

Application of the Flask Architecture to the X Window System Server

Eamon F. Walsh

National Information Assurance Research Laboratory

National Security Agency

`ewalsh@tycho.nsa.gov`

Abstract

This paper will outline the progress that has been made on extending the coverage of Security-Enhanced Linux access controls to the X Window System server, a major component of the Linux desktop. This has been accomplished by applying the Flask architecture to the X server and extending the reach of SELinux policy to cover X server objects. Modifications have been made to both SELinux library and the X.Org X server implementation in support of this goal. In the SELinux library, improved capabilities for obtaining policy decisions from the kernel were added. In the X server, a set of general security hooks was added, followed by a Flask module which makes use of them. This module extends the enforcement of kernel-based security policy to the X server in userspace, providing fine-grained access and information flow control to this vital desktop component using the existing SELinux policy store and toolchain.

1 Introduction

The Flask architecture, as currently implemented in SELinux [5], lacks control over key subsystems of modern GNU/Linux distributions that are implemented in userspace. These subsystems manage security-relevant objects, provide information transfer facilities, or both, but for historical or technical reasons are not part of the Linux kernel. Extending the Flask architecture to these subsystems is necessary to fully secure the desktop system. These extensions or “userspace object managers” are a current area of development activity, involving both the userspace subsystems themselves and supporting infrastructure in the kernel and in the SELinux userspace libraries. This paper describes the progress on perhaps the most important userspace subsystem to receive such attention, the X Window System.

The X Window System serves as the graphical user interface (GUI) foundation for desktops such as GNOME and KDE [15][16]. A userspace program, the X server, manages the desktop system’s display, draws windows, cursors, and other GUI constructs, delivers mouse and keyboard events to GUI applications, and provides basic clipboard and drag & drop support. Client applications, usually with the help of an X library and a toolkit such

as GTK or Qt, connect to the X server and issue commands to create and manipulate their GUI [17][18]. The protocol used by X clients and servers to communicate is extensible, and new command sets have been added to support shared memory, advanced rendering, hardware-accelerated graphics, and other features.

Existing security mechanisms in X mostly involve authenticating clients at connection time. Once connected, however, there are few controls in place to prevent client applications from engaging in malicious behavior. The X protocol allows client applications to manipulate windows belonging to other clients, including moving them, reading their contents, drawing into them, changing their focus, or listening for keyboard and mouse events being sent to them [7]. Clients are also permitted to send events, including fabricated keyboard or mouse input, to other clients. Mechanisms exist for arbitrary data transfer between clients, such as by setting and reading window properties. Clearly, access controls are needed on windows and other X server objects to prevent malicious behavior and control the flow information between X clients.

The Linux kernel provides limited support to the X server, in the form of access to video RAM, but otherwise has no knowledge of the X server’s internal objects such as windows. To the kernel, X client applications are simply normal processes which have a socket connection to the X server over which opaque data is being transferred. Current SELinux policy can prevent client applications from connecting to the X server entirely, but this solution is too coarse-grained from a usability perspective. What is needed are fine-grained controls on individual X server objects. An application of the Flask architecture to the X server itself will provide these controls, after which SELinux policy can be written for X objects and the X server made a policy enforcer in the same manner as the kernel.

Several steps are required to achieve this goal. A mechanism for obtaining policy decisions in userspace must be established, and supporting infrastructure must be provided in the SELinux libraries. Flask object classes and permissions must be selected to allow natural, comprehensive policy expressions governing X protocol operations, including protocol extensions. En-

forcement logic of some form must be implemented in the X server and Flask semantics implemented. Finally, appropriate policy must be written. The following sections will discuss completed and ongoing work in all of these areas.

2 Background

2.1 Library Support

The Flask architecture requires a separation between policy decisions and their enforcement [10]. In the SELinux runtime environment, policy decisions are obtained from the “security server” which is part of the kernel [5]. One of the key decisions made in the planning stages for the X work was to continue using the kernel security server to obtain policy decisions. This allows the existing SELinux policy toolchain to be reused and keeps policy centralized in a single place [8]. However, it required the development of supporting infrastructure to allow efficient retrieval of policy decisions from userspace, as well as notification of significant policy events such as reloading or invalidation.

The `security_compute_av` SELinux library call may be used to obtain policy decisions from the kernel. However, this interface returns raw decision vectors rather than a simple yes/no answer, provides no auditing, and does not cache decisions, requiring a resource-intensive trap into kernel space on each use. These issues were resolved by porting the existing access vector cache (AVC) from the kernel into userspace, making it a part of the SELinux library. The userspace AVC uses `security_compute_av` internally but provides an improved interface to the user, including automatic logging and caching of decisions in the same manner as the kernel AVC.

The `security_compute_av` call allows synchronous retrieval of policy decisions, but some policy events, such as policy reloads, are asynchronous. The userspace AVC must be made aware of these events in order to discard any cached decisions that may have been rendered invalid. The existing `selinuxfs` interface being insufficient for asynchronous communication, a new, netlink-based interface was introduced. An SELinux message family was created along with message types for enforcement mode changes and policy reloads. The userspace AVC was then enhanced to listen for these messages, optionally on a dedicated background thread, updating its cached decisions as appropriate.

The userspace AVC is substantially complete, and has been a part of the SELinux library (`libselinux`) since version 1.8. However, userspace object manager support in `libselinux` is not yet complete; refer to Section 5.1 for a discussion of the proposed labeling interface.

2.2 X Server Security Hooks

SELinux in its current form relies on the Linux Security Modules (LSM) project to provide access to key decision points throughout the kernel [9]. LSM provides general-purpose security hooks and a number of security projects besides SELinux have made use of them. Since X is a general-purpose windowing system just as Linux is a general-purpose operating system, the provision of LSM-style security hooks was determined to be the best method to provide enforcement logic in the X server. These hooks were added by what is now known as the X Access Control Extension (XACE).

The development of XACE was simplified by an existing security extension (“Security”) developed by the X Consortium in 1996 [12][13]. The Security extension provides a simple two-level trust hierarchy for client connections, where “untrusted” clients are restricted in several ways. The two-level trust model is too coarse-grained for general use, but the extension authors had conducted an analysis of the core X protocol, identifying places in which untrusted clients should be restricted and introducing checks into the X server code at those places. Much of the XACE development process consisted of simply replacing those checks with more generic callbacks. The Security extension was then rewritten to use the new callbacks, maintaining full backwards compatibility.

XACE has been accepted into the X.org mainline as of xserver release 1.2. A full discussion of the capabilities of XACE is beyond the scope of this paper [11], and the set of hooks may change slightly as work progresses (refer to Section 5.2). However, two specific areas where XACE provides coverage are worth discussing in detail, because they together provide the basis for nearly all of the enforcement activity of the Flask module discussed in the next section.

The first of the two areas is control over access to X server resources. Most X server objects or “resources,” including windows, pixmaps, cursors, fonts, and colormaps, are assigned unique ID numbers and stored together in a large hash table. The ID numbers include space for a client index number, allowing resources to be assigned a client “owner,” usually the client responsible for creating the resource. The resource system is extensible, allowing X protocol extensions to create new resource types for objects that they introduce [1]. Clients refer to resources by ID number when making protocol requests; the request handling code then calls a lookup function to retrieve a pointer to the object itself. An XACE hook is present in the lookup code. The hook includes as parameters the resource ID (from which the client owner and type of resource can be ascertained) and the client on whose behalf the lookup is being made [13]. This important hook allows security modules to vet any

and all resource lookups.

The second of the two areas is control over the X protocol dispatch table. Briefly, X protocol requests include major and minor codes. The major code specifies the protocol extension, with the first several major codes reserved for “core” X protocol requests. Major codes are assigned to extensions dynamically, but extensions have fixed names which can be checked to determine which extension is assigned to a given major code [7]. The minor code, by convention, specifies the individual request within the protocol extension. All incoming requests are dispatched through an array indexed by major code, containing function pointers to request handlers. For core protocol requests, the request handler processes the request immediately. For extension protocol requests, by convention, the request handler switches on the minor code, calling a second handler which performs the specific request.

XACE fills the server’s dispatch array entirely with calls to an XACE intercept function. This function calls an XACE hook, allowing security modules to examine and potentially reject all incoming requests before they are dispatched to the actual request handlers. Using the published protocol specifications, security module authors can write code to parse the incoming requests, checking resource ID numbers, flags, and other parameters. This powerful capability is used extensively by the Flask module.

2.3 X Server State Storage

XACE does not provide a mechanism for attaching state (labels) to X server objects. This mechanism is provided through a separate subsystem, `devPrivates`, which was originally intended for device driver writers to use for storing private data [1]. Certain server objects possess a `devPrivates` field, which points to a dynamically allocated array of generic value/pointer unions. At initialization time, drivers, extensions, and other modules can register for a slot in this array, and can additionally specify an amount of memory which will be allocated and referred to by the generic pointer in that slot. When object instances are created, the registrations are used to allocate the array and any extra space requested. In this way, `devPrivates` provides the ability to hang extra data from certain server objects.

The objects supported by `devPrivates` include the client structure itself, which is the main structure created when a new client connection is made. Device-related structures, including the `ScreenRec` object which represents a single screen, and the `DeviceIntRec` object which represents an input device, are also supported. The `ExtensionEntry` structure that represents each protocol extension is supported (this support was added as part of the XACE work). Finally, some resource objects, in-

cluding window and pixmap objects, are supported.

However, many resource types do not include a `devPrivates` field in their structure definitions. Minor objects such as individual window properties, and ephemeral objects such as event messages are likewise not supported. Extending the reach of the `devPrivates` mechanism to these objects is a priority; refer to Section 5.2.

3 Flask Module

The first step in the development of a Flask module for X was to determine the appropriate set of object classes and permissions for X. In 2003, Kilpatrick, Salamon, and Vance described a set of object classes and permissions¹ for use in securing the core X protocol [4]. This set is the one in use today, with the following deviations:

- A GC object class was introduced, representing graphics contexts, one of the basic server resource types. This class was added mainly for completeness, so that all the base resource types would be covered.
- A Property object class was introduced, representing named window properties. Objects of this class are meant to be labeled with a type derived from the name, so that for example, the standard `WM_CLASS` property would be labeled `wm_property_t`. The class works in conjunction with the `chprop` permission on the window class to allow fine-grained control over which properties client applications can read and write on given windows.
- An Extension object class was introduced, representing named X protocol extensions. Like the property class, objects of this class are meant to be labeled with a type derived from the name. For example, the `XKEYBOARD` and `XInput` extensions would be labeled `input_ext_t`. Client applications are denied use of any protocol requests belonging to an extension unless they have the `use` permission on that extension, and may not query the extension’s existence unless they possess the `query` permission. Restricting entire extensions at a time is coarse-grained, but is a useful mechanism particularly for extensions that have not yet been analyzed to determine what fine-grained checks to perform on each protocol request.
- Several permissions were added to the Window object to control certain core protocol requests; these are `setfocus`, `transparent`, and `mousemotion`. Having the `transparent` permission allows clients to create windows with no background or otherwise unfilled so that windows be-

¹The full list of classes and permissions is not enumerated here but can be found in reference [4].

neath it show through. Another new permission `extensionevent` was added to control permissions for sending events defined by protocol extensions. This is likely too coarse-grained to cover all such event types, but can serve as a fallback for those events which have not been categorized into the other event-sending permissions.

- A few other miscellaneous permissions were added.

The next step was to begin implementing the Flask module as an extension to the X server. Extensions make use of internal X server API's to initialize themselves, obtain a major code and provide a dispatch handler for new protocol requests, and register callbacks for various events, including XACE hooks. Extensions can be built as loadable modules provided that they rely only on the published API's and do not change the core X server code; the Flask module meets these requirements. All loading and initialization of extensions occurs before the server begins accepting client connections; clients cannot interfere with this process.

One of the first operations performed in the module initialization code is to register for space in server objects through the `devPrivates` mechanism described above. Of the eleven X object classes, ten correspond directly to internal objects, and ideally all instances of those objects would carry labels. However, due to the current limitations of `devPrivates`, many of the objects in question, including resource objects such as font, cursor, and color, do not have `devPrivates` support and cannot be directly labeled. For the time being, the Flask module maintains a list of labels attached to the main client object, one for each object class. The trouble with this scheme is that it limits to one label all objects of a given class belonging to a given client. Thus it is not currently possible for a client to have, for example, two windows of different types.

When a new client connection is made, the Flask module is notified via a client state callback. The callback function obtains the security context of the new client by calling the `getpeercon` SELinux library call on the client's socket descriptor (if the connection is from a remote machine, a default context is used). Then, successive `security_compute_create` library calls are used to determine labels for each of the client's object classes and these are stored in the list attached to the client structure. Once `devPrivates` support is available, this step will be performed at object creation time, and the label will be stored with the object.

A peculiar concept in the X server is that of the "server client." The X server itself owns resources and other objects, notably the root windows on each screen. A fake client object called the server client is present as a stub in the resource system to server as the owner of

these objects. The Flask module treats this client the same way as regular clients, but uses the context of the X server process as the starting point for computing labels.

The property and extension object classes are not associated with a particular client, but rather are labeled based on the name of the object instance as described above. When labeling these objects, the Flask module combines the user and role fields from a base context with a type looked up from the name. The base context used for extensions is that of the server process, and the one used for properties is that of the client object owning the window to which the property is attached. The type-to-name mapping is kept in a configuration file read by the Flask module on startup, although ideally this information would be maintained as part of the SELinux policy configuration in the same manner as the `file_contexts` database, which serves a similar purpose.

Also at initialization time, the Flask module registers callbacks on XACE hooks, the major one being the dispatch intercept hook described above. After this is done, initialization is complete and the X server proceeds into its dispatch loop, waiting for client requests to arrive. For the most part, the permissions on each object class correspond closely to the core protocol requests involving that class, and this is reflected in the dispatch-oriented structure of the module's enforcement code.

For example, suppose a `ChangeWindowAttributes` request arrives at the server. This request contains the resource ID of a window along with new window attributes to be set. The XACE dispatch intercept hook calls the Flask module callback, which drops through a switch statement to the `ChangeWindowAttributes` handler. That handler parses the incoming request, first determining the context of the client making the request (the source), then extracting the window ID and from it determining the client owner of the resource. The list of labels in the owner's client structure is consulted to determine the context of the window (the target). Finally, a call to `avc_has_perm` is made, passing the source and target contexts, window class value, and the `setattr` permission. If the result is a denial, XACE returns a `BadAccess` error to the client, otherwise the request continues on to its "real" handler. All of the core protocol requests are handled in this manner, and extension protocol requests will likely be handled in the same manner.

The Flask module also includes preliminary support for window labeling. Via an XACE hook, a callback function is called whenever new windows are created. That callback function sets a property on the window containing the window's security context, and in the future additional properties may be added to communicate other contexts, such as that of the client connection own-

```

0  # App can use cut buffers
1  allow app_t cut_buffer_property_t:property { read write };
2  allow app_t root_window_t:window { listprop chprop };

```

Figure 1: Example cut buffer access policy.

```

0  # App can capture window contents
1  allow app_t domain:drawable copy;
2  # App can create windows with no background
3  allow app_t self:window transparent;

```

Figure 2: Example screen capture policy.

ing each window. The property name begins with an `_SELINUX` prefix and it can be protected from modification via SELinux policy. Window managers can be modified to display the property contexts alongside or in lieu of the normal window title; Figure 4 (located on the last page) shows a screen shot of a modified `twm` that does this. Secure window labeling is discussed further in Section 5.4.

4 Policy

In the following sections, sample policy is presented to address some different security goals in X. Some assumptions are made in the policy fragments:

- Some type definitions and other statements are omitted to save space.
- Window objects are labeled directly with the owning process domain.
- The configuration file that maps extension and property names to the associated types is not shown.

The policy fragments are meant to be examples and are not comprehensive. As with all SELinux policy, all actions not explicitly allowed are denied. The examples thus consist of allow rules expressing the actions which we wish to permit.

4.1 Clipboard Access

The policy in Figure 1 allows a client application to access the cut buffers, which consist of 8 properties on the root window. They are intended for use as a simple 8-slot clipboard [3].

Line 1 grants the domain access to cut buffer properties, while line 2 allows reading and writing of properties on the root window. Removing the `chprop` and `write` permissions would permit reading from the clipboard only.

X Windows provides another clipboard mechanism, selections, which is more complex than cut buffers. Policy coverage of selections is possible but not covered

here.

4.2 Screen Capture

The policy in Figure 2 allows a client application to capture screen contents.

Line 1 by itself allows the domain to use the `GetImage` and `CopyArea` core protocol requests on all application windows, allowing their contents to be captured. This is the mechanism used by the GIMP application’s window capture feature. However, there is a “back door” method for capturing contents: create a window with no background, which causes the windows beneath it to show through, and then copy its contents [12]. Line 3 allows the application to create such no-background windows; denying this permission would prevent their creation.

The three capture methods described here are the only ones the author is aware of in the core protocol. There may be other methods available through extensions; it is hoped that the `copy` permission will cover all direct capture methods while `transparent` will cover any indirect methods such as alpha blending or translucency, which are becoming common.

4.3 Window Managers

Window managers in X are regular client applications that control the surroundings and placement of other windows on the screen. Window “decorations” such as title bars, borders, and resize handles are drawn by the window manager. Window managers have a great deal of control over other application windows, reparenting, moving, hiding, and resizing them as necessary. Because of this, it would be beneficial to run window managers in a separate domain from regular clients. Refer to Figure 3.

Line 1 allows the domain to access X protocol extensions designed specifically for window managers. Line 4 grants access to window manager related properties while line 5 grants permission to read and write properties on any application window; together these two lines

```

0  # Extensions
1  allow wm_t windowmgr_ext_t:xextension use;
2
3  # Properties
4  allow wm_t wm_property_t:property { read write };
5  allow wm_t domain:window { listprop chprop };
6
7  # Windows
8  allow wm_t domain:drawable getattr;
9  allow wm_t root_window_t:window { enumerate setattr };
10 allow wm_t domain:window { enumerate getattr setattr };
11 allow wm_t domain:window { map unmap move ctrlldlife };
12 allow wm_t domain:window { windowchangeevent clientcomevent };
13 allow wm_t domain:window { chparent chstack };
14
15 # Input
16 allow wm_t domain:window setfocus;
17 allow wm_t root_window_t:window setfocus;
18 allow wm_t xserver_domain:xinput setfocus;
19 allow wm_t xserver_domain:xinput warppointer;
20 allow wm_t xserver_domain:xserver { grab ungrab };
21 allow wm_t xserver_domain:xinput { activegrab passivegrab ungrab };

```

Figure 3: Example window manager domain policy.

grant access to window manager related properties (and only those properties) on all application windows. An example of such a property would be `WM_NAME`.

Lines 8-13 grant extensive control over application windows, including the ability to move, hide, and re-parent them, change the stacking order, and send notification messages of these activities to clients. Lines 16-21 grant permission to change the input focus to any window, move the mouse cursor, and create “grabs” on the server, which are used to redirect or temporarily interrupt input events.

The power required by window managers warrants a thorough review of any candidate before admitting it to the domain. However, the limited number of such programs in common desktop use should make this a tractable task.

4.4 Events

Another area of concern is the malicious sending of X events to windows belonging to other clients. This sort of behavior enables the “shatter” attack which has been demonstrated on Microsoft Windows systems. This attack involves sending malicious configuration events to a window owned by a privileged process, causing a buffer overflow that allows arbitrary code to be run as the privileged user. [6].

It should be noted that X events are part of a wire pro-

tol, not an internal API such as in Windows, nor can X events used to configure individual text fields or other widgets as in done in the shatter attack. However, there is a `SendEvent` protocol request which allows clients to send arbitrary events, including fabricated keyboard and mouse input or other unexpected notification normally generated only by the server.

In Kilpatrick et. al., the various core protocol events were grouped into rough categories which are expressed as permission bits on the window object class [4]. Thus, sending a `KeyPress` event to a window requires the `inputevent` permission on that window. In Figure 3, line 12, the window manager domain is granted permission to send events from two different categories to all application windows.

In this manner, SELinux policy may be used to control the sending of events. In the future, however, the category model may be discarded in favor of X events as a distinct object class, labeled based on the event name (or number) in a similar manner to the property class.

5 Future Work

5.1 Library Interfaces

As discussed previously, the property and extension object classes are labeled with a type that is derived from the name of the object. Types are defined in the SELinux

policy, but currently, the mapping from names to types is kept in a configuration file that is part of the X.org code base and installed as part of the X server. The X Flask module must load this file and parse its contents on server startup, and the parsing code constitutes a large part of the module at present.

The X server is not the only instance of a userspace object manager needing a string to type mapping. The D-Bus daemon uses such a mapping to label D-Bus messaging channels from their names, and the file contexts configuration consists essentially of such a mapping with regular expression matching semantics. In general it is likely that there will be a need for more such mappings as more userspace object managers are developed.

The author is developing a standard mechanism for use in querying such mappings, which would be a part of libselinux and store the mapping data in the policy configuration in a similar manner to the file_contexts data. This would work well with a modular policy that ships policy modules with each application, and would allow userspace object manager code to call a streamlined API instead of having to load and parse configuration data on a per-manager basis. This “labeling API” is a work in progress and early versions of it have been posted to the SELinux mailing list for review.

5.2 Security Hooks

As the XACE framework matures through use in the X community, the set of security hooks it provides will likely change to meet the needs of new and existing X server extensions and drivers, as well as other projects which may find it useful. For example, it is possible that the Solaris Trusted Extensions for X may make use of XACE to assist in an upstreaming to X.org [2]. XACE hooks that were introduced to support the legacy extensions Security and Appgroup may be removed, and others may be moved out of XACE and made a part of other X server mechanisms, such as the commonly used client state callback [1].

As discussed in Section 2.3, the devPrivates state storage mechanism must be extended to additional server structures to support full labeling of server objects and resources. Another problem with devPrivates is that the mechanism is not consistent from object type to object type. For example, the devPrivates support for the colormap object includes an initialization callback function, while other objects do not. Work is necessary to unify and extend this important supporting interface.

5.3 Policy

The author is preparing patches which will add support for X to the Reference Policy. However, in light of the continuing work on the supporting infrastructure and the

extension of the Flask module’s coverage to additional extensions, it is unlikely that the current set of object classes and permissions will remain unchanged. Policy development itself also reveals areas in which the set needs refinement.

Providing MLS support for X is an area that has not yet been seriously investigated. Because of data transfer mechanisms within the X server, notably the clipboard, an MLS desktop system will almost certainly require support within X and other desktop layers.

5.4 Trusted Labeling & Input

The preliminary labeling scheme discussed in Section 6 relies on the window manager to obtain the label from a property attached to each window and display it in the window’s decoration. This scheme is subject to spoofing attacks, since a malicious client application could recreate window decorations itself, misleading or confusing the user.

A more secure method would be to reserve an area of the screen for displaying labels. This area would be off-limits to client drawing; the server itself would be responsible for drawing the labels as the input focus changes from window to window. This scheme is employed by Solaris Trusted Extensions for X [2].

Secure input, input event labeling, and trusted path are areas that need addressing. However, the input subsystems in the X.org X server are in a state of churn as new features are added. For example, recently improved device hotplugging support was added, which has resulted in deep changes to the server. Other proposals on the table include support for multiple concurrent mouse pointers and new ways for selecting input focus on windows for use in 3D environments. This author does not plan to study the X input model in depth until development has settled down.

6 Conclusion

Application of the Flask architecture to the X server is a key security development that provides a foundation for securing the Linux desktop. Vigorous work on the projects described in this paper and in other userspace object manager and desktop-related areas is expected in the coming months of 2007.

As mentioned, the XACE framework has been accepted into the X.org mainline for release 1.2 of the X server. The current, tentative target for acceptance of the Flask module is release 1.3, which is scheduled for mid-2007.




Figure 4: Screen capture of simple labeling demonstration

References

- [1] S. Angebrandt et. al. "Definition of the Porting Layer for the X v11 Sample Server." X Consortium, Inc, and X.org Foundation (2004).
- [2] G. Faden. "Solaris Trusted Extensions: Architectural Overview." Sun Microsystems white paper (2006). Available URL: <http://opensolaris.org/os/community/security/projects/tx/>.
- [3] J. Gettys et. al. "Xlib - C Language X Interface." The Open Group (1996).
- [4] D. Kilpatrick, W. Salamon, and C. Vance. "Securing the X Window System with SELinux." NAI Labs Report #03-006 (2003). Available URL: <http://www.nsa.gov/selinux/info/docs.cfm>.
- [5] P. Loscocco and S. Smalley, "Integrating Flexible Support for Security Policies into the Linux Operating System." In *Proc. 10th USENIX Conference (FREENIX Track)* 2001.
- [6] C. Paget. "Exploiting design flaws in the Win32 API for privilege escalation." (2002).
- [7] R. Scheifler. "X Window System Protocol." X Consortium, Inc. (2004).
- [8] S. Smalley. "Configuring the SELinux Policy." NAI Labs Report #02-007 (2002). Available URL: <http://www.nsa.gov/selinux/info/docs.cfm>.
- [9] S. Smalley, C. Vance, and W. Salamon. "Implementing SELinux as a Linux Security Module." NAI Labs Report #01-043 (2001). Available URL: <http://www.nsa.gov/selinux/info/docs.cfm>.
- [10] R. Spencer, S. Smalley, P. Loscocco, M. Hibler, D. Andersen, and J. Lepreau. "The Flask Security Architecture: System Support for Diverse Security Policies" In *Proc. 8th USENIX Conference (Security Symposium)* 1999.
- [11] E. Walsh. "X Access Control Extension Specification." X.org Foundation (2006).
- [12] D. Wiggins. "Security Extension Specification: Version 7.1." X Consortium, Inc. (1996).
- [13] D. Wiggins. "Security Extension Server Design (Draft Version 3.0)." X Consortium, Inc. (1996).
- [14] The X.org Foundation. Available URL: <http://www.x.org>.
- [15] GNOME: The Free Software Desktop Project. Available URL: <http://www.gnome.org>.
- [16] K Desktop Environment. Available URL: <http://www.kde.org>.
- [17] GTK+: The GIMP Toolkit. Available URL: <http://www.gtk.org>.
- [18] Qt: Trolltech. Available URL: <http://www.trolltech.com/products/qt>.