


# Breeze: A Language For Writing Secure Software

Cătălin Hrițcu



**Penn**  
UNIVERSITY of PENNSYLVANIA





SAFE



Northeastern

**BAE SYSTEMS**

# The SAFE team



**May 2011 @ San Jose:** Sumit Ray, Howard Reubenstein, Andrew Sutherland, Tom Knight, Olin Shivers, Benjamin Pierce, Ben Karel, Benoit Montagu, Jonathan Smith, Catalin Hritcu, Randy Pollack, André DeHon, Gregory Malecha, Basil Krikeles, Greg Sullivan, Greg Frazier, Tim Anderson, Bryan Loyall.

**Missing or added since May 2011:** Greg Morrisett, Peter Trei, David Wittenberg, Amanda Strnad, Justin Slepak, David Darais, Robin Morisset, Chris White, Anna Gommerstadt, Marty Fahey, Tom Hawkins, Karl Fischer, Hillary Holloway, Andrew Kaluzniacki.



# Sony Makes it Official: PlayStation Network Hacked

By Thomas, PCWorld | April 23, 2011 7:35 AM

PlayStation Network was taken offline three days ago, all eyes fell on the Anonymous group. Sony has a dislike to Sony over its treatment of hardware hacker George Hotz. The group has taken down PlayStation 3 consoles and boasts 70 million users, so this is no joke. Last night Sony confessed that an "external intrusion" had taken down the PlayStation Network. Sony thought the smooth and seamless transition was "going forward". However, the scale as to why this is a fraudulent and substantial infrastructure breach raises questions.

SEPTEMBER 07, 2011

# Debacle deepens for hacked SSL certificates issuer

Report indicates widespread compromise of DigiNotar's infrastructure as questions industry's response to breaches

By Robert Lemos | InfoWorld

The [attack on DigiNotar](#), an issuer of certificates, appears more serious than thought. A report by the breach raises questions about the fraudulent and substantial infrastructure breach.

Article Comments

Tweet

Email Order Reprints

Text Size +

# HTC admits to 'serious' security flaw on smartphones

October 5, 2011

Smartphone maker HTC has admitted that its smartphones are vulnerable to a security flaw that could allow attackers to find out where you are, intercept text messages and disable apps.

On Tuesday Fairfax Media reported that **Android smartphones** that use the HTC Sense app that is granted access to location data require. Once granted access to location data that was shown

## The Stuxnet outbreak

### A worm in the centrifuge

An unusually sophisticated cyber-weapon is mysterious but important

Sep 30th 2010 | from the print edition



...y, the Predator and Reaper drones, may have been used to track which alleged al-Qaida terrorist Anwar al-Awlaki.

U.S. Air Force declined comment on the report, which was first reported by Wired.

...keep wiping it off and it keeps coming back," an official said to the publication.


...the Pentagon would unlikely confirm its has been hit by yet another cyber-attack, the virus could trigger one of the most deadly and popular malware attacks, which have proved especially effective in Iran, Pakistan, as well as other venues like Yemen.

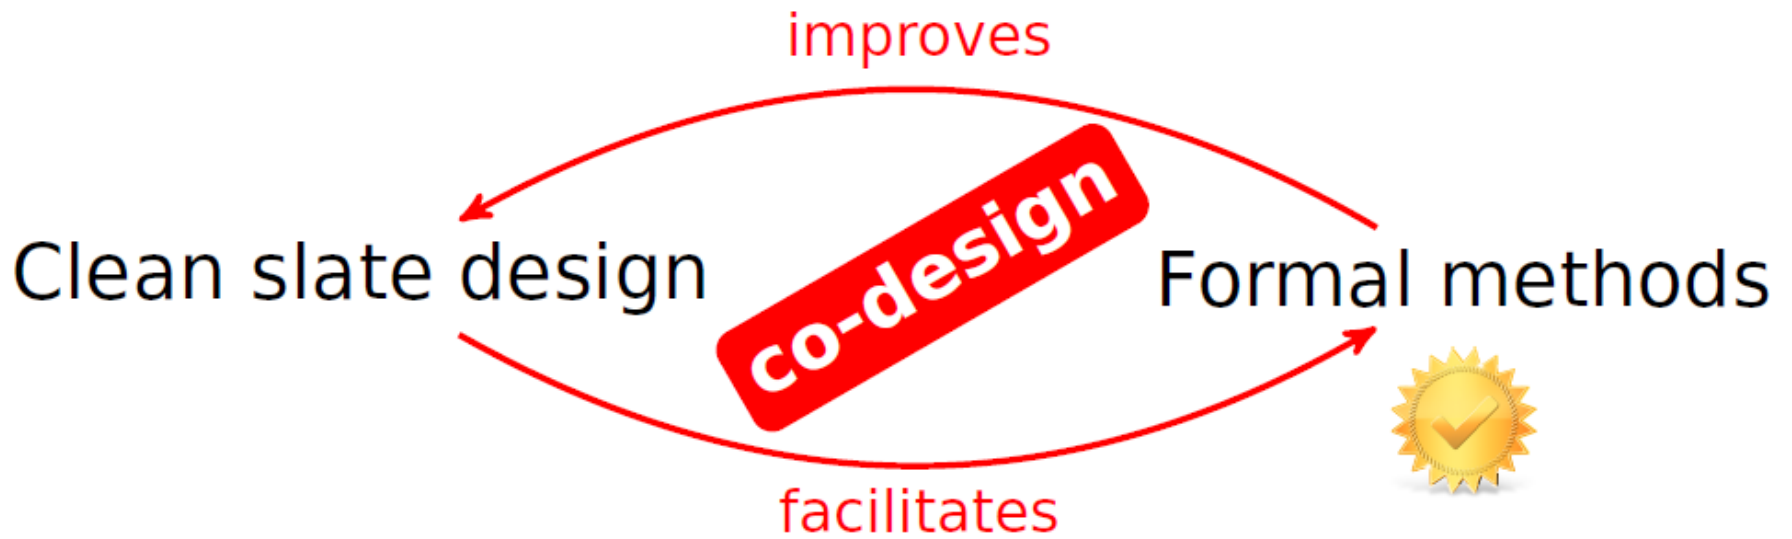
# Common Weaknesses Enumeration: Top 25

Rank	Score	ID	Name
[1]	93.8	<a href="#">CWE-89</a>	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection')
[2]	83.3	<a href="#">CWE-78</a>	Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection')
[3]	79.0	<a href="#">CWE-120</a>	Buffer Copy without Checking Size of Input ('Classic Buffer Overflow')
[4]	77.7	<a href="#">CWE-79</a>	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting')
[5]	76.9	<a href="#">CWE-306</a>	Missing Authentication for Critical Function
[6]	76.8	<a href="#">CWE-862</a>	Missing Authorization
[7]	75.0	<a href="#">CWE-798</a>	Use of Hard-coded Credentials
[8]	75.0	<a href="#">CWE-311</a>	Missing Encryption of Sensitive Data
[9]	74.0	<a href="#">CWE-434</a>	Unrestricted Upload of File with Dangerous Type
[10]	73.8	<a href="#">CWE-807</a>	Reliance on Untrusted Inputs in a Security Decision
[11]	73.1	<a href="#">CWE-250</a>	Execution with Unnecessary Privileges
[12]	70.1	<a href="#">CWE-352</a>	Cross-Site Request Forgery (CSRF)
[13]	69.3	<a href="#">CWE-22</a>	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal')
[14]	68.5	<a href="#">CWE-494</a>	Download of Code Without Integrity Check
[15]	67.8	<a href="#">CWE-863</a>	Incorrect Authorization
[16]	66.0	<a href="#">CWE-829</a>	Inclusion of Functionality from Untrusted Control Sphere
[17]	65.5	<a href="#">CWE-732</a>	Incorrect Permission Assignment for Critical Resource
[18]	64.6	<a href="#">CWE-676</a>	Use of Potentially Dangerous Function
[19]	64.1	<a href="#">CWE-327</a>	Use of a Broken or Risky Cryptographic Algorithm
[20]	62.4	<a href="#">CWE-131</a>	Incorrect Calculation of Buffer Size
[21]	61.5	<a href="#">CWE-307</a>	Improper Restriction of Excessive Authentication Attempts
[22]	61.1	<a href="#">CWE-601</a>	URL Redirection to Untrusted Site ('Open Redirect')
[23]	61.0	<a href="#">CWE-134</a>	Uncontrolled Format String
[24]	60.3	<a href="#">CWE-190</a>	Integer Overflow or Wraparound
[25]	59.9	<a href="#">CWE-759</a>	Use of a One-Way Hash without a Salt

<http://cwe.mitre.org/top25/index.html#Listing>

# The SAFE project





Occasion to try new abstractions

Simpler, cleaner design


Fully correct wrt. specs

Precise specifications



Global guarantees

Machine checked proofs


***Userware***  
(written in Breeze)



***Concreteware***  
(written in Tempest;  
formally verified)



***Hardware***  
(written in BlueSpec)



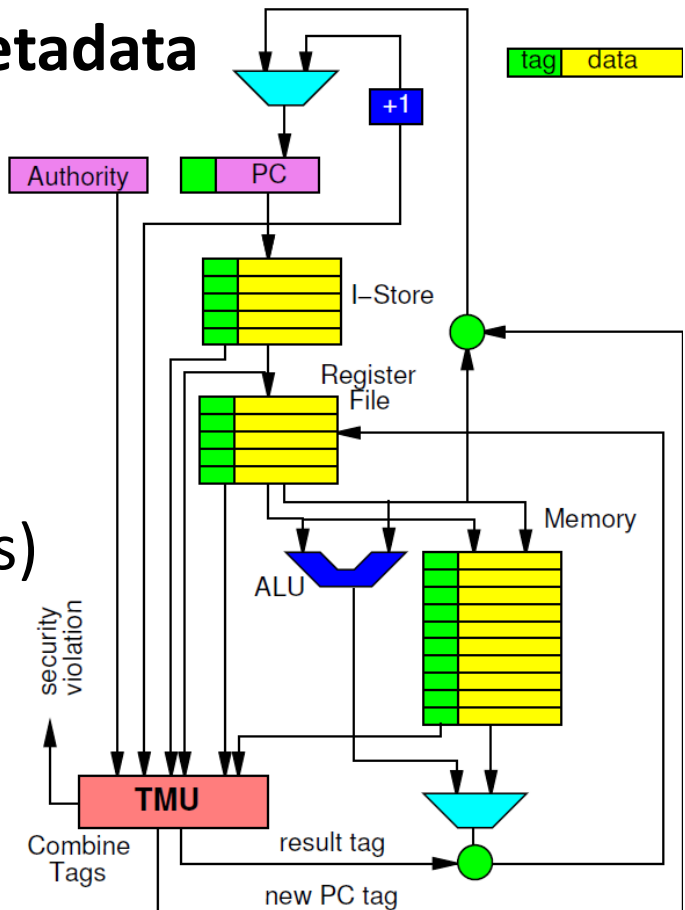


# Breeze

- High-level PL for writing **secure** software
  - user programs & system services (+ security policies)
- **functional core** ( $\lambda$ ) + concurrency ( $\pi$ ) + state (!) ...
- **dynamically typed** (for now)
  - easier to experiment with
  - directly reflects capabilities of HW
  - dynamically-checked first-class **contracts**
- **fine-grained dynamic information flow control**
- **access control** (more later on why we need both)

# Fine-grained dynamic IFC


- Every Breeze value has an IFC label:  $\langle \rho, \lambda x.x \rangle @ l$
- At the SAFE HW level, **every word is tagged**
  - tags are pointers to **arbitrary metadata**
  - **checking happens in parallel**  
(HW rule cache speeds this up)
  - TMU is **programmable**  
(rules defined in software)
  - **composable** rule models  
(IFC labels = one of many models)
  - primitive TMU rule models:
    - can't add two pointers,
    - can't execute an integer, etc.
    - only scheduler can create threads



# SAFE hardware prototype

FPGA

DRAM (2GB)



PCI2e 2 GB/s each direction

# IFC Labels in Breeze

- Started with Decentralized Label Model (DLM) by Liskov & Meyers ( $A \rightarrow [B, C], D \rightarrow [C, D]$ )
- Later realized **DLM is not the only option**
  - Asbestos, HiStar/DStar, FLUME, HAILS, ...  
all have their own decentralized ... label models
- Breeze has “generalized label model”
  - parameterized by join, meet, principals, authority, etc.
  - hope to also support HAILS DC labels in the future
- How labels look like = the boring part of the story
  - in the **dynamic IFC** setting:  
**the interesting part is how labels behave at runtime**



# Dynamic IFC: is it even possible?

- Non-interference is not a property (2-hyperproperty)
- Until around 2007 **folklore in PL community:**  
**non-interference can't be enforced dynamically**  
because of implicit flows
  - at least not without multi-execution or static information about the branches that are not executed
- Dynamic techniques pioneered IFC [Fenton, CJ '74]
- The lack of non-interference proofs didn't stop the OS community from building dynamic IFC systems (Asbestos, HiStar/DStar, FLUME ...)

# Purely dynamic IFC: Yes We Can!

- “From dynamic to static and back: Riding the roller coaster of IFC research” [Sabelfeld & Russo, PSI 2009]
- purely dynamic analysis (monitor) for termination-insensitive non-interference
  - termination-sensitivity hard to achieve in any way
- “flow insensitive” analysis
  - labels of mutable variables can’t change
  - precise flow sensitive dynamic analysis proved impossible [Sabelfeld&Russo, CSF 2010]
  - flow insensitivity is not such a big deal for a new language
    - ML references / Java variables + fields only have weak updates
    - Breeze channels have label for contents fixed at creation time

# Purely-dynamic IFC: functional setting

- “Efficient purely-dynamic information flow analysis”  
[Austin & Flanagan, PLAS 2009]

$$\frac{}{\rho, pc \vdash \lambda x.e \Downarrow \langle \rho, \lambda x.e \rangle @ pc} \quad \frac{\rho(x) = v @ l}{\rho, pc \vdash x \Downarrow v @ (l \vee pc)}$$

$$\frac{\begin{array}{l} \rho, pc \vdash e_1 \Downarrow \langle \rho', \lambda x.e \rangle @ l_1 \\ \rho, pc \vdash e_2 \Downarrow v_2 @ l_2 \\ \rho'[x \mapsto v_2 @ l_2], pc \vee l_1 \vdash e \Downarrow v_3 @ l_3 \end{array}}{\rho, pc \vdash e_1 e_2 \Downarrow v_3 @ l_3}$$

- Termination-insensitive non-interference  $\left. \begin{array}{l} \rho_1, pc \vdash e \Downarrow v_1 @ l_1 \\ \rho_2, pc \vdash e \Downarrow v_2 @ l_2 \\ \rho_1 \simeq_l \rho_2 \end{array} \right\} \Rightarrow v_1 @ l_1 \simeq_l v_2 @ l_2$

# The “infectious pc” problem

- pc automatically “tracks” branch labels
  - in order to prevent implicit flows
  - `copy = (if secret@H then true@L else false@L)`
- pc “infects” all values created on high branch
- this leads to **deeply high-labeled values**:
  - `if true@H then [1,2] else [] ==>`  
`(cons (1@H) (cons (2@H) (nil@H)))@H@H`
- we needed “infectious pc” because of **automatic pc lowering** on control flow merge points



# First attempt to “manual pc lowering”

- Using pc declassification, really bad idea (adds tons of completely spurious audit points)



Catalin Hritcu <catalin.hritcu@gmail.com>

---

## [Safe-breeze] Manual PC declassification considered harmful

6 messages

---

Benjamin C. Pierce <bcpierce@cis.upenn.edu>

Wed, Aug 17, 2011 at 1:24 PM

To: safe-breeze@lists.crash-safe.org

A few months ago, we made the decision that it was better to remove the "automatic declassification of the PC" at the ends of conditionals and functions in Breeze (and at return instructions in the ISA) and, instead, demand that programmers lower the PC manually, if it becomes higher than they want it. Over the past few days, we've finally made this change to Breeze and have been experimenting with programming in this style. Our conclusion, sadly, is that it doesn't work.

# The “poison pill” problem

- Fine-grained, dynamic IFC with decentralized LM
- Any code can classify data
  - `(P,_,_) = newPrin “P”; pill = 42@(P -> [P])`
- High data can be hidden under low labels
  - `x = [1,2,pill]@L; send cpub x`
- IFC violations are dynamic errors
  - `recv cpub x; map ((+) 1) x`
- Labels themselves are an IF channel
  - `x = (if secret@H then ()@H else ()@top);  
copy = (labelOf x == H)`

# Non solution(s)

- Labels are an IF channel, so **hide labels**

$$\frac{\rho, pc \vdash e \Downarrow v@l}{\rho, pc \vdash \text{labelOf } e \Downarrow l@l}$$

- Threads get killed on IFC errors
  - critical components (scheduler, allocator, drivers, ...) get killed on reading poison pill
- Make channels unreliable?
  - contract on channel, silently discards poison pills
  - restore reliability using threads and timeouts?
  - so we can still write labelOf using a timing channel

# Brackets: killing two birds with one stone

- Brackets = construct for **manual pc lowering / restoring**
  - automatic pc lowering played big role in creating our 2 problems
- Programmer has to **predict pc + label on all c.f. branches**

$$\frac{\rho \vdash e, pc \Downarrow v@l, pc' \quad l \sqsubseteq lb \quad pc' \sqsubseteq lb}{\rho \vdash lb \langle e \rangle, pc \Downarrow v@lb, pc}$$

- Brackets are **not a declassification construct**
  - always safe, unlike manual pc declassification
- Final label cannot depend on secrets
  - `top<if secret@H then ()@H else ()@top> ==> ()@top`
- Labels are now public, **no more “poison pills”**
- **No more “infectious pc”**
  - `H<if true@H then [1,2] else []> ==> [1,2]@H`



# Attacker model


- Why does Breeze also have access control?
  - Isn't IFC enough?
- **IFC doesn't protect against malicious code!**
  - there will always be covert channels
  - malicious code can extremely easily exfiltrate secrets
  - I can write exfiltrator in 5 minutes, without timer
- Big part of IFC community seems to ignore the problem and accept **bogus attacker model**
  - We won't!

# Authorities in Breeze

- Unprivileged Breeze code **cannot**
    - read and compute with secrets  
(and thus leak them over covert channels)
    - declassify secrets / endorse tainted values
  - Authorities are first-class values = **capabilities**
    - $(P, P_e, P_d) = \text{newPrin } "P"; \text{ setAuth } P_e \text{ in } \dots$
  - Ambient authority makes this workable
    - otherwise pass authority on each little ISA instruction?
    - authority can also be attenuated (least privilege)
  - Access control decision based on:
    - ambient read authority + IFC label of the data
- } separate capabilities

# Ambient authority propagation strategies

- Q: Do closures capture creation-time authority? (S = lexical propagation)
- Q: Is authority passed from caller to callee? (D = dynamic propagation)
- Capable experiment by G. Mallecha & G. Morrisett
- Breeze-du-jour does “S”
- What’s the relation with SBAC, HBAC, IBAC?



Q: What do “least privilege”, “well-compartmentalized” mean?

# Status of Breeze

- Interpreter with lots of different flags & externals
  - standard library in different variants
  - programming experiments; exploring design space
  - compiler to SAFE architecture planned for the future
- Various subsets formalized in Coq & Isabelle
  - proved termination-insensitive non-interference
  - big-step semantics (concurrency not formalized yet)
- Ongoing effort on releasing a 1<sup>st</sup> stable version