

Type Safety for TAL-1

Progress: If $\vdash M$ then there is some M' such that $M \rightarrow M'$.

Preservation: If $\vdash M$ and $M \rightarrow M'$ then either M' is `StackOverflow`, or $\vdash M'$.

The Java Security Manager

Allows or disallows various operations.

Various kinds of operations (reading or writing files, connecting to another machine) requires asking the security manager for permission.

Security managers are objects of the [SecurityManager](#) class.

```
public class BadClass {
    public static void main(String args[]) {
        try {
            Runtime.getRuntime().exec (" /bin/rm /path/to/filexyz");
        } catch (Exception e) {
            System.out.println ("Deletion command failed: " + e);
            return;
        }
        System.out.println ("Deletion command successful!");
    }
}
```

```
public class BadClass {
    public static void main(String args[]) {
        try {
            Runtime.getRuntime().exec (" /bin/rm /path/to/filexyz");
        } catch (Exception e) {
            System.out.println ("Deletion command failed: " + e);
            return;
        }
        System.out.println ("Deletion command successful!");
    }
}
```

Deletion command successful!

The local file gets deleted, if the user has permissions from the operating system.

What if such code is present in some applet loaded by a web-browser?

What if such code is present in some applet loaded by a web-browser?

```
import java.applet.Applet; import java.awt.Graphics;

public class BadApplet extends Applet{

    String text;

    public void init() {
        try { Runtime.getRuntime().exec("/bin/rm -rf /path/to/filexyz");
        } catch (Exception e) { text = "Deletion command failed: " + e; return; }
        text = "Deletion command successful!";
    }

    public void paint(Graphics g){ g.drawString(text, 15, 25); }
}
```

This applet is used in the following HTML page.

```
<html><body>  
<applet code="BadApplet.class" width=750 HEIGHT=50></applet>  
</body></html>
```

This applet is used in the following HTML page.

```
<html><body>  
<applet code="BadApplet.class" width=750 HEIGHT=50></applet>  
</body></html>
```

Loading this page in a web browser shows:

```
Deletion command failed: java.security.AccessControlException:  
access denied (java.io.FilePermission /bin/rm execute)
```


This applet is used in the following HTML page.

```
<html><body>  
<applet code="BadApplet.class" width=750 HEIGHT=50></applet>  
</body></html>
```

Loading this page in a web browser shows:

```
Deletion command failed: java.security.AccessControlException:  
access denied (java.io.FilePermission /bin/rm execute)
```

The web browser automatically give restricted permissions to applets.

The sandbox associated with a class depends upon the source from where it was loaded.

The typical sequence used for potentially dangerous operations:

- **User program** makes some request to the Java API.
- The **Java API** asks the security manager for permissions.
- If the **security manager** doesn't want to allow this operation, it throws back an **exception** which is thrown back to the user program.
- Otherwise the security manager does nothing and the Java API completes the operation.

In the previous example, the user program calls the **exec** method, which calls the **checkExec** method on the security manager to check for permission.

The code executed on calling `exec` is similar to this:

```
public process exec (String command) throws IOException {  
    ...  
    SecurityManager sm = System.getSecurityManager();  
    if (sm != null) {  
        sm.checkExec();  
        // security exception can be raised here  
    }  
    // remaining code follows  
    ...  
}
```

Another example: reading files.

```
// open a file
FileInputStream fis = new FileInputStream ("somefile");
// read a byte
int x = fis.read();
```

The code executed on calling [FileInputStream](#) is similar to

```
public FileInputStream (String name) throws FileNotFoundException {
    SecurityManager sm = System.getSecurityManager();
    if (sm != null) { sm.checkRead(name); }
    try { open (name);
    } catch (IOException e) {
        throw new FileNotFoundException (name);
    }
}
```

The `System` class has various useful data and functions which are global for the whole virtual machine.

The security manager is obtained by `getSecurityManager` method, and `null` is returned if no security manager has been set.

The security manager is set by `setSecurityManager` method, and an exception is raised if the security manager has already been set.

Hence once the security manager has been set, it cannot be modified.

In particular, java applications can set the security manager before executing remote applets, so that these applets don't try to set their own security manager.

Defining one's own security manager: we extend the `SecurityManager` class and override the functions as required.

```
public class NewSecurityManager extends SecurityManager {  
    public void checkExec (String cmd) {  
        // always disallow exec  
        throw new SecurityException ("exec not allowed")  
    }  
}
```

Modifying the `BadClass` to use this security manager.

```
public class NewBadClass {
    public static void main(String args[]) {
        SecurityManager sm = new NewSecurityManager();
        System.setSecurityManager(sm);
        try {
            Runtime.getRuntime().exec ("/bin/rm /path/to/filexyz");
        } catch (Exception e) {
            System.out.println ("Deletion command failed: " + e);
            return;
        }
        System.out.println ("Deletion command successful!");
    }
}
```

Modifying the `BadClass` to use this security manager.

```
public class NewBadClass {
    public static void main(String args[]) {
        SecurityManager sm = new NewSecurityManager();
        System.setSecurityManager(sm);
        try {
            Runtime.getRuntime().exec ("/bin/rm /path/to/filexyz");
        } catch (Exception e) {
            System.out.println ("Deletion command failed: " + e);
            return;
        }
        System.out.println ("Deletion command successful!");
    }
}
```

Deletion command failed: java.lang.SecurityException: exec not allowed

Examples of methods of the security manager.

- `checkRead (String file)`: called e.g. by `FileInputStream (String file)`.
- `checkWrite (String file)`: called by `FileOutputStream (String file)`.
- `checkDelete (String file)`

Examples of methods of the security manager.

- `checkRead (String file)`: called e.g. by `FileInputStream (String file)`.
- `checkWrite (String file)`: called by `FileOutputStream (String file)`.
- `checkDelete (String file)`

Note that while creating a `FileInputStream` object requires a `checkRead` call, the actual `read()` operations on the file input stream requires no permission.

- A trusted class can choose to deliver the `FileInputStream` object to an untrusted class which can then read from the file.
- It is efficient to check permissions only once.

The Access Controller

- Has functions similar to the security manager.
- Provides easy enforcement of fine grained security policies.
- The security manager works most of the time by calling the access controller.
- Implemented by the `AccessController` class, accessed through its static methods.

Involves the following four classes.

- The **CodeSource** class: represents the source from which a certain class was loaded, an an optional list of certificates which was used to sign that code.
- The **Permission** and **Permissions** classes: represent various kinds of permissions.
- The **Policy** class: a policy maps code source objects to permission objects. Only one policy can be associated with the JVM at any point of time, like the security manager. But the policy can be modified.
- The **ProtectionDomain** class: a protection domain represents all the permissions granted to a particular code source.

A permission has three properties:

- **A type**: what kind of permission is this?
- **A name**: the object that this permission talks about.
- **Actions**

A permission has three properties:

- **A type**: what kind of permission is this?
- **A name**: the object that this permission talks about.
- **Actions**

Permission objects for accessing files are members of the **FilePermission** class (subclass of the **Permission** class).

- The type is **FilePermission**.
- The name is the name of the file.
- Possible actions are "read", "write", "delete" and "execute".

A permission has three properties:

- **A type**: what kind of permission is this?
- **A name**: the object that this permission talks about.
- **Actions**

Permission objects for accessing files are members of the **FilePermission** class (subclass of the **Permission** class).

- The type is **FilePermission**.
- The name is the name of the file.
- Possible actions are "read", "write", "delete" and "execute".

Permission objects are used for requesting permissions as well as for representing granted permissions.

The security manager, on receiving the `checkExec("/bin/rm")` call, would normally construct the following permission object

```
FilePermission fp = new FilePermission ("/bin/rm", "execute");
```

and then query the access controller.

```
AccessController.checkPermission (fp);
```


The security manager, on receiving the `checkExec("/bin/rm")` call, would normally construct the following permission object

```
FilePermission fp = new FilePermission ("/bin/rm", "execute");
```

and then query the access controller.

```
AccessController.checkPermission (fp);
```

Other examples:

```
FilePermission fp1 = new FilePermission ("/bin/*", "execute");
```

```
FilePermission fp2 = new FilePermission ("/home/userx", "read, write");
```

```
SocketPermission sp1 = new SocketPermission ("hostname:port", "connect");
```

```
SocketPermission sp1 = new SocketPermission ("hostname:port", "accept, listen");
```

Policies are specified by objects of `Policy` class.

It can be obtained and set using `getPolicy ()` and `setPolicy (Policy p)`.

Policy objects can be created by reading from a file which lists the policy rules.

Typically done at startup time:

```
java -Djava.security.manager -Djava.security.policy=<policyfilename> <class> <args>  
appletviewer -J-Djava.security.policy=<policyfilename> file.html
```

The policy file have rules mapping code sources to sets of permissions.

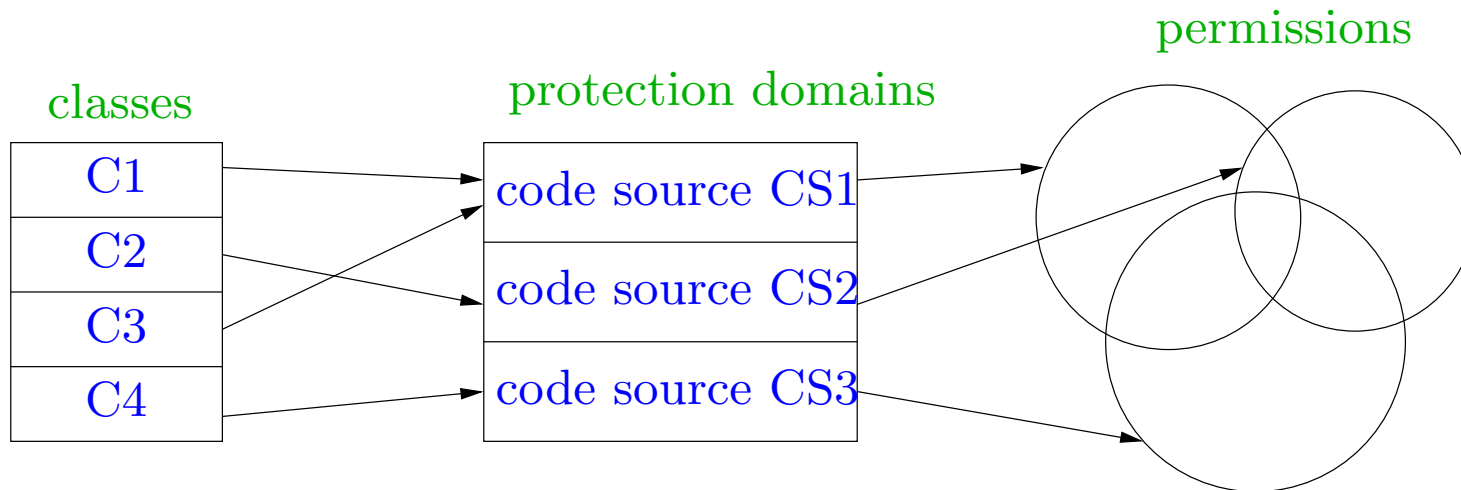
```
grant codeBase "file:/home/userxyz/classes" {  
    permission java.io.FilePermission "/bin/rm" "execute";  
    permission java.net.SocketPermission "localhost:1024-" "listen, accept";  
};
```

```
grant signedBy <signer>, codeBase "http://www.xyz.com" {  
    permission ...  
    ...  
};
```

A **protection domain** groups a **code source** with a set of permissions.

The **class loader** is supposed to associate a protection domain with a class when it loads the class.

The protection domain associated with each class is used by the access controller when it is called to check a permission using the **checkPermission()** method.



Stack inspection

Allowing or disallowing a permission depends on the context in which the `checkPermission` method was called.

The access controller needs to examine the protection domains associated with all the classes on the stack.

The permission is granted only if all the protection domains on the stack have this permission.

In our old example, the `BadClass.main()` method for deleting a file calls the `Runtime.exec()` method which calls the `AccessController.checkPermission()` to check execute permission on `/bin/rm`.

Further, the `BadClass.main()` method itself may be called by some other method `m()` of class `C`.

We get the following stack.

<code>AccessController.checkPermission()</code>
<code>Runtime.exec()</code>
<code>BadClass.main()</code>
<code>C.m()</code>
...

The execute permission should be granted only if all the classes on the stack have that permission in their protection domain.

Hence the access controller checks that all frames from the top of the stack to the bottom have this permission in the protection domains of the respective classes.

Sometimes a trusted class may choose to give its permissions to lower frames on the stack.

E.g. an untrusted applet may call some routine to draw something on the screen, and the routine requires some local font file.

This is done using the `doPrivileged()` method.

```
untrustedclass { f() { ... trustedclass.draw() ...}}
trustedclass {
    public void draw {
        ...
        AccessController.doPrivileged (new PrivilegedAction () {
            public Object run () {
                // privileged code here
                ... <read font file> ...
            } }); }}

```

Instead of the `doPrivileged()` method

```
AccessController.doPrivileged (new PrivilegedAction () {  
    public Object run () {  
        <privileged code>  
    }  
});
```

earlier versions used `beginPrivileged()` and `endPrivileged()` calls.

```
AccessController.beginPrivileged();  
<privileged code>  
AccessController.endPrivileged();
```


To understand the **stack inspection** algorithm let us assume the following operations.

- `enablePrivilege(T)`
- `disablePrivilege(T)`
- `checkPrivilege(T)`
- `revertPrivilege(T)`

where T is a **target** (permission in the Java terminology) we wish to protect.

Actions taken by these operations:

- `enablePrivilege(T)` puts an `enabledPrivilege(T)` flag on the current stack frame if the current class has access to T according to the policy.
- `disablePrivilege(T)` puts a `disabledPrivilege(T)` flag on the current stack frame (and removes `enabledPrivilege(T)` flag if present).
- `revertPrivilege(T)` removes `enabledPrivilege(T)` and `disabledPrivilege(T)` flags from the current stack frame if present.
- `checkPrivilege(T)` examines the stack as follows ...

```
checkPrivilege (T) {  
    for SF from top stack frame to bottom stack frame {  
        if (policy doesn't allow the class in SF to access T) throw ForbiddenException;  
        if (SF has enabledPrivilege (T) flag) return;  
        if (SF has disabledPrivilege (T) flag) throw ForbiddenException;  
    }  
    return; // reached bottom of stack  
}
```

The ABLP Logic

Abadi, Burrows, Lampson and Plotkin, 1993

We will model stack inspection using the (subset of) ABLP logic described below. The language contains

- **Principals**, modeling persons, organizations as well as cryptographic keys.
- **Targets**, modeling resources we wish to protect.
- **Statements**, modeling utterances of principals.

The ABLP Logic

Abadi, Burrows, Lampson and Plotkin, 1993

We will model stack inspection using the (subset of) ABLP logic described below. The language contains

- **Principals**, modeling persons, organizations as well as cryptographic keys.
- **Targets**, modeling resources we wish to protect.
- **Statements**, modeling utterances of principals.
 - The statement $P \text{ says Ok}(T)$ means that the principal P is authorizing access to target T .

The ABLP Logic

Abadi, Burrows, Lampson and Plotkin, 1993

We will model stack inspection using the (subset of) ABLP logic described below. The language contains

- **Principals**, modeling persons, organizations as well as cryptographic keys.
- **Targets**, modeling resources we wish to protect.
- **Statements**, modeling utterances of principals.
 - The statement $P \text{ says Ok}(T)$ means that the principal P is authorizing access to target T .
 - $P \mid Q \text{ says } s$ means $P \text{ says } (Q \text{ says } s)$, i.e. P quotes Q as saying s .

The ABLP Logic

Abadi, Burrows, Lampson and Plotkin, 1993

We will model stack inspection using the (subset of) ABLP logic described below. The language contains

- **Principals**, modeling persons, organizations as well as cryptographic keys.
- **Targets**, modeling resources we wish to protect.
- **Statements**, modeling utterances of principals.
 - The statement $P \text{ says Ok}(T)$ means that the principal P is authorizing access to target T .
 - $P \mid Q \text{ says } s$ means $P \text{ says } (Q \text{ says } s)$, i.e. P quotes Q as saying s .
 - $P \wedge Q \text{ says } s$ means that both P and Q say s .

The ABLP Logic

Abadi, Burrows, Lampson and Plotkin, 1993

We will model stack inspection using the (subset of) ABLP logic described below. The language contains

- **Principals**, modeling persons, organizations as well as cryptographic keys.
- **Targets**, modeling resources we wish to protect.
- **Statements**, modeling utterances of principals.
 - The statement $P \text{ says Ok}(T)$ means that the principal P is authorizing access to target T .
 - $P \mid Q \text{ says } s$ means $P \text{ says } (Q \text{ says } s)$, i.e. P quotes Q as saying s .
 - $P \wedge Q \text{ says } s$ means that both P and Q say s .
 - $P \Rightarrow Q$ means that P speaks for Q , i.e. P has at least as much authority as Q .

We assume a set of atomic statements and atomic principals.

principal $P ::=$

AtomicPrincipal

$P_1 \wedge P_2$

$P_1 \mid P_2$

statement $s ::=$

AtomicStatement

$s_1 \wedge s_2$

$s_1 \rightarrow s_2$

P says s_1

$P_1 \Rightarrow P_2$

Example Given some s we define following new statements.

$$s_1 \equiv (Alice \wedge Bob) \text{ says } (Charlie \Rightarrow (Alice \wedge Bob))$$

Alice and *Bob* declare *Charlie* to be their representative.

Example Given some s we define following new statements.

$$s_1 \equiv (Alice \wedge Bob) \text{ says } (Charlie \Rightarrow (Alice \wedge Bob))$$

Alice and *Bob* declare *Charlie* to be their representative.

$$s_2 \equiv Charlie \mid Alice \text{ says } s$$

Charlie quotes *Alice* as saying s .

Example Given some s we define following new statements.

$$s_1 \equiv (Alice \wedge Bob) \text{ says } (Charlie \Rightarrow (Alice \wedge Bob))$$

Alice and *Bob* declare *Charlie* to be their representative.

$$s_2 \equiv Charlie \mid Alice \text{ says } s$$

Charlie quotes *Alice* as saying s .

$$s_3 \equiv (Alice \text{ says } s) \rightarrow s$$

If *Alice* says s then it must be true.

Example Given some s we define following new statements.

$$s_1 \equiv (Alice \wedge Bob) \text{ says } (Charlie \Rightarrow (Alice \wedge Bob))$$

Alice and *Bob* declare *Charlie* to be their representative.

$$s_2 \equiv Charlie \mid Alice \text{ says } s$$

Charlie quotes *Alice* as saying s .

$$s_3 \equiv (Alice \text{ says } s) \rightarrow s$$

If *Alice* says s then it must be true.

Intuitively, from $s_1 \wedge s_2 \wedge s_3$ we should be able to prove s .

Example Given some s we define following new statements.

$$s_1 \equiv (Alice \wedge Bob) \text{ says } (Charlie \Rightarrow (Alice \wedge Bob))$$

Alice and *Bob* declare *Charlie* to be their representative.

$$s_2 \equiv Charlie \mid Alice \text{ says } s$$

Charlie quotes *Alice* as saying s .

$$s_3 \equiv (Alice \text{ says } s) \rightarrow s$$

If *Alice* says s then it must be true.

Intuitively, from $s_1 \wedge s_2 \wedge s_3$ we should be able to prove s .

For this we require certain rules (axioms) for making proofs.

Axioms about statements

- 1 *If s is an instance of a theorem of propositional logic then s is true in ABLP logic.*

Axioms about statements

- 1 *If s is an instance of a theorem of propositional logic then s is true in ABLP logic.*

E.g. the ABLP statement

$$(P \text{ says } s) \rightarrow (P \text{ says } s)$$

is an instance of the propositional logic statement

$$X \rightarrow X$$

Axioms about statements

- 1 *If s is an instance of a theorem of propositional logic then s is true in ABLP logic.*

E.g. the ABLP statement

$$(P \text{ says } s) \rightarrow (P \text{ says } s)$$

is an instance of the propositional logic statement

$$X \rightarrow X$$

The ABLP statement

$$(P \text{ says } s) \wedge ((P \text{ says } s) \rightarrow s) \rightarrow s$$

is an instance of the propositional logic statement

$$(X \wedge (X \rightarrow Y)) \rightarrow Y$$

Axioms about statements

- 1 *If s is an instance of a theorem of propositional logic then s is true in ABLP logic.*

E.g. the ABLP statement

$$(P \text{ says } s) \rightarrow (P \text{ says } s)$$

is an instance of the propositional logic statement

$$X \rightarrow X$$

The ABLP statement

$$(P \text{ says } s) \wedge ((P \text{ says } s) \rightarrow s) \rightarrow s$$

is an instance of the propositional logic statement

$$(X \wedge (X \rightarrow Y)) \rightarrow Y$$

Hence both ABLP statements are true.

2 If s and $s \rightarrow s'$ then s' .

2 If s and $s \rightarrow s'$ then s' .

3 $(P \text{ says } s \wedge P \text{ says } (s \rightarrow s')) \rightarrow P \text{ says } s'$

We can draw conclusions from statements made by principals.

2 *If s and $s \rightarrow s'$ then s' .*

3 *$(P \text{ says } s \wedge P \text{ says } (s \rightarrow s')) \rightarrow P \text{ says } s'$*

We can draw conclusions from statements made by principals.

4 *If s then $P \text{ says } s$ for every principal P .*

True ABLP statements are supported by all principals.