

Enforcing Security in Semantics Driven Policy Based Networks

Palanivel Kodeswaran, Sethuram Balaji Kodeswaran, Anupam Joshi, Tim Finin

*Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
1000 Hilltop Circle, Baltimore, MD 21250
{palanik1,kodeswar,joshi,finin}@cs.umbc.edu*

Abstract—Security is emerging as an important requirement for a number of distributed applications such as online banking, social networking etc. due to the private nature of the data being involved. Further more, the wide spread use of portable devices such as laptops, PDAs etc. allows users to make meaningful ad hoc collaborations. Traditional security solutions are not feasible for these scenarios due to the varying nature of the collaborations in terms of entities involved and their roles, available resources etc. Under these circumstances, we need generic solutions that take into account the semantics of the collaborations in determining the set of allowable operations. In this paper, we propose an extensible framework that uses semantics driven policies for enforcing security. Our policies are rooted in semantic web languages which makes amenable to interoperability, and also enables high level reasoning for conflict resolution and policy adaptation. We describe our policy based network that uses packet content semantics to best handle different streams, and show how our framework can be used to secure enterprise networks and the BGP routing process.

I. INTRODUCTION

Security is emerging as an important requirement for a number of distributed applications such as online banking, social networking etc. due to the private nature of the data being involved. Further more, the wide spread use of portable devices such as laptops, PDAs etc. allows users to make meaningful ad hoc collaborations. Traditional security solutions are not feasible for these scenarios due to the varying nature of the collaborations in terms of entities involved and their roles, available resources etc. Under these circumstances, we need generic solutions that take into account the semantics of the collaborations in determining the set of allowable operations. In this paper, we propose an extensible framework that uses semantics driven policies for enforcing security. We describe our policy based network that uses packet content semantics to best handle different streams, and show how our framework can be used to secure enterprise networks and the BGP routing process.

In our system, we use policies for enforcing security as policies provide a generic and flexible framework which can later be easily modified based on changing requirements. Given the dynamicity of emerging computing environments, we want to be able to specify our policies at a high level such that we can focus on the abstract conditions and constraints that need to be maintained in the system. Also, given the heterogeneity of available devices, we expect that policy

specifications should be as device independent as possible. In these cases, to enforce policies, an adaptation layer would be used to translate high level policy specifications into low level device specific primitives. Allowing automated reconfiguration of devices on the fly would require that the system be able to reason about policies and adapt them based on the new requirements.

We propose that policies specified in semantic web languages can satisfy the above requirements. In our system, policies are specified using a combination of OWL and SWRL. In particular, we choose OWL DL as it is complete and decidable, and therefore all conclusions are guaranteed to be computed in a finite amount of time. The combination of OWL and SWRL can be used to define ontologies using which one can declaratively define facts, policies and rules in terms of what needs to be true or false for a policy to hold. In our system, policy specifications are in terms of SWRL rules which use high level concepts defined in appropriate ontologies, thus making the policy specifications generic, device independent and extensible. We can also specify meta policies for guiding the interaction among policies. For example, we can use meta policies to prioritize policies when multiple policies are applicable in a context. Further, we envisage that different organizations would have different policies at different granularities for the same device. By specifying policies in semantic web languages, devices would be able to reason over the policies and arrive at a configuration that meets the overall combined requirements. Also, rooting policies in semantic web languages makes dynamic reconfiguration automatic and easy, as new facts can be inferred from the policies.

The rest of this paper is organized as follows. Section 2 describes our content based tagging scheme. In Section 3, we present our semantics driven policy based network. Section 4 describes the rationale behind using semantic web languages for policy specification. Section 5 describes how security policies can be enforced in our framework. In section 6, we present related work and finally we conclude in section 7.

II. CONTENT BASED SEMANTIC TAGGING

This section presents our packet level semantic tagging framework that enables intermediary routers to reason over the tags to determine how to best handle the data streams

flowing through them. Our approach is to provide the network routers visibility into the semantic content of data streams passing through them. This content level information can then be used by the routers to make more intelligent routing and data handling decisions. Our approach differs from active networks in that the data streams merely provide additional meta data while the network has complete control on how to use this metadata. Thus network operators still retain complete over their network operations. In our framework, we use RDF for labelling the semantic content. We choose RDF as it is very flexible, generic and its growing acceptance as the de-facto standard for meta-data markup. By utilizing RDF as the mechanism to markup flows/packets, intermediary intelligent routing entities can use this metadata to reason over their knowledge base to determine how best to handle a given flow. Also, inferences can be made to generalize or specialize a given flow to best meet its demands.

Now we present our system architecture. We break it down into two components; at a node level and at a system level that spans the network.

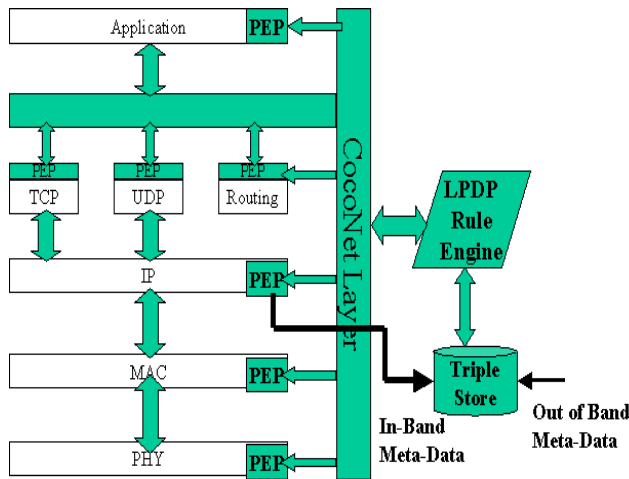


Fig. 1. CoCoNet Node Framework

At the node level, the architecture we propose introduces our Node Framework as an additional layer called the CoCoNet layer between the application and the transport layer. This layer is responsible for intercepting socket calls made by applications to the transport layer. The API is enhanced to allow the application to provide semantic level information for messages transmitted over this interface. A Local Policy Decision Point (LPDP) is used to determine what policies to enforce based on the content. In our framework, each Policy Enforcement Point (PEP) is at every layer in the networking stack while [1] treats the PEP at a node level. Placement of the PEP at every level of the stack allows us to implement coordinated cross layer interactions initiated and controlled by our framework. The PEP exposes the interlayer optimization points that any particular layer supports. The framework utilizes the policies stored in the LPDP to drive the settings to be applied to each of the PEPs in the stack.

Essentially, we are proposing to expose a network stack as a collection of switches and dials and allow an external policy to determine the exact settings of each of these dials (based on content and context). We want to expose functionality, not necessarily the mechanism of how it is achieved (this falls under intra-layer optimization). For example, a MAC can advertise two different data rates and their associated packet error probabilities without exposing the FEC scheme used to achieve these rates. The policies can be specified as production rules (if (condition) then (action)) or event-condition-action rules (on event if (condition) then (action)). In essence, the Node Framework provides a rich, extensible option for realizing policy controlled cross-layer interactions within a node's network stack. By parameterizing the possible set of interactions that are permissible, the cross layer interactions are kept tractable without making the implementation overly convoluted [2].

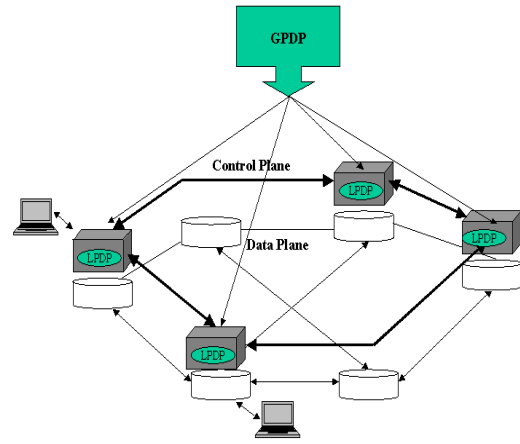


Fig. 2. Overlay Network

At the network level, we envision that there will be an overlay network comprised of routers that run the CoCoNet Router Framework. Client machines running our Node Framework communicate over this overlay. The overlay comprises of two components;

- A control plane component that involves interactions between the CoCoNet Router Layers at the routing elements.
- A data plane component through which the data packets are flowing.

Over the CoCoNet Router control plane, routers can exchange traditional management information such as link states, buffer lengths etc. In addition, information such as content types currently being handled, adaptations currently available can be advertised. An additional key piece of information exchanged is the local policies that are currently being applied to a data stream that is being routed. Local PEP settings for a given stream or flow have global implications. For example, unless every hop is reliable, a data packet cannot be reliably routed through a network. The data plane can be implemented

as either:

- A UDP connection between two routers.
- A TCP connection between two routers.
- An IP-in-IP tunnel between two routers.
- A layer 2 LSP.
- A DiffServ aware network.
- An IntServ aware network.

A CoCoNet Router Framework will perform the necessary mapping based on policy, content and context. For instance, suppose a packet arrives at a router indicating that it requires reliable transfer semantics. The data plane chosen to the next hop, in this case, could be over a TCP connection. Likewise, a data packet indicating that it is sensitive information (telnet logins for example) but currently not encrypted can be routed to the next hop over an IPSEC tunnel or dropped if none is available (if that is the policy). Where a CoCoNet Router Framework runs is very implementation dependent. For example, in case of a wireless adhoc network, every host is a router and hence can potentially run a (albeit simplified) CoCoNet Router Framework. Likewise, in an enterprise setting, the host machines within the enterprise will likely run only the Coconet Node Framework with only the exterior gateway routers running the Coconet Router Framework. A network service provider will most likely have only edge routers run the Coconet Router Framework leaving the core optimized for fast data flow handling. The role of the Global Policy Distribution Point (GPDP) is to disseminate any network wide policies that need to be enforced. This can include items such as preferential treatment that needs to be given to content originating from a particular domain, preferential treatment for a particular type of content, any content based adaptation techniques that need to be employed in the network etc. It is envisioned that the GPDP is controlled by the ISP to set forth global rules while the LPDP hosted at an enterprise location is possibly shared between the ISP and the enterprise. This can further be extended to say that the LPDP is under local user control (based on user policies and preferences) and can additionally, host user preferences.

The information conveyed in the metadata is really left up to the application. For example, an MP3 stream may have the following description which can be used to differentiate between official and entertainment video streams.

```
<?xml version="1.0"?>
<rdf:RDF
  xmlns:rdf=
    "http://www.w3.org/1999/
      02/22-rdf-syntax-ns#"
  xmlns:mmschema=
    "http://www.mySchema.org/mms#">
<rdf:Description
  rdf:about=
    "http://www.myContent.com/
      SalesReport.mp3">
<rdf:type rdf:resource=
  "http://www.mySchema.org/mms#audio"/>
```

```
<mms:LengthInMin>5</mms:LengthInMin>
<mms:LengthInMB>4</mms:LengthInMB>
<mms:technicalType>
  http://www.mySchema.org/mms#MP3
</mms:technicalType>
<mms:semanticType>
  http://www.mySchema.org/mms#Lecture
</mms:semanticType>
</rdf:Description>
```

Furthermore, providing content information so that a router can differentiate between, for example, video streaming from a surveillance camera and a streaming movie allows the network to make smart decisions on routing data streams across links with different reliability characteristics. Also, for our architecture, we are using RDF which provides a generic mechanism to convey metadata which can be reasoned over.

III. SEMANTICS DRIVEN POLICY BASED NETWORK

In this section, we present our policy based network built on top of the semantic tagging architecture.

Policy based networks employ mechanisms that allow network operators to specify at a high level, rules defining how packet flows are handled within a network, how network resources are allocated, define access control restrictions and levels of service. The policies are enforced by configuring the network devices with the requisite primitives so that the desired actions are performed on the data streams. One of the main challenges frequently faced is ensuring that network configuration settings are applied consistently throughout the network so that the correct actions are taken by the network devices; however, this is often error-prone and difficult to manage especially when there is a heterogeneity of network devices and management protocols. Additionally, policies that are commonly in use today are limited in their expressibility. Rules such as “allow traffic from A higher priority over B” and “permit user A” are easy to enforce but are limited in their expressibility. For networks to offer highly specialized services, administrators need to be able to specify more complex handling rules such as “allow security surveillance video streams higher priority than webcasts” (within an enterprise) or “downsample any video to user A so as not to exceed 128 kbps” (due to different levels of service or capabilities of the device associated with the user). For such policies, enforcement cannot be performed by packet header inspection alone as all the requisite details may not be directly accessible from the data packets as they are today.

To solve this issue, we define an alternate model to achieving policy based networks that provides fine grained services for network traffic, automates network configuration and eases network management. The model relies on two key components; namely a tagging mechanism as described in the previous section, that allows packets and/or streams to convey higher level semantic information that can be used in conjunction with the lower level information garnered from packet header inspection and a framework for specifying rules

in an easy to use, formal model that can be checked for consistency. The process of converting the rules to the lower level primitives understood by the network devices is also handled by the framework allowing the network administrators to focus just on defining the administrative policies. In our model, applications encode data packets with descriptions conveying content semantics using the W3C Web Ontology Language (OWL)[3] as explained in the previous section. Ideally, the ontology used for this is provided by the network service provider. This description is encoded as a special header that is embedded into the data stream. Our motivation for using OWL (specifically, OWL-DL) is capability of the language to express formal semantics, defining class hierarchies and their relationships, associated properties, cardinality restrictions while still retaining decidability and computational completeness. Using OWL for ontology specification makes the framework generic, flexible and more scalable than using proprietary labeling schemes that raise interoperability issues.

Utilizing the framework, interim routers that handle the data packets, run a reasoning engine that can reason over the OWL description and invoke rules depending on the correct set of actions that need to be enforced. Our framework utilizes the W3C Semantic Web Rule Language (SWRL)[4] as the rule language which provides an easy to use mechanism for specifying event-condition-action rules which is the majority of rules envisioned for a typical network. Using this framework, content providers now provide metadata to the network that can then be used by the network providers to determine how best to handle a given packet or flow that best suites that content. While our framework allows the deployment of specialized handling routines into the network, a key differentiator between our approach and that of active networking[5], [6] with respect to packet handling is that unlike active networks, the metadata is not a contract on how the data should be handled but rather what the data is. The network provider retains complete control of how the packets are handled within the network and can fine tune policies to offer the best service for that type of content.

A policy based network for a typical enterprise can be built using our proposed framework and additional components for storing policies, conflict resolution and adaptation as described in [7].

IV. DESIGN OF POLICIES ROOTED IN SEMANTIC WEB LANGUAGES

There are several reasons motivating us to root our policy specification and enforcement mechanisms in semantic technologies. Specific to the domain of networking, for any successful policy language, it must be universally interoperable considering the number of various organizations (enterprises, ISPs, networking vendors etc) that must interact to power a large scale network. In addition, if the system needs to be capable of automatically processing, reasoning over and responding as appropriate, the language must be machine-interpretable with understandable syntax and semantics for expressing data, rules and constraints on networks, networking

devices, hardware components, software protocols, user applications and end users. In addition, complexity of the policy specification mechanism is key to the ease of its acceptance. In this regards, a declarative policy language that enables each authority to draft abstract policies in a high-level language capturing policies and logic used for guiding activities of networking enterprise is a good candidate. Each authority can define only those objectives and constraints that are relevant to its needs. The policies represent rules and constraints that are necessary for a target network infrastructure to be valid. This information contained in the policies is defined in a manner that is as hardware, software, and protocol independent as possible. Therefore, the authorities do not focus on writing procedures for configuring a specific infrastructure; instead they focus on describing a generic infrastructure and its features without needing to master and understand each of the various device/protocol/system specific mechanisms. The policy software components embedded or in the vicinity of each of the networked devices can convert the specified policy into device specific settings and configurations.

We believe that the combination of the W3C Web Ontology Language (OWL) and W3C Semantic Web Rule Language (SWRL) standards is applicable for policy control as it is machine understandable, sound, complete, extensible through additional ontologies, and supports heterogeneous application domains. OWL has axiomatic and model-theoretic semantics, which allows for verification of knowledge expressed in OWL constructs. In our work, we have chosen to use a subset of OWL, namely OWL DL as it is complete and decidable. This is an important feature in order to guarantee that all conclusions are computable and that all computations finish in finite time. OWL + SWRL can be used to define ontologies, using which one can declaratively define facts, policies, and rules in terms of what needs to be true or false for a policy to hold. SWRL specifies OWL-based abstract syntax and vocabulary for representing Horn-like rules. SWRL defines a rule as an implication from a set of antecedent atoms to a set of consequent atoms. In our work, the policy language uses the antecedent atoms for representing policy constraints. The language uses the consequent atoms for defining directive actions that apply whenever the constraints are satisfied by evaluating information stored in a local knowledge base and by executing relevant attached procedures. The ontologies can be extended for declaratively capturing any concept or predicate without changes to the underlying system capable of processing OWL and SWRL. The language can be further extended by defining functions as procedural attachments and mapping them to predicates in OWL ontologies. This allows for the policy enforcement mechanisms to process functions while enabling the system to depend on low-level, optimized implementations which is particularly important in the domain of networking. Extending on our current framework, in order to support multiple policies, we can also define a vocabulary for creating meta-policies. Meta-policies are used for guiding the interaction among policies. The meta-level vocabulary defines constructs for resolving conflicting,

overlapping policies. For example, the meta-level vocabulary can be used to create a default conflict resolution rule such that a prohibitive policy overrides permissive policy. At the same time, the meta-level vocabulary allows one to define absolute and relative prioritization of policies, thus overriding the default rule. The meta-policies define an automatic conflict resolution diagnosis in order to respond to situations when policies presented to a network impose conflicting conditions on the overall infrastructure or on one specific component. Additionally, the policy software components embedded or in the vicinity of each networked device can use this meta-information to automatically merge policies from multiple authorities and generate a target configuration that meets the combined requirements. The components follow the semantics defined by the policy language. Consequently, their steps in merging policies can be formally verified using a theorem-proving model. In order to combine multiple policies, the language depends on closed-world assumption reasoning. In this case, the system assumes that all rules are to be evaluated only by the knowledge contained within a knowledge base. This allows a reasoning engine to yield a solution in a finite time.

By utilizing semantic technologies to drive our framework, we can now realize dynamic reconfiguration of knowledge as new facts can be inferred through the policies specified. Current relational technologies and those based on static schema are dependent on pre-existent knowledge and do not offer this flexibility. There are a plethora of tools available to drive the ontology specification, verification, reasoning engine, etc. that can be incorporated to build such a system that can be deployed on a large scale.

V. ENFORCING SECURITY POLICIES IN THE PROPOSED FRAMEWORK

A. Securing Enterprise Networks

The use case we consider in this paper is that of a secure enterprise that wants to enforce prioritization on types of content that can flow across the links comprising its network. The enterprise has profiled its network and assessed security credentials to all the links and routers. As an example, a link that is fully within the premises of the enterprise (physically secure) is assessed as a “safe” link, one that is a VPN running over an external service providers network may be assessed as “potentially unsafe” while a wireless RF hop may be assessed as “unsafe”. The enterprise applications are “smart” and can encode content level tags into the data packets that carry semantic information about the content as well as the application/user/device. For this example, as we are interested in the security semantics, applications additionally provide information about the sensitivity of the content (such as secret, top secret or normal), type of content and the security credentials of the context within which they run. For such an enterprise, the following policies would be appropriate:

- Only “Safe” links can be used to carry “TopSecret” data
- All data over “Open” links need to be encrypted

- Restrict multimedia flows in the network to max of 75% link capacity
- Allow admin traffic preferential service over network backups
- Allow user access to data only if user clearance is high enough

Simulation Toolkit: We used NS2 to simulate such an enterprise. The network topology considered was a random network with links classified with a “security” tag that defined their safety levels. We assume the nodes belong to a single Autonomous Domain (AD) and run a link state routing protocol. We modified the standard FTP/CBR applications to allow for the specification of semantic descriptions into the packet streams. For the network ontology, we used Protege as the editor for specifying our ontology. Jess was used as the reasoning engine. The choice of Jess was mainly motivated by its easy integration with Protege. Other reasoning engines can be used as a replacement if desired.

To begin, we defined an ontology to represent our enterprise. The ontology is available online at [8]. In our implementation, our ontology also contains special instances of classes representing the various actions that a PEP should take such as dropping data, encrypting data etc. These special instances also contain the low level primitive commands that need to be invoked to realize the necessary behavior. In our case, these commands are expressed as a snippet of Tcl code that can be evaluated by NS2. For example, a policy such as *All unencrypted secret data over “open” links need to be encrypted* can be expressed logically in SWRL as:

```
DataTraffic(?d) ^
datasensitivity(?d,?sensitivity) ^
Secret(?sensitivity) ^
encryptionstatus(?d,?encryptstatus) ^
UnEncrypted(?encryptstatus) ^
nextHop(?d,?nexthop) ^
securityLevel(?nexthop,?securitylevel) ^
Open(?securitylevel) ^
EncryptData(?action)
→ inferredAction(?d,?action)
```

The *EncryptData* instance has the following Tcl command encoded in it indicating the device understandable actions that need to be taken.

```
set clsfr [ get-classifier $interimRouterId ]
$ns at [$ns now] ``$clsfr install-interceptor
encryptdata $flowid $srcId $sport
$destId $dport $qdelay $overhead``
```

Using this methodology, we can now define the various actions that a Policy Enforcement Point (PEP) could take and assign to each of these actions, the corresponding primitive commands (Tcl snippets). The Policy Decision Point (PDP) was implemented as a Java process that received OWL streams from a client PEP (a network router within NS2), invoke the reasoner and send back the Tcl commands depending on the actions that needed to be invoked. The PEP (NS2) would then execute the commands received from the PDP.

B. Secure Routing

In this section, we show how security can be incorporated into the BGP routing process using our framework.

BGP Extensions: Border Gateway Protocol (BGP) was originally designed as a simple path vector protocol to share routing information between autonomous systems (AS) which has today, become the de-facto inter-domain routing protocol enabling the Internet. Autonomous systems (ISPs, enterprises etc) use policies to define how the routes are to be shared and among which peers. These policies can be driven by various factors such as commercial peering agreements, security considerations, load balancing requirements etc. These policies are then implemented in the network routers as configuration parameters to control the protocol behavior. One of the main challenges frequently faced is ensuring that network configuration settings are applied consistently throughout the network so that the correct actions are taken by the network devices both within an autonomous system and across boundaries. However, this is often error-prone and difficult to manage.

To apply our framework to provide BGP route dissemination that takes into account the security credentials and external relationships, we needed to make two modifications to the protocol. The first modification is aimed at establishing identity of the BGP peers in a secure and verifiable manner. For this purpose, we assume the BGP session establishment process is extended to include the sharing of signed credentials to validate the identity of the BGP peers and their affiliations. Prior work such as S-BGP [9] have shown that this is feasible using a public key infrastructure and signed certificates. This modification is necessary as it is important for a BGP router to establish the identity of its peer so that the routes learned from and advertised to this peer can be handled correctly. The second modification is to include with the route advertisement in the BGP update messages, an additional optional and transitive attribute that conveys semantic meta-data about that NLRI. The intent here is for the originating AS to provide this information to allow other nodes to handle this route appropriately. The interim routers are allowed to add to this description as necessary (keeping the original intact) in a manner that is secure and cannot be repudiated. In this work, we are concerned about the import/export policies in use in the BGP decision process. The modifications allow our framework to, for each route that is being advertised to or learned from, contact a PDP, the PDP will reason over the semantic information provided for that route and the policies that need to be enforced, and will communicate to the BGP node whether or not, that route can be shared or accepted.

Use Case: The use case we consider in this paper is that of a secure version of BGP where there are constraints on route exchanges between BGP peers. As with the real Internet, BGP nodes are owned by different agencies that have different affiliations. During the initial session establishment, nodes exchange their identity information to indicate the agencies to which they belong. These agencies or organizations have external socio-economic, political or financial relationships

that will influence the BGP nodes in their exchanges. Routes advertised by each AS is tagged with additional semantic information to describe aspects such as its confidentiality, sharing restrictions etc. For such a use case, the following policies would be appropriate:

- Routes marked as “ShareWithFriendly” can only be exchanged between routers that belong to organizations that have a collaborative relationship
- Routes marked as “Restricted” can only be shared between nodes that belong to the same parent organization (even if they are different divisions of that organization)
- Routes marked to be used only for data backup traffic are installed only during non-peak hours
- Allow a route to be used only for data traffic that has a specified or higher clearance level.

Simulation Toolkit: We used the ns-BGP [10] extension to NS2 to implement our framework. The network topology considered is a linear network with nodes grouped into various ASes. Each node is initialized with credentials that specify what organization the node belongs to. We modified the BGP session establishment process to allow the exchange of these credentials so that the BGP nodes can establish the identity and affiliation of the peers that they are interacting with. We added an additional optional transitive attribute to the BGP update protocol messages intended to convey additional semantic information about the route.

To begin, we defined an ontology to use for our BGP example. The ontology is available online at [8]. We modeled the various BGP protocol messages and constructs. Since we are dealing with import/export policies, we modeled special instances of classes representing the various actions that a BGP router (PEP) should take such as whether a route should be advertised or not, whether a route should be accepted or not etc. These special instances contain the low level primitive commands that need to be invoked to realize the necessary behavior. In our case, we implemented handlers in the NS2 implementation to handle the response coming back from the reasoner to determine whether a route should be included in an advertisement or whether a route that was received, should be accepted (these commands are expressed as snippets of Tcl code that are evaluated by NS2). For example, a policy such as *All routes are shareable with a peer as long as the peer and the originating router are owned by the same organization* can be expressed in SWRL as:

```
BGP_Update(?adv) ^
interimRouter(?adv, ?routeradvertising) ^
dest(?adv, ?peer) ^
owner(?routeradvertising, ?org) ^
owner(?peer, ?org) ^
AllowRouteAdvertisement(?allow)
  → inferredAction(?adv, ?allow)
```

The *AllowRouteAdvertisement* instance has the following Tcl command encoded in it indicating the device understandable actions that need to be taken.

```
set Response "OK"
```

In this case, if the reasoner asserts this rule, the corresponding Tcl command will be sent back as the reasoner’s response. Using this methodology, we can now define any arbitrary action that a PEP could take and assign to each of these actions, the corresponding primitive commands (Tcl snippets) to be executed. The PDP (reasoner) was implemented as a Java process that received RDF streams from a client PEP (a BGP agent within NS2), invoke the reasoner and send back the Tcl commands depending on the actions that needed to be invoked. The Protege IDE served the role of a Policy Editor. Using this framework, we implemented our typical use case scenario focusing on the import/export policies for BGP. For our example, we consider a network of four autonomous domains with five BGP routers. The Autonomous Domain AS0 belongs to UK forces. The Autonomous Domains AS1 and AS2 belong to two organizations within the US military. Finally, the last Autonomous Domain AS3 belongs to Russian military. During the initial BGP session establishment, the identity of each of the peers is established. This indicates the organization that the router belongs (US_{Milcom} , UK_{Milcom} , $Russian_{Milcom}$ etc) which is tracked in the “owner” property of the network devices. Some of these organizations have external relationships (such as NATO to which US_{Milcom} and UK_{Milcom} belong). Such external relationships are modeled through OWL restrictions on properties. For example, a device that is part of NATO is modeled as a one where there is a necessary and sufficient constraint that the owner is either an instance of US_{Milcom} , UK_{Milcom} or $France_{Milcom}$. Each router that originates a route includes a description that at the least, indicates the sharing restrictions for that route. In the current version, we have values such as None (which is similar to the “internet” community attribute in BGP), Restricted and ShareWithFriendly as examples. The intention here is that a route marked as “ShareWithFriendly” can only be shared with a peer who can be considered friendly. For example, if we considered forces within NATO to be friendly’s, a SWRL policy to permit the routes marked as “ShareWithFriendly” to be exchanged could be written as:

```

BGP_Update(?adv) ^
interimRouter(?adv, ?routeradvertising) ^
dest(?adv, ?peer) ^
NATO_Forces(?routeradvertising) ^
NATO_Forces(?peer) ^
routeRestriction(?adv, ?restriction) ^
ShareWithFriendly(?restriction) ^
AllowRouteAdvertisement(?allow)
  → inferredAction(?adv, ?allow)

```

Once the simulation starts, each router advertises its routes with its peers in order to compute its routing table. The simulation proceeds until all routes are computed and the routers settle on their tables. Note that when two routers belonging to UK_{Milcom} and US_{Milcom} (AS0 and AS1) are in a BGP session and while none of the routers have explicitly been identified as belonging to NATO, the reasoner can deduce this relationship and allow route exchanges between them. Similarly the reasoner can deduce that the route exchange

cannot be allowed between AS2 and AS3 as they do not have an explicit relationship that permits this.




Fig. 3. Topology

In this manner, we can now setup arbitrary relationships between routers and can specify policies through higher level rule based mechanisms to implement fine grained control over the protocol. This example can be easily extended to scenarios where the relationships are short lived and arbitrary such as in emergency response scenarios (where organizations may temporarily want to share information for providing quick response), application need driven (such as for supporting live event feeds) etc. by extending on the ontology and defining the desired policies.

VI. RELATED WORK

Policy based networks and approaches have been the focus of extensive research in recent years. Quality of service oriented initiatives such as Intserv [11] and Diffserv [12] rely on policies to drive flow classification, admission control, resource reservations etc. However, the policies used are limited in their expressibility and restricted to traffic forwarding semantics with little support for features such as content adaptation, specialized routing etc. In this respect, the Active Networks [5], [6] took the approach of allowing a more generic per packet handling semantic with the packets determining what the router should do with them, which differs from our approach in which the router (using its specified policies) controls how the packet is handled and not the other way.

There has been significant research on securing BGP. SBGP [9] proposes a comprehensive architecture for securing BGP using public key certificates. SBGP uses a pair of PKIs, one for address authentication and the other for route validation. SoBGP [13] provides more flexibility compared to SBGP. In addition to the above PKIs, a third type of certificate is used which provides routing policy and local topology. When a route is received, it is compared for consistency with the topology database and dropped if found to be inconsistent. The architecture is more flexible as there are no fixed structures of authority and ASes can decide on their own for accepting routing announcements and policies. RPSL [14] is an object oriented language for specifying routing policies from which router configurations can be automatically generated. RPSL generated router configurations can aid in preventing internet router misconfigurations but it does not support inference and is limited in expressibility.

There have been recent efforts in using the semantic web for security applications. The authors in [15] propose using a combination of conventional security mechanisms and the ability to reason about security at a semantic level for enforcing security in autonomous systems. Also, they describe a set

of requirements that need to be supported for implementing a semantic firewall. [16] proposes using context as the first principal for policy specifications governing access control in pervasive environments. Their approach stems from the fact that traditional subject/role based policies wouldn't work in pervasive environments due to the ad hoc mode of collaborations, where the roles and identities of the entities involved is not known ahead of the actual collaboration. They also propose using semantic languages for policy specification to aid in policy reasoning, conflict resolution and policy adaptation.

VII. CONCLUSION

In this paper, we have described an extensible security framework that is based on policies. Our policies are specified in semantic web languages which makes them amenable to interoperability, conflict resolution and reasoning. We described our policy based network built on top of semantically tagged packets. In our framework, applications semantically tag packets with meta-data describing the contents being carried. We then showed how our framework can be used for securing enterprise networks and BGP routing.

REFERENCES

- [1] R. Yavatkar, D. Pendarakis, and R. Guerin, "RFC 2753: A Framework for Policy-based Admission Control," Jan. 2000. [Online]. Available: <http://www.faqs.org/rfcs/rfc2753.html>
- [2] P. Kawadia, V. Kumar, "A cautionary perspective on cross-layer design," *Wireless Communications, IEEE [see also IEEE Personal Communications]*, vol. 12, no. 1, pp. 3–11, Feb. 2005.
- [3] D. L. McGuinness and F. van Harmelen, "Owl web ontology language overview," W3C Recommendation 10 February 2004, Tech. Rep., 2004. [Online]. Available: <http://www.w3.org/TR/owl-features/>
- [4] I. Horrocks, P. F. Patel-Schneider, H. Boley, S. Tabet, B. Grosz, and M. Dean, "Swrl: A semantic web rule language combining owl and ruleml," W3C Member submission 21 May 2004, Tech. Rep., 2004. [Online]. Available: <http://www.w3.org/Submission/SWRL/>
- [5] D. Wetherall, J. Gutttag, and D. Tennenhouse, "Ants: A toolkit for building and dynamically deploying network protocols," 1998. [Online]. Available: citeseer.ist.psu.edu/wetherall98ants.html
- [6] D. S. Alexander, W. A. Arbaugh, M. Hicks, P. Kakkar, A. Keromytis, J. T. Moore, C. A. Gunter, S. M. Nettles, and J. M. Smith, "The SwitchWare active network architecture," *IEEE Network Magazine*, vol. 12, no. 3, pp. 29–36, 1998, Special issue on Active and Controllable Networks. [Online]. Available: <http://www.cis.upenn.edu/~switchware/papers/switchware.ps>
- [7] S. Kodeswaran, O. Ratsimor, A. Joshi, and F. Perich, "Using semantic tags for policy based networking," in *Globecom '07: Proceedings of the IEEE Global Communications Conference*, Washington, DC, USA, 2007.
- [8] "<http://www.cs.umbc.edu/~kodeswar/ontologies/NetworkOnto.owl>." [Online]. Available: <http://www.cs.umbc.edu/~kodeswar/ontologies/NetworkOnto.owl>
- [9] S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol(s-bgp)," *IEEE Journal on Selected Areas in Communication*, vol. 18, pp. 582–592, 2000.
- [10] T. Feng, R. Ballantyne, and L. Trajkovic, "Implementation of bgp in a network simulator," in *Proc. Applied Telecommunications Symposium, ATS'04*, April 2004, pp. 149–154.
- [11] R. Braden, D. Clark, and S. Shenker, "RFC 1633: Integrated Services in the Internet Architecture: an Overview," June 1994. [Online]. Available: <http://www.faqs.org/rfcs/rfc1633.html>
- [12] S. Blake, D. L. Black, M. A. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," December 1998, status: INFORMATIONAL.
- [13] "Secure Origin BGP (SoBGP) Certificates. Internet Research Task Force, June 2003. (draft-weis-sobgp-certificates-00.txt)."
- [14] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra, "Routing Policy Specification Language (RPSL)," Internet Engineering Task Force: RFC 2622, June 1999.
- [15] R. Ashri, T. Payne, D. Marvin, M. Surrudge, and S. Taylor, "Towards a semantic web security infrastructure," in *Semantic Web Services 2004 Spring Symposium Series*, ["lib/utis:month_9040" not defined] 2004. [Online]. Available: <http://eprints.ecs.soton.ac.uk/9040/>
- [16] A. Toninelli, R. Montanari, L. Kagal, and O. Lassila, "A semantic context-aware access control framework for secure collaborations in pervasive computing environments," in *ISWC'05 : Proceedings of the 5th International Semantic Web Conference*, 2005.