

# Structure and Authentication of Policy Loads for Policy Defined Radio Systems

Stephen Berger  
TEM Consulting  
512-864-3365  
stephen.berger@ieee.org

John Chapin  
Vanu, Inc.  
617-864-1711  
jchapin@vanu.com

## Abstract

*A policy defined radio is a device whose operating parameters are constrained by downloaded policy rather than hardware limitations. Correct operation of these devices depends both on the contents of the policy and on the authentication mechanisms that assure only approved policies will be accepted. This paper explores the new requirements created by policy-defined radios and suggests a structure for policy loads that naturally fits those requirements. In addition a means of satisfying these requirements by adopting tools being used by law enforcement for computer forensics and offered by the NIST National Software Reference Library (NSRL) is explored.*

## Introduction

The core concept of policy defined radio is to break the link between the capabilities of a radio device and its operating limits. Under the policy defined radio concept, a device may have a great deal of capability but be limited in its operation to a set of policies, which are downloadable to the device.

Policy defined radio significantly improves both regulatory control and operational flexibility of radio systems. Regulatory control is improved because policies can depend on factors such as location and time of day, and policies can be changed over time as justified by changing market, societal or technical conditions. Operational flexibility is improved because network managers can exploit a single hardware device for a variety of different functions requiring different capabilities.

However, the flexibility and adaptability of policy defined radio presents new challenges. This paper describes and analyzes these challenges and suggests a policy structure to help solve them. The use of cryptographically signed files to authenticate downloaded policies is considered in detail.

## Background

### *What is a policy*

A policy is an artifact that provides a yes/no answer to the question “is it acceptable to perform a specified transmission.” A transmission is specified in part by the traditional radio parameters of center frequency, power, and modulation. There can also be nontraditional policy inputs, including time and date, geographic location of the radio, output of a spectrum sensor, and information provided by other collaborating radios.

A policy is represented as a data set that can be downloaded to the radio device, typically as a file or group of files. Within this data set there can be data (tables, strings) to be interpreted by a policy engine within the radio, or software for direct execution.

In implementing a policy-defined radio, the policy engine or software need not run before each transmission. If a series of transmissions will use the same center frequency, power, and modulation, for example, the policies need only be consulted once. For more sophisticated policies, the device may need to monitor information like time of day or geographic location against certain boundaries provided by the

policy, to know when the policy must be checked again.

### ***Comparison to SDR and other radio types***

A critical attribute of a true policy-defined radio is the ability to download multiple policies from independent sources. A transmission is only permitted if all active policies permit it. This is called *policy mixing*. The importance of policy mixing will become clear later in the paper.

A Software Defined Radio (SDR) is a device where the signal processing functions of the radio are implemented in software rather than in hardware. This enables a single hardware device to support multiple radio standards.

An SDR is not necessarily policy-defined. The purpose of SDR software is to generate and receive signals rather than to permit or reject transmission. If multiple SDR software loads are downloaded from different sources, the single load that is running at any given time normally has complete control over the transmitted signal. Therefore a traditional SDR does not perform policy mixing. Of course, a radio that uses SDR to implement its signal processing can also use policy-defined radio techniques for management.

A radio with Dynamic Access Spectrum Management (DASM) is a device which, instead of requiring a static frequency reservation, implements access rules that permit transmission under specific conditions. A policy-defined radio does not necessarily support DASM. For example, only policies that represent static frequency reservations may be allowed on the radio. Similarly, a radio with DASM is not necessarily policy-defined. The DASM rules can be implemented in local software rather than being downloadable or mixable. The term Policy-Based Adaptive Radio (PBAR) has been identified in the current draft of IEEE 1900.1 to identify implementations that are both policy-defined and adaptive.

A Cognitive Radio uses autonomous goal-directed reasoning, also known as artificial intelligence, for local control. This enables the radio to learn from events and thereby adapt more effectively to user needs and environmental challenges. A radio can be cognitive without being policy-defined and vice versa. However, cognitive radio techniques are widely considered to be a promising approach for implementing sophisticated policy engines for policy-defined radio, so there is a close link between the two approaches. However, the goal-directed reasoning is bounded such that regulatory authorities may confirm

that the adaptability will never configure the radio for a disallowed transmission state.

### ***Conformance assessment***

Most regulatory structures follow the general requirements of ISO 17011 in designing conformity assessment systems. It is assumed that policy defined radio systems will be required to comply with regulations and other requirements generally structured under the guidance provided in ISO 17011 and its companion documents. Hence, to receive regulatory approve some key questions must be satisfactorily answered before these systems will be permitted. Among these questions are:

1. What are the requirements for a minimal acceptable system?
2. Are the testing lab/testers/lab assessors qualified to effectively evaluate designs?
3. Will the vendor deliver units within manufacturing tolerances to those evaluated?
4. How will regulatory officials know if non-compliant units are delivered and what corrective actions can it take?
5. Are there adequate safeguards that the systems will be used as intended?

Within question 1 of this general framework, policy-defined radio creates a challenging set of new problems for conformance assessment, especially when compared to traditional radios.

- *Correctness of a policy artifact*: Does the artifact actually represent a legal and desirable policy? Does the artifact comply with local regulatory requirements?
- *Correctness of a policy engine*: If the policy artifact is data rather than code, does the policy engine in the device behave correctly when interpreting that data?
- *Correctness of policy support in the device*: Does the device provide correct inputs, e.g. do all sensors behave as required. Also, does the device correctly respond to policy decisions, e.g. does it halt transmissions when rejected by the policy?
- *Correctness of policy use by the device*: Does the device check the active policies at the required operational points?
- *Device integrity*: Does the device sufficiently defend against a malicious user or third party who seeks to corrupt or bypass policy processing?
- *Policy authentication*: Does the device sufficiently defend against attempts to download unapproved or corrupted policies?

This paper only considers the last requirement, assuring the radio operates only with approved policies. The other conformance requirements are equally critical and are the subject of active work by multiple researchers, for example in the DARPA XG program.

## Structure of policy loads

### *Rationale*

Regulators have different levels of concern for different aspects of device behavior. Some requirements, such as those that relate to operator safety from RF exposure, require high levels of assurance. Other requirements, while still important, require lower levels of assurance. In particular if effective means of field management are provided some risk of violation of this class of lower level parameters may be allowed.

An example of a less stringent requirement might be international unintentional emission requirements under the IEC standards. When considering the issue of manufacturing tolerance of devices the IEC standards give guidance for statistical sampling of production lines. This guidance finds that a statistical assurance that 80% of a production line will be compliant with the unintentional emission requirements 80% of the time is considered satisfactory compliance. The issue in this case is to keep the required sample testing to reasonable limits. If higher levels of assurance were required large numbers of devices would require testing simply to prove compliance. Such additional testing may have arguably little value. The point being made is that there are different levels of assurance required by regulators for different requirements.

### *Structure*

As the requirements of policy loads for software defined radio are considered there would seem to be three levels of policies:

1. Some policies are critical and require strong assurance that they will never be changed in deployed devices. Examples might be that a handheld device will never exceed the RF safety limits. Another might be that a device intended for use on planes will never transmit on radio navigation frequencies. For these policies, called in this paper Critical Policies, there must be very strong safeguards. A characteristic of some and perhaps all of these policies is that they apply widely, even globally. As examples the RF safety and unintentional emission

requirements are widely implemented using international standards.

2. Of only slightly less concern are what may be termed Regulatory Policies. These policies contain the basic operating constraints established by regulatory authorities. These also require significant safeguards and the ability to verify that deployed devices in fact are operating ONLY with approved policy sets.
3. The third group of policies is called here Management Policies. These policies are available to network managers for use in optimizing the systems under their supervision.

A logical priority is apparent in this organization. A Regulatory or Management policy must never override a Critical policy. Similarly a Management policy must always operate within the bounds of both Critical and Regulatory policies.

Policy Level	Scope	Assurance	Flexibility
<b>Critical</b>	<b>Global or Regional</b>	<b>Very High</b>	<b>Low</b>
<b>Regulatory</b>	<b>National</b>	<b>High</b>	<b>Moderate</b>
<b>Management</b>	<b>Local</b>	<b>Moderate</b>	<b>High</b>

### *Policy mixing*

One of the goals of this structure is to allow great flexibility to network managers while giving regulators the assurance that devices will only operate within prescribed parameters.

In contemplating this goal it becomes clear that there are tendencies to separate these levels of policies by their scope of application. Critical policies tend to have very wide, even global scope. As an example, all nations require that RF exposure limits be observed and despite national differences, there are international recommendations on those limits. In contrast Management policies would normally only apply to local systems and be modified between systems to optimize performance.

This organization suggests the possibility of levels of control and flexibility. Critical policies will require the highest level of control and assurance while Management policies may require only moderate control and offer great flexibility. If there is good

confidence that Critical and Regulatory policies can be relied upon to properly restrict device behavior then great flexibility and freedom may be given at the level of management policies.

The policy mixing mechanism defined earlier becomes valuable in this context. Critical, regulatory and management policies will be authored independently and loaded into the device at different times. Policy mixing assures that each policy load has the independent ability to consider and potentially reject each specified transmission.

## Authentication of policy loads

### *Threat model for policy authentication*

Policy authentication in a policy-defined radio defends the radio against download of corrupt or unapproved policies. As a basis for assessing policy authentication mechanisms, we categorize the threats that could lead to incorrect download attempts.

- *Accident:* The owner or operator of a device accidentally presents a policy to the wrong device, or there is a configuration management error somewhere.
- *Cheating by user:* The user seeks to improve performance or access capabilities beyond what is permitted by approved policies.
- *Cheating by manufacturer:* The manufacturer seeks to improve performance or access capabilities beyond what is permitted by approved policies.
- *Local attack:* The user or other person with physical access to the device seeks to cause it to behave in a harmful manner.
- *Remote attack:* A person without physical access to the device seeks to cause it to behave in a harmful manner.

Clarifying the threat model enables considering which attacks should be considered and which ignored when designing the policy authentication mechanisms.

For example, a cheating manufacturer could design the device to internally modify or ignore a downloaded policy. Possibilities like these mean that cheating by manufacturer is better addressed through legal deterrents than through any technical mechanism. As a result this threat is not considered in the policy authentication mechanism.

While a local attack will be rare, cheating by the user is fundamentally the same case and is likely to be a significant problem. Therefore the policy authentication mechanism must defend against this.

In particular, the traditional authentication mechanism where a policy is accepted if and only if it is presented across a hardwired local connection (rather than delivered remotely across the network) is not an effective mechanism for policy-defined radios.

The most common threats are likely to be accidents and remote attacks, both of which need to be considered.

### *Public key cryptography*

Public key cryptography offer basic mechanisms that can be leveraged for policy authentication. With public key cryptography, the approved policy author encrypts the policy using their secret key. Their public key is well known, for example stored in the radio device, and successful decryption using that public key provides proof that the policy could only have come from the approved policy author.

A public key approach could be incorporated into a regulatory or certification scheme whereby only a highly trusted source could encrypt a policy, especially a critical or regulatory policy. It is possible that the only method for having a policy properly encrypted for transmission would be at the end of a regulatory approval or certification process. Used in this way the regulatory or certifying authority would have assurances that disallowed policies could not be deployed.

## Interrogation of device policies

In addition to authentication of policy downloads by the devices receiving them, cryptographic mechanisms can also be used to learn externally what policies a device has loaded. This supports a number of user requirements.

Cooperating devices will require a means for validating that other devices in their local system are using identical or compatible policy loads.

System or network managers will require a means for knowing what devices are operating on their system and what policy loads they are using. When system operation must be modified the manager will require a means for confirming that updates have been successfully received and implemented.

### *Digital signatures*

The basic mechanism supporting policy interrogation is the use of digital signatures produced by HASHing algorithms. A digital signature is a short number, typically 128 to 1024 bits, computed from an original file, with the following characteristics:

1. The original file cannot be reproduced from the digital signature. Thus the original file is secure and remains confidential even though the signature may be made widely available.
2. It is computationally infeasible to create a new file that computes to the same digital signature as an existing file. These algorithms truly produce a unique signature or “fingerprint” of the file.

**Policy interrogation based on signatures**

By use of digital signatures policy sets can be uniquely identified with high confidence. To achieve this, a list of signatures of validated policy sets must be available from a trusted source.

This could be arranged by having regulatory authorities escrowed policy files as part of the approval process and issue signatures on the escrowed files. Further, policy defined radios can be required required to have a capability to HASH their policy sets and transmit those codes (either over the air or over a local connection) when requested.

This would provide a means of verifying a radio’s policy set without knowing the details of the policy set. Regulators and network managers would be able to check deployed systems and gain assurance that they were operating with approved policy sets.

If devices transmitted with HASH codes when given an authorized command to do so then regulators and network managers would be able to verify the systems operating in their vicinity and use this tool when dealing with network or regulatory issues.

**Parallel of the NIST NSRL**

An important model to be studied for application to policy defined radio is the NIST National Software Reference Library (NSRL). The NSRL is a project supported by the U.S. Department of Justice's National Institute of Justice (NIJ), the FBI and other federal, state, and local law enforcement agencies, with the National Institute of Standards and Technology (NIST) to use computer technology in the investigation of crimes involving computers.

The NSRL is designed to collect software, produce profiles of those files and incorporate the profiles computed from this software into a Reference Data Set (RDS). The RDS can be used by law enforcement, government, and industry organizations to review files on a computer by matching file profiles in the RDS. This information is used by law enforcement when analyzing computers or file systems that have been seized as part of criminal

investigations. Standard program files and other files that are known and unaltered can be eliminated, leaving only files that may contain evidence.

The RDS is a collection of digital signatures of known, traceable software applications. There are application hash values in the hash set which may be considered malicious, i.e. steganography tools and hacking scripts. Hence the mechanism may be used to separate known trusted files, from known malicious files and unknown files.

Reference Data Set version 2.10 was released in September 2005 containing 10,663,650 unique digital signatures for 33,860,009 files. The signatures are produced using SHA-1, MD5 and CRC32 algorithms.

HASH Examples of Different Microsoft Notepad Versions <sup>1</sup>		
Version	Bytes	SHA-1
NT4\ALPHA	68368	F1F284D5D757039DEC1C44A05A C148B9D204E467
NT4\I386	45328	3C4E15A29014358C61548A981A4 AC8573167BE37
NT4\MIPS	66832	33309956E4DBBA665E86962308F E5E1378998E69
NT4\PPC	68880	47BB7AF0E4DD565ED75DEB492 D8C17B1BFD3FB23
WINNT31.WK S\I386	57252	2E0849CF327709FC46B705EEAB5 E57380F5B1F67
WINNT31.SRV \I386	57252	2E0849CF327709FC46B705EEAB5 E57380F5B1F67

At this time NIST believes that the SHA-1 has not been broken, however there are known MD5 collisions and weaknesses. The NSRL data provides an MD5 to SHA-1<sup>2</sup> mapping to facilitate the migration away from MD5. The SHA-1 algorithm will be superseded in 2010 by FIPS 180-2, Secure Hash Standard, which contains SHA-224, 256, 384 and 512. The NSRL plans to provide a SHA-1 to SHA-256 mapping.

<sup>1</sup> From Douglas White’s White, Douglas, Presentation to the EAC TGDC, July 9, 2004,  
<sup>2</sup> Secure Hash Algorithm (SHA-1) is specified in FIPS 180-1. It is a 160-bit hashing algorithm which performs 1045 combinations of 160-bit values to produce a unique digital signature or “fingerprint”.

Currently the NSRL is used by ISPs to track application sharing on servers. System administrators are also using this tool to confirm valid operating system file states on machines in their network. These applications have many similar characteristics to the needs of policy defined radio systems and suggest the possibility of a positive adoption of a modified version for the needs of policy defined radio systems.

## Conclusions

### *Potential for standardization*

If a standard in the IEEE 1900 series were to be developed to support this approach to the structure and authentication of policy loads, it might be outlined as follows:

1. Policy categories and contents of each category.
2. Protections required of each category of policy.
3. Public key cryptographic algorithms to be used for authentication
4. Digital signature algorithms to be used.
4. Format for requesting transmission of HASH codes and format for their transmission.
5. Informative annexes outlining how this functionality might be utilized by regulators, manufacturers and network managers to facilitate their work and assure the integrity of the system.

### *Closing Summary*

This paper has reviewed the requirements for conformance of policy defined radio systems, with a particular focus on authentication of policy downloads. A hierarchy of policies is proposed containing three categories of radio policies. Parallels with the needs of law enforcement in computer forensics, as addressed in the NIST NSRL, are explored for application to the needs of policy defined radio. Mechanisms for providing the required security, verification and independence of various policy levels and components of the verification system are then discussed. Taken together, the hierarchy, use of cryptographic mechanisms and implementation protections are proposed as a possible means of addressing the requirements of policy provisioning systems.

## References

White, Douglas, "NIST National Software Reference Library (NSRL)", Presented at the Mid-Atlantic Chapter HTCIA Meeting , September 28, 2005, VA<sup>3</sup>

White, Douglas, "National Software Reference Library (NSRL)", Presentation to the EAC TGDC, July 9, 2004, Washington D.C.

---

<sup>3</sup> NIST NSRL presentations are available at: <http://www.nsrl.nist.gov/Presentations.htm>