

888 888888b. 888 .d8888b. d8b 888
888 "88b 888 d88P"88b 888 888 "88b d88P"88b d8P Y8b 888 d88P Y88b Y8P 888
888 .88P 888 Y88b. 888
888 88888b. 88888b. .d88b. 888 888 8888b. .d88b. .d88b. 8888888K. 8888b. .d8888b .d88b. .d88888
888 "88b 888 "88b d88P"88b 888 888 "88b d88P"88b d8P Y8b 888 "Y88b. .d8P Y8b d88P" 888 888 888888 888 888
888 .d888888 888 888 888 888 888 888 .d888888 888 888 88888888 888888 888 888 .d888888 "Y8888b. 88888888 888 888
888 888 888 888 888 Y88b 888 Y88b 888 888 888 Y88b 888 Y8b. 888 d88P 888 888 X88 Y8b. Y88b 888 Y88b d88P Y8b. Y88b. Y88b 888 888 888 Y88b. Y88b 888
88888888 "Y888888 888 888 "Y88888 "Y88888 "Y888888 "Y88888 "Y8888 8888888P" "Y888888 88888P" "Y8888 "Y88888
888 888
Y8b d88P Y8b d88P
"Y88P" "Y88P" 888 888
888 888888 888 888
888 Y88b. Y88b 888
"Y88888 "Y88888
888 Y8b d88P
"Y88P"

Capabilities: principles and object capabilities

aslan@cs.au.dk

acknowledgements: Lau Skorstengaard

Capabilities

- We discuss the topic of capabilities in two classes
 - Today: capability principles and object capabilities
 - Next week: capability machines

Secure design principles [Saltzer & Schroeder' 1975]

Economy of mechanism simple design makes it feasible to evaluate all possible paths

Fail-safe defaults conservative defaults make security errors observable

Complete mediation check every access to every object; skeptically examine performance gains that sacrifice security

Open design decoupled design

Every program and every user of the system should operate using the least set of privileges necessary to complete the job

Separation of privilege if feasible, avoid single point of compromise

Least privilege limit the damage from error (cf. "need-to-know")

Least common mechanism every shared mechanism presents a new path to leak information

Psychological acceptability design for human use

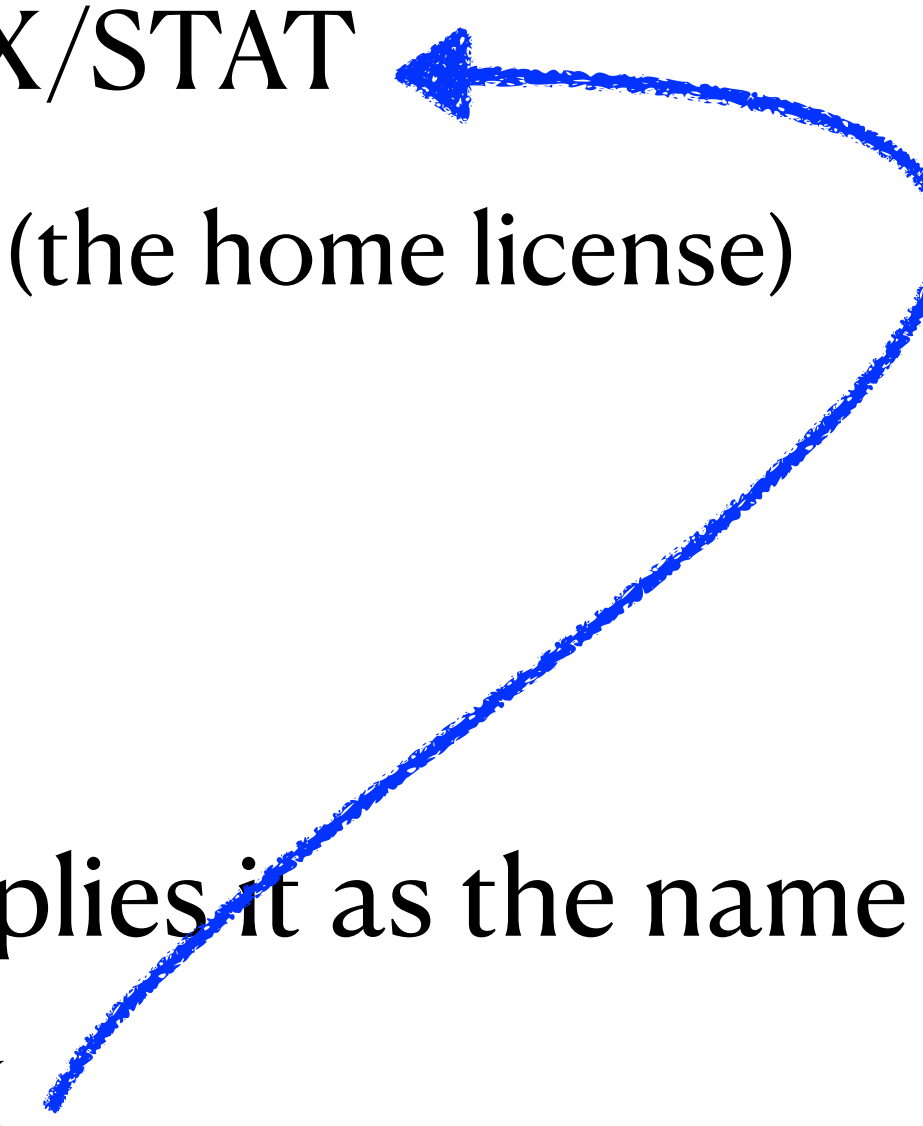
Access control

a form of authorization

- A mechanism for enforcing confidentiality and integrity
- Predefined security-sensitive *operations*
 - e.g., Read/Write/Execute
- Reference monitor (e.g., OS) consulted whenever predefined operations are invoked
- *Access control policy* specifies which operations a principal (subject) can do on an object

Confused deputy

[Hardy, 1988]

- Compiler installed in directory `/SYSX/`
 - It has a flag for outputting debug information into a user-specified file
 - Compiler also writes usage statistics info into `/SYSX/STAT`
 - The compiler therefore has write permission to `/SYSX` (the home license)
 - There is also a billing file `/SYSX/BILL`
 - Integrity attack
 - User learns about the name of `/SYSX/BILL` and supplies it as the name of debugging info
 - The attack works because compiler can write to `/SYSX`
- 

Who is to blame in this example?

- The compiler should have scanned the specified output destination to make sure that the debug information was not written to a different directory?
 - **No**, it is useful to allow the user to write the debug info in a different directory.
- The compiler should have checked the the directory was not SYSX.
 - **Again no**, the SYSX folder may not have been invented when the compiler was written, and a user may have a legitimate reason to use SYSX.
- The compiler should have checked that the specified file for output was not SYSX/BILL!
 - **No**, that may not be the only sensitive file in SYSX, and it is not the compiler's job to enforce access control.

“The fundamental problem is that the compiler runs with authority stemming from two sources (that’s why the compiler is a confused deputy)”



the user

the home
license

Hardy, 1988

Permission and Authority

[Miller, 2006]

- *Permissions* specify direct access rights
- *Authority* is the ability to cause effects
 - Directly by permission, or
 - Indirectly by permitted interactions with other programs
- In confused deputy, the user has authority to write to SYSX/BILL despite having no permission
- Gap between permissions and authority \Rightarrow leakage of access rights
 - Can we make it so that “reasoning about permissions” ~ “reasoning about authority”?

Capability model

[Dennis and van Horn, 1965]

- Traditional *access control policies* separate designation (knowing a name) from permission
 - (a user may have no permission to SYSX/BILL but they know the file name)
- Main idea in capability model:
 - Combine designation and access right into *a capability*
 - Possession of a capability grants the access
 - It is therefore crucial that capabilities are *unforgeable*
 - Prevents creating authority out of thin air
 - Capabilities may be shared between components
 - They can also be *attenuated* to comply with the principle of least privilege
- Simplifies reasoning about authority
 - No capabilities means no access

Capability is an unforgeable token of authority

The confused deputy example and capabilities

- In a capability system
 - Compiler would have capability for the file containing statistics
 - User supplies capability instead of location
- Confused deputy problem is no problem with capabilities :-)

Capabilities in operating systems

- KeyKOS [1992], EROS [1999], Coyotos [2004], PSOS[2003]
 - Pure-capability systems that use both OS capability and hardware capabilities
- seL4[2009]
 - Based on L4 microkernel
 - Formally verified in Isabelle
 - Uses capabilities as access control mechanisms
 - Does not rely on special hardware
- CheriBSD [2014]
 - Adoption of FreeBSD to run on CHERI architecture
 - Uses CHERI's capability-based memory protection

Object Capabilities

Object-capability model

based on [Miller, 2006]

- In object-oriented model, there is no need to describe access control as a separate concern
 - All we need is abstraction and modularity (that are highly emphasized already)
- Compare:
 - Principle of information hiding:
 - abstractions should hand out information only on a need to know basis
 - Principle of least privilege
 - authority should be handed out only on a need to know basis
- Access abstractions through encapsulation

Object-capability model

- Object model (recall from prior courses)
 - No distinction between subjects and objects – “everything is an object”
 - Objects consist of code and state
 - *State* is a mutable collection of *references* objects
 - Objects behave according to their code
 - Interaction between objects by sending messages over references
- Object-capability model
 - No distinction between subjects and objects – “everything is an object”
 - Objects consist of code and state
 - *State* is a mutable collection of *references* objects
 - Objects behave according to their code
 - Interaction between objects **only** by sending messages over references
 - Reference graph == access graph

Terminology

Object	Instance or data
Instance	Combination of data and code <ul style="list-style-type: none">- Behavior of instance described by code
Reference	Access to an object indivisible combining <ul style="list-style-type: none">- designation of object- access permission- means of access
Capability	Reference to non-data

Object-capability model: graph dynamics

References can only be obtained according to the following rules

Connectivity by

- Initial conditions
 - An initial access graph dictates the initial references in the system.
- Parenthood
 - At the moment of object creation, the creator holds the sole reference for the newly created object.
- Endowment
 - At object creation, the creator may make some of their references available
- Introduction
 - References can be sent in messages.
 - This is the only way to obtain existing references

Object orientation and object capabilities

- Not every language with objects follows the object-capability model
- Languages that do not follow object capability model
 - C++
 - possible to forge references by casting integers to pointers
 - C# and Java
 - public interface is circumvented when instance variables are assigned indirectly
- What does it take to make Java satisfy object-capability model?

Differences between Joe-E and Java

Joe-E programs are Java programs with added restrictions

- Part of reflection API disabled
 - Reflection API can circumvent private access modifiers
- Prohibit overwriting finalize ()
- Limit global scope
 - Java programs have *ambient authority* over outside world
 - In Joe-E
 - No ambient authority
 - Dependency injection required
- Tamed Java library
 - Library contains constructors, methods, and fields that grants arbitrary authority.
 - Remove these from library or make safe versions when possible
 - getParentFile()
 - File (File dir, String child) - File (null, path)

Differences between Joe-E and Java

Further restrictions

- Only allow static final Immutable field
 - Immutable
 - Joe-E tagging interface
 - Deeply immutable
 - Static fields not associated with object – ambiently available
- All exceptions must be Immutable
 - Means for transferring control and references unexpectedly
 - Immutable exceptions cannot transfer capabilities
- Prevent catching Errors
 - Prevent finally keyword – can be used to catch Errors
 - finally for Exceptions can be encoded with try and catch

Revoking and limiting authority

- In Joe-E and other object capability languages:
 - Capabilities cannot be restricted
 - Solved by facet patterns
 - Defines additional interface for manipulating state
 - Capabilities are irrevocable
 - Solved by: Revocable capability patterns
 - e.g., Java decorator pattern

Facet pattern in Joe-E

Restricting authority

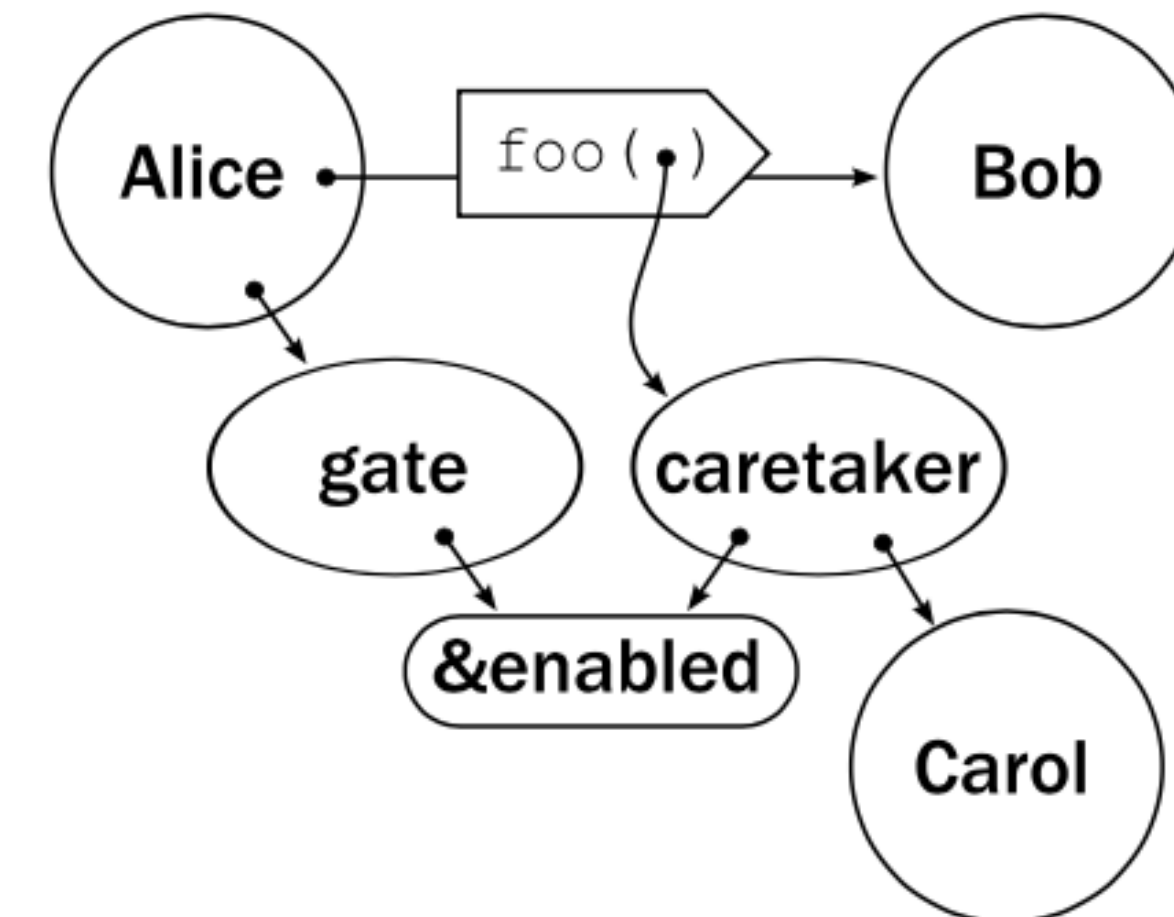
```
class Queue {
    public Object dequeue() {
        ...
    }
    public void enqueue(Object o) {
        ...
    }
}

class Queue {
    public Object dequeue() {
        ...
    }
    public void enqueue(Object o) {
        ...
    }
    public Receiver enqueuer() {
        return new Receiver() {
            public void receive(Object x) {
                enqueue(x);
            }
        };
    }
}
```

Caretaker pattern in E

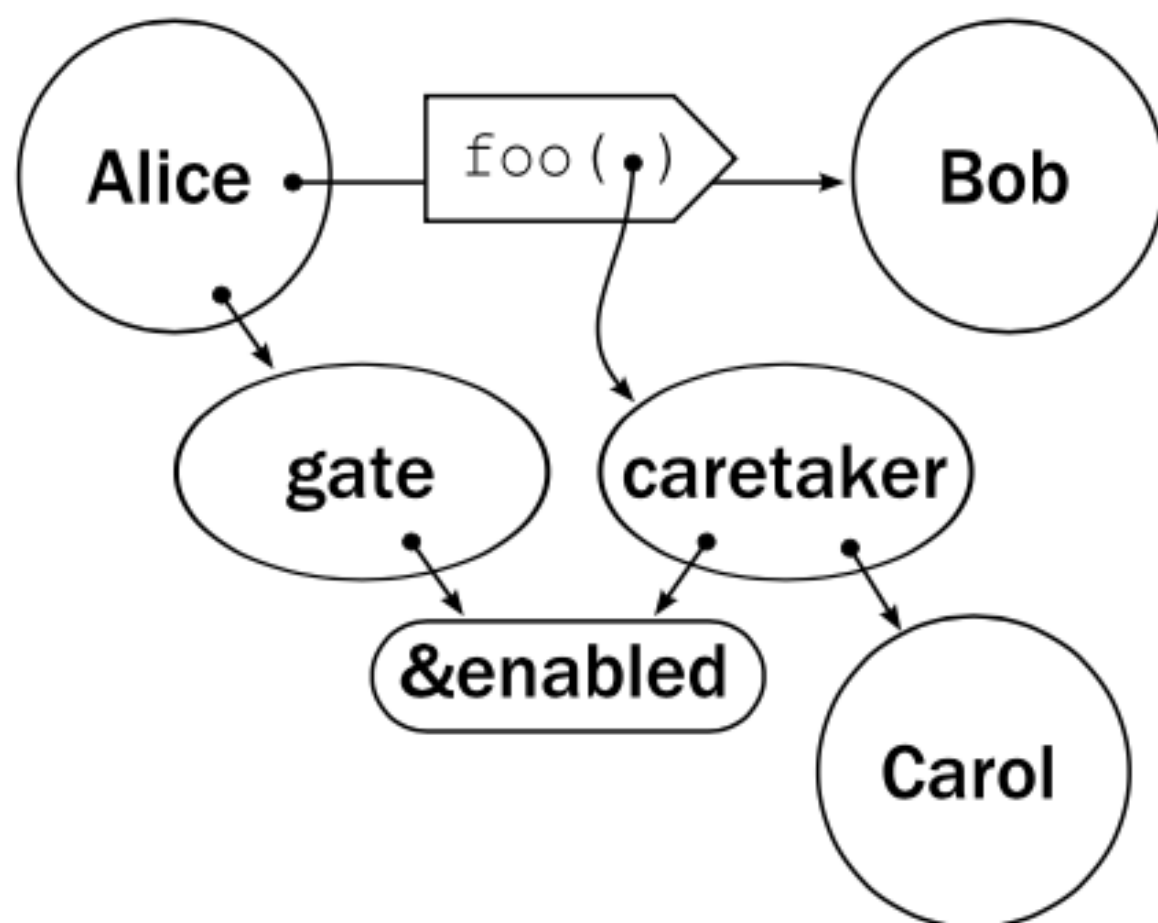
revocation with capabilities

```
def makeCaretaker(target) {  
  var enabled := true  
  def caretaker {  
    match [verb, args] {  
      if (enabled) {  
        E.call(target, verb, args)  
      } else {  
        throw("disabled")  
      }  
    }  
  }  
  def gate {  
    to enable() { enabled := true }  
    to disable() { enabled := false }  
  }  
  return [caretaker, gate]  
}
```



Membrane pattern in E

```
def makeCaretaker(target) {
  var enabled := true
  def caretaker {
    match [verb, args] {
      if (enabled) {
        E.call(target, verb, args)
      } else {
        throw("disabled")
      }
    }
  }
  def gate {
    to enable() { enabled := true }
    to disable() { enabled := false }
  }
  return [caretaker, gate]
}
```




Revocation with capabilities

```
def makeMembrane(target) {
  var enabled := true
  def wrap(wrapped) {
    if (Ref.isData(wrapped)) {
      # Data provides only irrevocable knowledge, so don't
      # bother wrapping it.
      return wrapped
    }
  }
  def caretaker {
    match [verb, args] {
      if (enabled) {
        def wrappedArgs := map(wrap, args)
        wrap(E.call(wrapped, verb, wrappedArgs))
      } else {
        throw("disabled")
      }
    }
  }
  return caretaker
}
def gate {
  to enable() { enabled := true }
  to disable() { enabled := false }
}
return [wrap(target), gate]
}
```

Object-capability languages and systems

- E, Joe-E
- Google Caja
 - Capability safe subset of JavaScript
 - SafeEcmaScript
 - <https://github.com/tc39/proposal-ses>
- Shill
 - Capability-based shell scripting language
- Midori



Founder/CEO [Pulumi](#) • Cloud, languages, and developer tools guy • Eat, sleep, code, repeat

Popular:

[Pulumi](#)

[Midori](#)

[Software Leadership](#)



© Joe Duffy, 2018

November 3, 2015

Blogging about Midori

Enough time has passed that I feel safe blogging about my prior project here at Microsoft, “Midori.” In the months to come, I’ll publish a dozen-or-so articles covering the most interesting aspects of this project, and my key take-aways.

Midori was a research/incubation project to explore ways of innovating throughout Microsoft’s software stack. This spanned all aspects, including the programming language, compilers, OS, its services, applications, and the overall programming models. We had a heavy bias towards cloud, concurrency, and safety. The project included novel “cultural” approaches too, being 100% developers and very code-focused, looking more like the Microsoft of today and hopefully tomorrow, than it did the Microsoft of 8 years ago when the project began.

I worked on Midori from 2009 until we transitioned the teams to their respective new homes during 2012-2014. I led the groups focusing on the developer experience: language, compilers, core frameworks, concurrency models, and IDEs/tools. And I wrote lots of code the whole time.

Although we started with C# and .NET, we were forced to radically depart in the name of security, reliability, and performance. Now, I am helping to bring many of those lessons learned back to the shipping products including, perhaps surprisingly, C++. Most of my blog entries will focus on the key lessons that we’re now trying to apply back to the products, like asynchrony everywhere, zero-copy IO, dispelling the false dichotomy between safety and performance, capability-based security, safe concurrency, establishing a culture of technical debate, and more.

I’ll be the first to admit, none of us knew how Midori would turn out. That’s often the case with research. My biggest regret is that we didn’t OSS it from the start, where the meritocracy of the Internet could judge its pieces appropriately. As with all big corporations, decisions around the destiny of Midori’s core technology weren’t entirely technology-driven, and sadly, not even entirely business-driven. But therein lies some important lessons too. My second biggest regret is that we didn’t publish more papers. This blog series may help to recitify some of this.

I shall update this list as new articles are published:

1. [A Tale of Three Safeties](#)
2. [Objects as Secure Capabilities](#)
3. [Asynchronous Everything](#)
4. [Safe Native Code](#)
5. [The Error Model](#)
6. [Performance Culture](#)
7. [15 Years of Concurrency](#)

Midori was a fascinating journey, and the most fun I’ve had in my career to-date. I look forward to sharing some of that journey with you.

Summary

- Capability concept
 - Motivation: example of confused deputy
 - Capabilities
 - unforgeable token of authority
 - combine designation and access
- Object Capability Model
 - Principle of encapsulation coincide with least privilege
 - Capability safety requires restrictions on many language features
- Patterns for programming with capabilities