in real time, they cannot manipulate the instantaneous signal strength gradient of the signal they
20 emit. The latter property is a fundamental characteristic of radio signal propagation.

Forming a constellation of devices requires the devices be introduced to each other securely and with minimal user intervention. Using an out-of-band (OOB) mechanism is one way of bootstrapping trust. The OOB channel may be used to pre-configure the devices with authentication information and other information, which helps them to
25 connect using the regular (in-band) channels.

Media used for OOB communication can be broadly classified into one-way (e.g. USB flash storage, floppy disks etc) and two-way (IR).

When two-way OOB media are used for the OOB exchange, both participating devices can be pre-configured with the information that will be used during the in-
30 band authentication phase. This simplifies the later in-band message exchanges.

Using one-way media, one device "pushes" its authentication information, while the other device only reads this information. This poses some challenges

during the in-band authentication because now only one device has the authentication information pre-configured before connection.

ONE-WAY OOB CHANNEL

5 Embodiments of the present method and apparatus provide bootstrap authentication information using a one-way OOB channel and subsequent in-band messages.

 When a one-way OOB channel is used to convey authentication information, there are two possible scenarios: One where the sender is the server and the other
10 where the sender is the client. Here, client refers to a device that wishes to be enrolled in the device constellation, and server is the device that has the authority to authenticate and subsequently authorize the client to become a part of the constellation.

 FIG. 7 depicts a scenario where the server initiates the connection.

15 The embodiment in FIG. 7 has the following elements:

 CERTs - Server's certificate

 S1, S2, S3 - Randomly generated secrets

 Info - Server's connection info

 E(S1)Ps, E(S2)Ps, E(S3)Ps - Secrets encrypted using the Server's Public key

20 CERTc - Client's certificate

 E(S2)Pc, E(S3)Pc - Secrets encrypted using the client's public key

 Regarding the OOB message in FIG. 7:

 The server sends its certificate, a randomly generated secret (S1) and some
25 connection information to the client over the OOB channel.

 Regarding the in-band messages in FIG. 7:

 The client connects to the server over the in-band channel and sends the secret (S1) encrypted with the server's public key, another randomly generated secret (S2) also encrypted with the Server's Public Key and the Client's Public Key.

30 The server decrypts the secret (S1) and verifies that this is the same client to which it had sent the OOB data. If the server decides to allow the client to join the constellation, it creates a certificate for the client with the Client's Public Key. Using the Client's Public

Key, it encrypts the secret (S2), and another randomly generated secret (S3). The server then sends these items to the client.

The client decrypts the secret (S2) with its public key and verifies that its previous message (#2) had reached the server. The client certificate can now be used for future transactions. The client, after decrypting S3, encrypts it again using the Server's Public Key, and sends it to the server as a guarantee that the server's last message (#3) reached it.

After this, normal communication between the client and the members of the constellation can start.

FIG. 8 depicts a scenario where the client initiates the connection.

The embodiment in FIG. 8 has the following elements:

Pc - Public key of the client

S1, S2, S3, S4 - Randomly generated secrets

Info - Client's contact information

CERTs - Server's certificate

E(S1)Pc, E(S2)Pc, E(S3)Pc, E(S4)Pc - Secrets encrypted using the Client's

Public key

E(S2)Ps, E(S3)Ps, E(S4)Ps - Secrets encrypted using the server's Public key

CERTc - Client's certificate

Regarding the OOB message in FIG. 8:

The client sends its public key and a randomly generated secret (S1) and some connection information to the server over the OOB channel.

Regarding the In-Band messages in FIG. 8:

If the server decides to deny entry of the client into the constellation, it need not do anything further. If the server decides to allow the client into the constellation, it connects to the client over the in-band channel and sends the secret (S1) encrypted with the Client's Public Key, the server's certificate and another randomly generated secret (S2) also encrypted with the Client's Public Key.

After verifying the secret (S1), the client is assured that this is the same server to which it sent the OOB message. The client encrypts the secret (S2), and a newly generated secret (S3) with the Server's Public Key and sends it to the server.

The secret (S2) assures the server that the client received its previous message (#2). The server creates a new certificate for the client. It encrypts the secret (S3),

and a newly generated secret (S4) with the Client's Public Key. It then sends the certificate and the two secrets to the client.

Receiving S3 assures the client that the server received its previous message (#3). It then sends a message to the server with the secret (S4) encrypted using the Server's
5 Public Key, as an assurance that it received message #4 from the server.

Normal communication between the client and the members of the constellation can begin after this point.

TWO-WAY OOB CHANNEL

10 Embodiments of the present method and apparatus bootstrap trust between two communicating devices using a two-way OOB channel with very little or no user intervention. Examples of two-way OOB channels are Infrared ports on laptop PCs or a cross-over Ethernet cable connecting the two devices.

FIG. 9 depicts an embodiment with two devices that want to conduct secure
15 communications using the network that is also shared by many other devices/users. This network, which could be wireless LAN or wired LAN, is referred as the in-band channel. In order to bootstrap trust, they use another connection, which, by nature of its physical properties is deemed secure. This connection may be a two-way channel. For example, this could be two-way Infra Red or a small cross-over Ethernet cable and is referred to as
20 the OOB channel.

FIG. 9 shows laptop personal computers as devices, Device1 and Device2, but they could be any communicating device. The two-way OOB channel used is secure from any kind of attack because of its physical properties. Basically, a human must be able to infer from just looking at the physical connection that a third device
25 cannot intercept the communication and that the two devices are the intended peers to each other. If a cross-over Ethernet cable is used, it can be seen by human eyes that it is connected to the two devices. If an infrared channel is used, by its directionality and short range, humans can be sure that a third device cannot interpose itself with a man-in-the-middle attack. Since public key techniques are used to protect the OOB
30 exchange, passive eavesdroppers on the OOB channel cannot learn any secrets that would enable them to violate the security of the system.

First, the two devices are connected via the two-way OOB channel for the bootstrapping. They exchange some information between them on the OOB channel.

Then, the two devices connect to each other on the in-band channel. Using the information exchanged on the OOB channel, the two devices are able to prove to each other that they are the intended peers and also able to exchange encrypted messages between them for secure communication.

5 FIG. 10 shows an embodiment of the detailed message exchange between the two devices. The initiator device is shown as client and the responder device is shown as server in FIG. 10. This embodiment operates as follows.

 The client sends its public key to the server over the OOB channel.

 The server sends its digital certificate to the client over the OOB channel. In
10 addition to sending its digital certificate, the server sends a randomly generated secret (S1), and information required for the client to connect to the server. The connection information could be IP address and TCP port number in most cases. In the case of WLAN in-band channel, SSID (Service Set Identifier) and WEP (Wired Equivalent Privacy) key may also be included.

15 The client can now connect to the server over the in-band channel and present the secret (S1) encrypted by the server's public key. In addition, the client sends a randomly generated secret (S2), also encrypted by the server's public key.

 If the server can decrypt the secret (S1) with its private key and verify that it is same as the one it sent previously to the client, the identity of the client is established. The
20 server decrypts secret (S2) using its private key and uses it in its next message to the client.

 The server can now create a digital certificate for the client and send it to the client on the in-band channel along with secret (S2) encrypted with the Client's Public Key. The server signs the digital certificate. In addition, the server sends another randomly
25 generated secret (S3), also encrypted by the Client's Public Key.

 The client decrypts the secret (S2) with its private key and if it is same as the one it sent previously to server, then the client is assured of the identity of the server. The digital certificate for the client can be used for future communications with the server. The client decrypts secret (S3) using its private key and uses it in its next message to the server.

30 The client encrypts the secret (S3) with the Server's Public Key and sends it to the server. The server decrypts the secret (S3) using its private key. If it is the same as the one it sent previously to the client, then the server is assured of receipt of the client certificate by the intended client.

The server's certificate may be a self-signed certificate or a certificate signed by a certificate authority. The client's certificate may be signed by the server.

Embodiments of the present system and method significantly enhance the value of wireless devices by making it easy to establish trusted groups of devices with
5 secure connections and mutual access control.

Implementing the apparatus, systems, and methods disclosed herein may enhance the security of networks. Embodiments of the present system and method provide improved operation regarding: a framework for out-of-band introduction followed up by in-band credential and application policy establishment phase (configuration of application
10 policies will not be part of our initial standardization effort); the typical user experience for introducing new wireless devices to a network; support for multiple different out-of-band and in-band communication channels; simple role-based access control; and binding domain credentials and policies to specific networks.

The accompanying drawings that form a part hereof, show by way of illustration,
15 and not of limitation, specific embodiments in which the subject matter may be practiced. The embodiments illustrated are described in sufficient detail to enable those skilled in the art to practice the teachings disclosed herein. Other embodiments may be utilized and derived therefrom, such that structural and logical substitutions and changes may be made without departing from the scope of this disclosure. This Detailed Description, therefore,
20 is not to be taken in a limiting sense, and the scope of various embodiments is defined only by the appended claims, along with the full range of equivalents to which such claims are entitled.

Such embodiments of the inventive subject matter may be referred to herein, individually and/or collectively, by the term "invention" merely for convenience and
25 without intending to voluntarily limit the scope of this application to any single invention or inventive concept if more than one is in fact disclosed. Thus, although specific embodiments have been illustrated and described herein, it should be appreciated that any arrangement calculated to achieve the same purpose may be substituted for the specific embodiments shown. This disclosure is intended to cover any and all adaptations or
30 variations of various embodiments. Combinations of the above embodiments, and other embodiments not specifically described herein, will be apparent to those of skill in the art upon reviewing the above description.

The Abstract of the Disclosure is provided to comply with 37 C.F.R. §1.72(b), requiring an abstract that will allow the reader to quickly ascertain the nature of the technical disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing

5 Detailed Description, it can be seen that various features are grouped together in a single embodiment for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter lies in less than all features of a single disclosed embodiment.

10 Thus, the following claims are hereby incorporated into the Detailed Description, and each claim stands on its own as a separate embodiment.

What is claimed is:

1. A method for bootstrapping applications and services at different layers in a communications stack, comprising:

registering applications and network services for notification of an out-of-band introduction; and

5 using the out-of-band introduction to bootstrap secure in-band provisioning of credentials and policies to be used to control subsequent access and resource sharing on an in-band channel.

2. The method according to claim 1, wherein the credentials are domain credentials.

10 3. The method according to claim 2, wherein domain credentials issued to newly-added device include implicit role-based authorization so that no additional access control configuration is needed after the introduction is complete.

4. An apparatus that provides a device introduction framework that is extensible to bootstrap multiple applications and services at different layers in a communications stack of a platform, comprising:

15 an out-of-band (OOB) introduction mechanism and data structure for conveying keying material, names for user display, and in-band addressing information for a next phase of a credential and policy bootstrapping process for newly-introduced devices;

20 an extensible data structure for conveying across an in-band channel additional credentials, domain policy, and configuration data associated with a domain to which one of the newly-introduced devices belongs;

a protocol that confirms knowledge of the out-of-band keying material and derives a secure in-band channel for exchanging data structures used to provision credentials and domain policy; and

25 a processing framework structured to maintain credentials and associated domain policy and to make credentials and associated domain policy available to a trusted application and driver components on the device platform.

5. The apparatus according to claim 4, wherein the out-of-band introduction mechanism is one of a wire, a USB token, RFID (radio frequency identifier), infrared, and short-range radio.

6. The apparatus according to claim 4, wherein multiple applications and services at different layers in a communications stack run on a platform, and wherein applications running on the platform during subsequent operations consult the processing framework for credentials and policy associated with a currently active domain.

7. The apparatus according to claim 4, wherein the apparatus further comprises a centralized policy store using an extensible data structure, and wherein service providers that own or manage devices easily provision policy for a variety of different applications and services by updating the policy data for a domain of the service providers.

8. The apparatus according to claim 4, wherein globally-unique domain identifiers are used to prevent policy updates from one domain owner from overwriting policy associated with a different domain.

9. The apparatus according to claim 4, wherein, at runtime, the processing framework provides the appropriate policy to applications and services according to a currently active domain.

10. The apparatus according to claim 4, wherein the hardware-based structure of the OOB introduction resists attackers present on the in-band channel, and never requires entering or confirming passwords or other cryptographic secrets.

11. The apparatus according to claim 4, wherein integration of the OOB introduction is structured such that role-based access control and policy framework increases user convenience and remote device manageability for service providers and application developers.

12. The apparatus according to claim 4, wherein the OOB introduction is structured to effect bootstrap secure in-band provisioning of credentials and policies that are used to control subsequent access and resource sharing on the in-band channel.

5 13. A method for an introduction process that provides a device introduction framework that is extensible to bootstrap multiple applications and services at different layers in a communications stack, comprising the steps of:

registering applications and network services for notification of out-of-band (OOB) introductions;

10 OOB introducing at least one of shared secret, in-band addressing info, and public key;

receiving, by the applications, callback with OOB channel characteristics, domain identifier, and peer device name, and providing an opportunity to pass policy data to peer over an in-band channel;

15 conducting at least one of in-band credential establishment and in-band policy exchange;

establishing subsequent secure interaction based on identities and keys; and

consulting, by multiple software components, with introduction framework to determine domain and use appropriate credentials and policies.

20 14. The method according to claim 13, wherein the domain credentials issued to a newly added device includes implicit role-based authorization so that no additional access control configuration is needed after the introduction process is complete.

25 15. The method according to claim 13, wherein the OOB introduction is used to bootstrap secure in-band provisioning of credentials and policies that are used to control subsequent access and resource sharing on the in-band channel.

16. The method according to claim 13, wherein the method further comprises, in conducting at least one of in-band credential establishment and in-band policy

exchange, proving knowledge of keys previously exchanged OOB and generating new keys.

17. The method according to claim 13, wherein multiple applications and services at different layers in a communications stack run on a platform, and wherein
5 applications running on the platform during subsequent operations consult the processing framework for credentials and policy associated with a currently active domain.

18. The method according to claim 13, wherein the method further comprises maintaining a centralized policy store using an extensible data structure, and wherein
10 service providers owning/managing devices easily provision policy for a variety of different applications and services by updating the policy data for domains of the service providers.

19. The method according to claim 13, wherein the method further comprises using globally-unique domain identifiers to prevent policy updates from one domain
15 owner from overwriting policy associated with a different domain.

20. The method according to claim 13, wherein at runtime, the processing framework providing appropriate policy to applications and services according to a currently active domain.

21. The method according to claim 13, wherein integration of the OOB
20 introduction with role-based access control and policy framework further increases user convenience and remote device manageability for service providers and application developers.

22. The method according to claim 13, wherein the OOB introduction is used to bootstrap secure in-band provisioning of credentials and policies that are used to
25 control subsequent access and resource sharing on the in-band channel.

23. A method for introducing two devices, comprising the steps of:

activating each device of two devices to initiate an introduction;
reducing a respective strength of each device to an effective range of about
five inches or less;
moving the devices toward each other to form an out-of-band channel;
5 each device detecting a signal of the other device, and simultaneously
measuring a signal gradient by simultaneously measuring the signals received on
both antennas;
determining if any of the signals diminish according to an acceptable profile;
and
10 rejecting the introduction if any of the signals are not so diminished, and if a
signal does have a correct strength gradient, receiving cryptographic material
(such as public keys or key identifiers, but not precluding symmetric keys) across
the out-of-band channel and subsequently using the cryptographic material to
establish secure connection on an in-band channel after introduction.

15 24. The method according to claim 23, wherein operation of in-band hardware is
modified such that it approximates an out-of-band channel.

25. The method according to claim 24, wherein the devices being introduced
reduce their own in-band signal strengths and also measure a signal strength
gradient of signals received from the other device in real-time, to detect
20 physical proximity.

26. A method for establishing trust between a server and a client, comprising the steps
of:

initiating an introductory process for the server and the client using an out-of-band
(OOB) message in an OOB channel; and

25 using in-band messages in an in-band channel to complete the introduction process
for the server and the client.

27. The method according to claim 26, wherein the server sends its certificate,
randomly generated secret and connection information to the client over the OOB
channel.

28. The method according to claim 27, wherein:

the client connects to the server over the in-band channel and sends a first secret encrypted with a server's public key, another randomly generated second secret also encrypted with the server's public key and a client's public key;

5 the server decrypts the first secret and verifies that this is the same client to which it had sent the OOB data;

if the server decides to allow the client to join a constellation, the server creates a certificate for the client with the client's public key, and using the client's public key, the server encrypts the second secret, and another randomly generated third secret, the server then sending the certificate, the second secret and the third secret to
10 the client;

the client decrypts the second secret with its public key and verifies that its previous message reached the server; and

the client, after decrypting the third secret, encrypts the third secret again
15 using the server's public key, and sends the third secret to the server as a guarantee that the server's last message reached the client.

29. The method according to claim 26, wherein the client sends its certificate, randomly generated secret and connection information to the server over the OOB channel.

20 30. The method according to claim 29, wherein:

if the server decides to deny entry of the client into the constellation, it need not do anything further, and if the server decides to allow the client into a constellation, the server connects to the client over an in-band channel and sends a first secret encrypted with a client's public key, a
25 server's certificate and a randomly generated second secret also encrypted with the client's public key;

after verifying the first secret, the client is assured that this is the same server to which it sent the OOB message, the client encrypts the second secret, and a newly generated third secret with the server's public key and
30 sends the third secret to the server;

the second secret assures the server that the client received its previous message, the server creates a new certificate for the client, the server encrypts the third secret and a newly generated fourth secret with the client's public key, and the server sends the certificate and the third and fourth secrets to the client; and

5

receiving the third secret assures the client that the server received its previous message, the client then sends a further message to the server with the fourth secret encrypted using the server's public key, as an assurance that the client received the previous message from the server.

10 31. A method for bootstrapping trust between two communicating devices, comprising the steps of:

connecting the two devices via a two-way out-of-band (OOB) channel for the bootstrapping;

exchanging information between the devices on the OOB channel;

15 connecting the two devices to each other on an in-band channel;

proving to the two devices, using the information exchanged on the OOB channel, that the two devices are intended peers; and

exchanging encrypted messages between the two devices for secure communication.

20 32. The method according to claim 31, wherein one device is a server and the other device is a client, and wherein the client sends its public key to the server over the OOB channel.

33. The method according to claim 32, wherein:

25 the server sends to the client over the OOB channel the server's digital certificate, a randomly generated first secret, and information required for the client to connect to the server;

the client connects to the server over the in-band channel and sending the first secret encrypted by the server's public key, a randomly generated second secret that is encrypted by the server's public key;

an identity of the client is established, if the server can decrypt the first secret with its private key and verify that it is same as the one it sent previously to the client, the server decrypting the second secret using its private key and using the second secret in its next message to the client;

5 the server creates a digital certificate for the client and sends the digital certificate to the client on the in-band channel along with the second secret encrypted with the client's public key, the server sending another randomly generated third secret also encrypted by the client's public key;

10 the client decrypts the second secret with its private key and if it is same as the one the client sent previously to server, then the client is assured of an identity of the server, the client decrypting the third secret using its private key and using the third secret in its next message to the server;

 the client encrypts the third secret with the server's public key and sends it to the server; and

15 the server decrypts the third secret using its private key, and if it is the same as the one the server sent previously to the client, then the server is assured of receipt of the client certificate by the intended client.

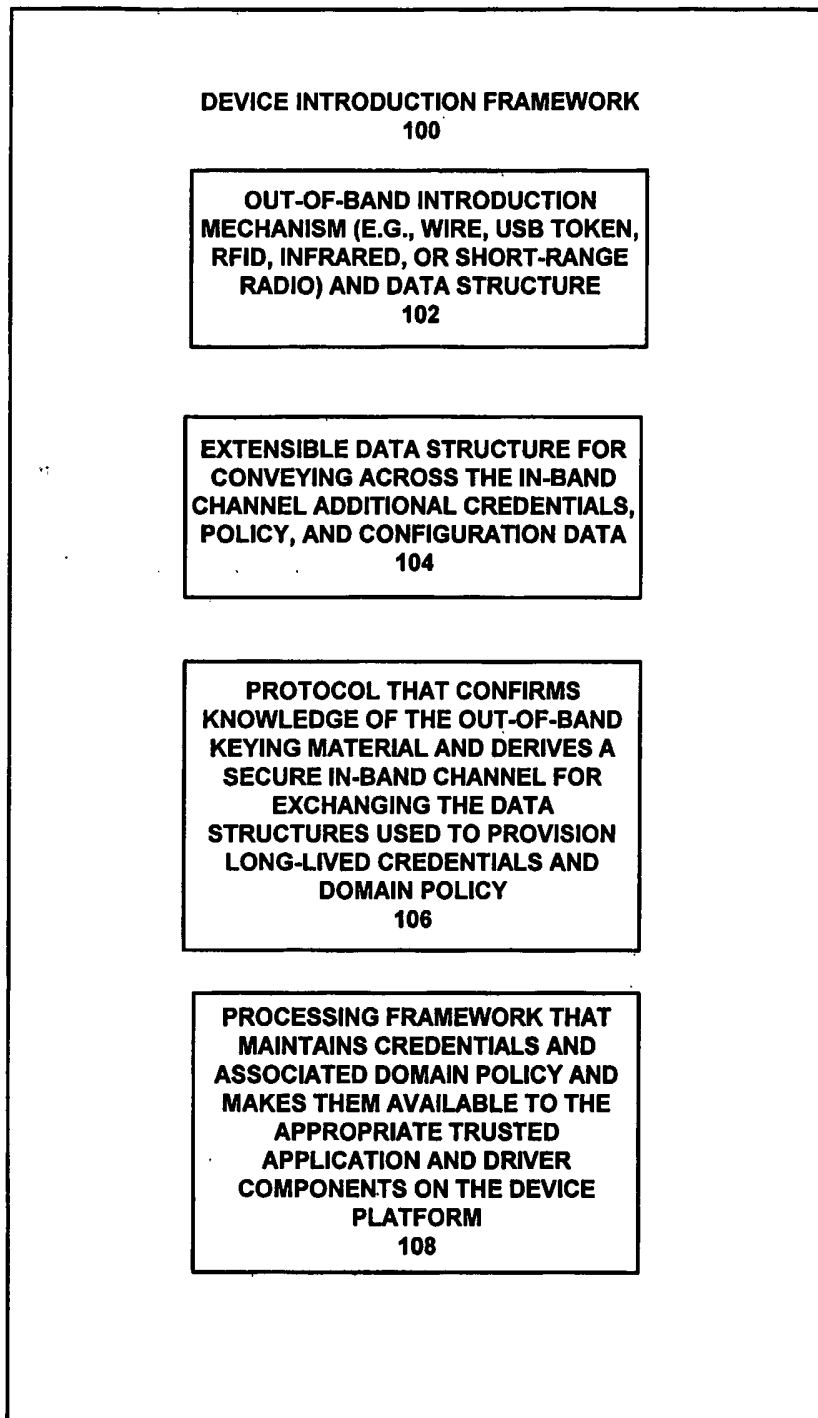


FIG. 1

2/10

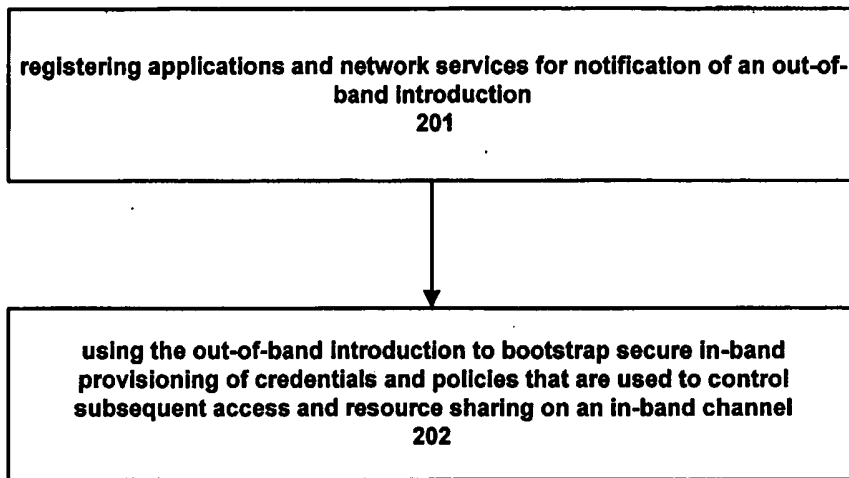


FIG. 2

3/10

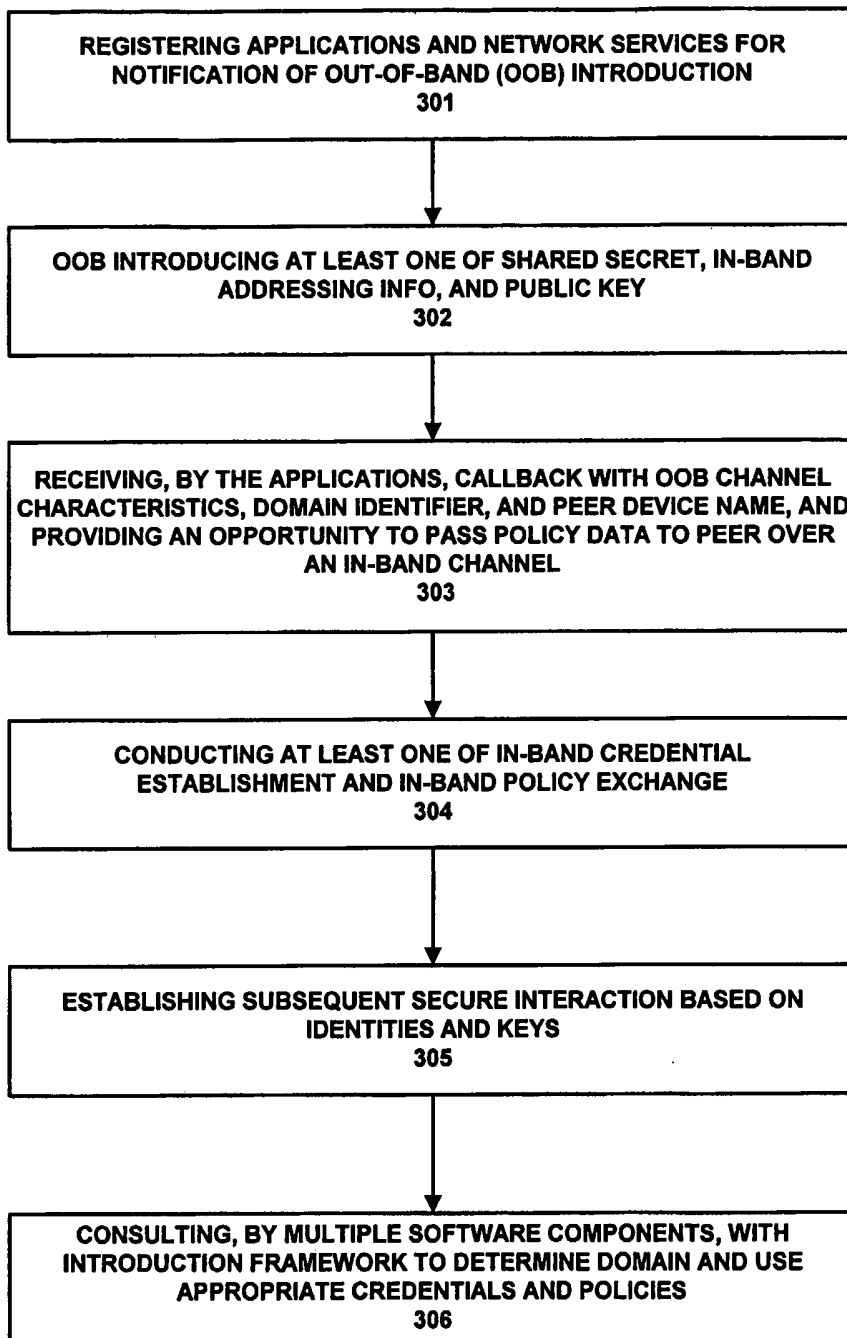


FIG. 3

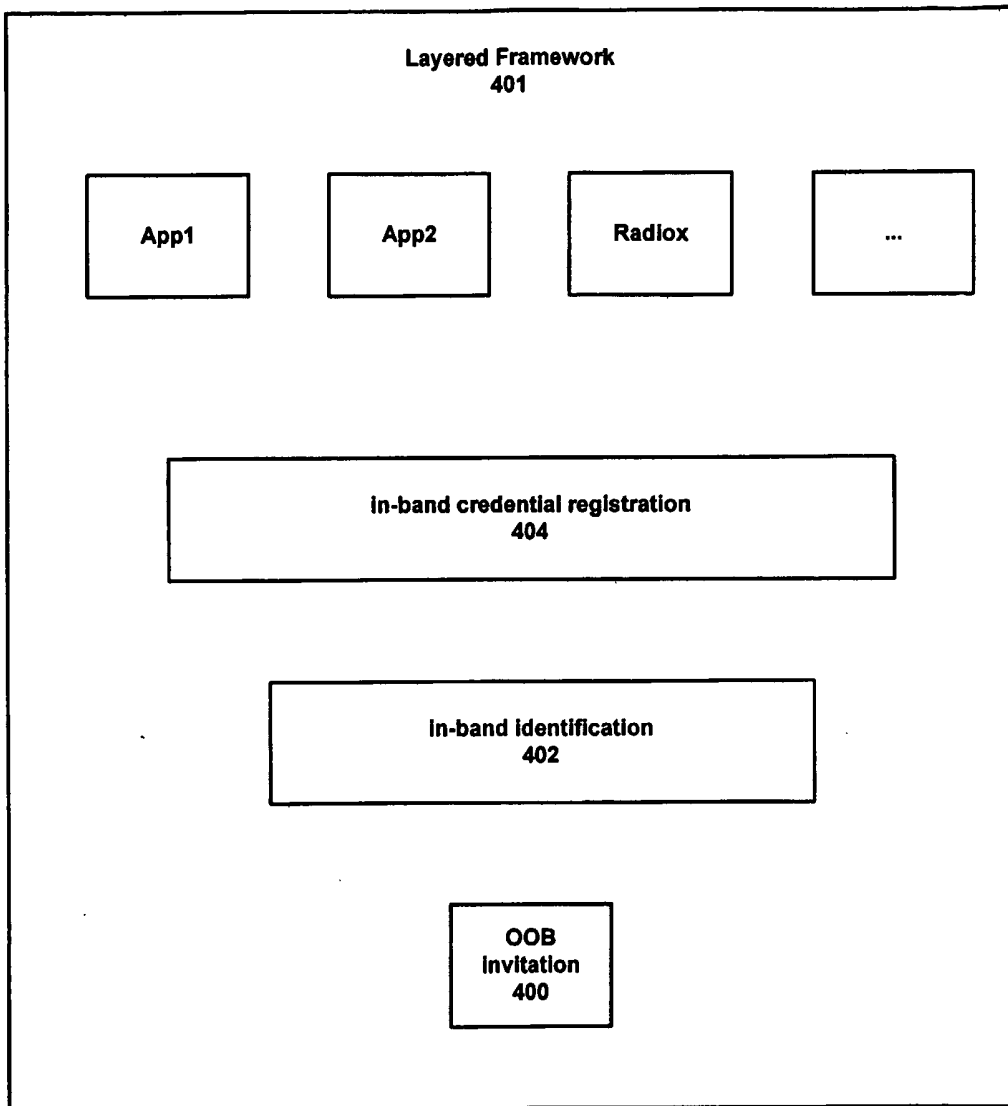


FIG. 4

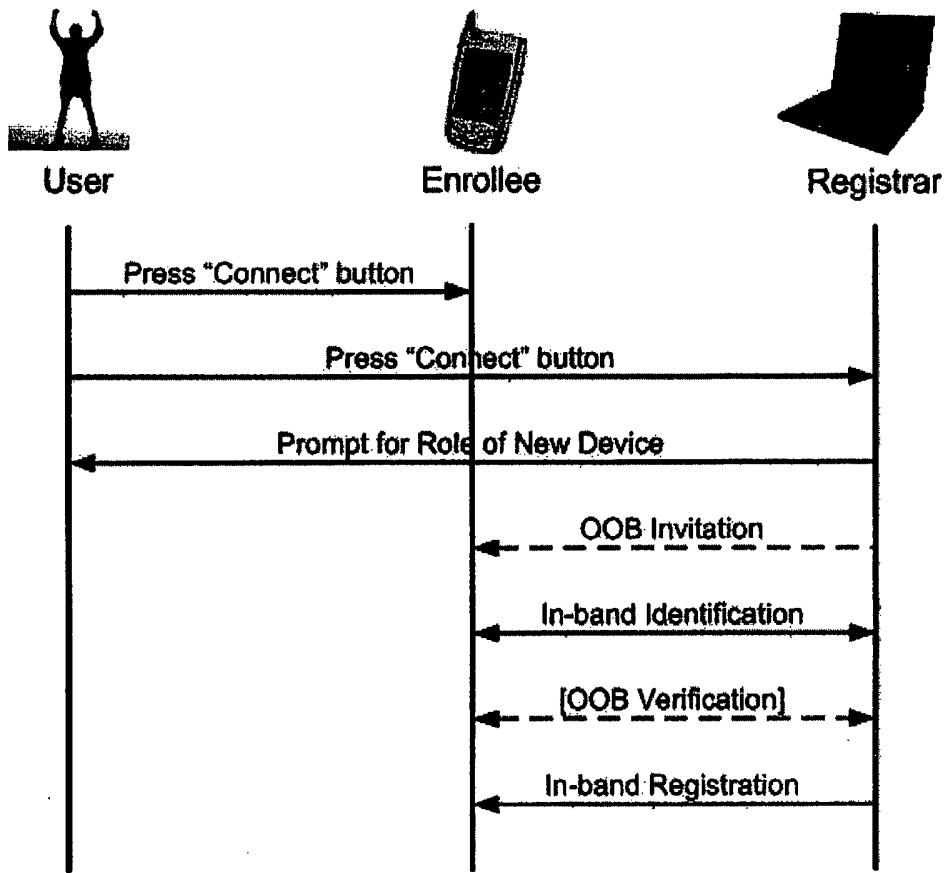



FIG. 5

Secure Introduction Example

- Out-of-Band Invitation



- In-Band Registration (over 802.11)

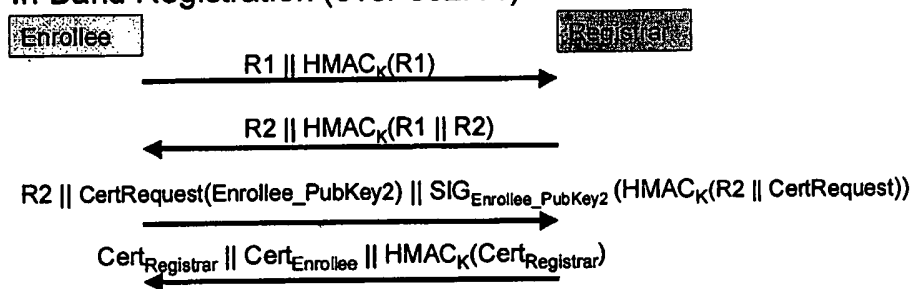


FIG. 6




FIG. 7




FIG. 8




FIG. 9

FIG. 10

