

Evidence-based Audit



Jeff Vaughan
Limin Jia, Karl Mazurak, and Steve Zdancewic

Department of Computer and Information Science
University of Pennsylvania


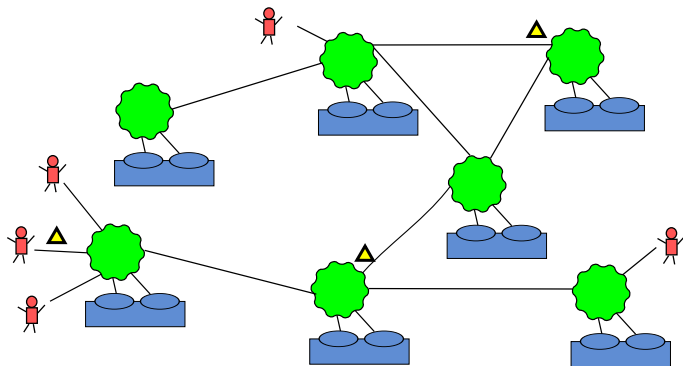
CSF/LICS Joint Session
June 24, 2008





Our Setting: Distributed Access Control





Our Setting: Distributed Access Control




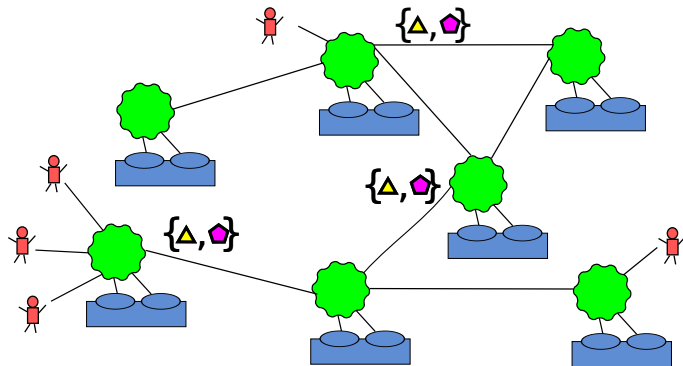
Our Setting: Distributed Access Control



Our Setting: Distributed Access Control



Key Idea: Proofs attest to data integrity.



The Aura Project

- A programming language called Aura
 - A propositional fragment, modeled here by $Aura_0$
 - An ML-like computation language [Jia+ 08]
- A security aware programming model
 - active, potentially malicious principals
 - mutual distrust between applications and principals
 - emphasis on access control and audit
- An implementation including compiler and .Net-based runtime


The Aura Project

- A programming language called Aura
 - A propositional fragment, modeled here by $Aura_0$
 - An ML-like computation language [Jia+ 08]
- A security aware programming model
 - active, potentially malicious principals
 - mutual distrust between applications and principals
 - emphasis on access control and audit
- An implementation including compiler and .Net-based runtime


Today's Talk

Analyzing the local security of Aura applications.

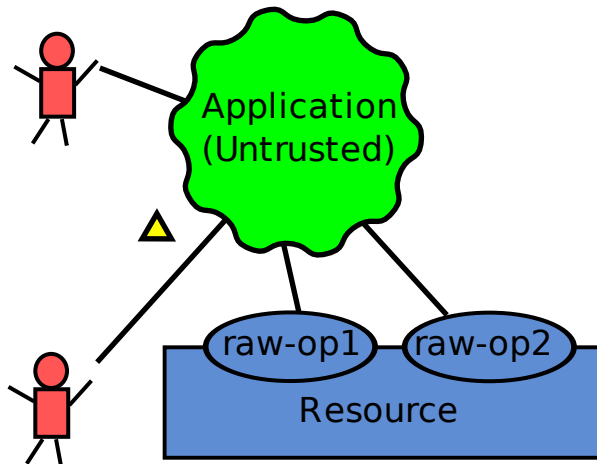
Access control, locally.




Access control, locally.




Access control, locally.




Access control, locally.




Access control, locally.




Access control, locally.




Access control, locally.




Access control, locally.




Access control, locally.



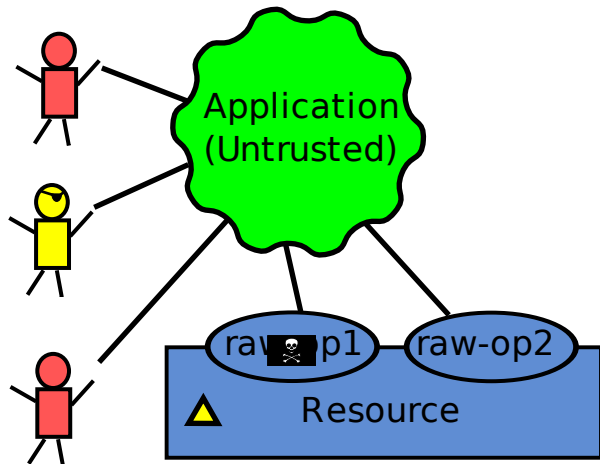
Access control, locally.



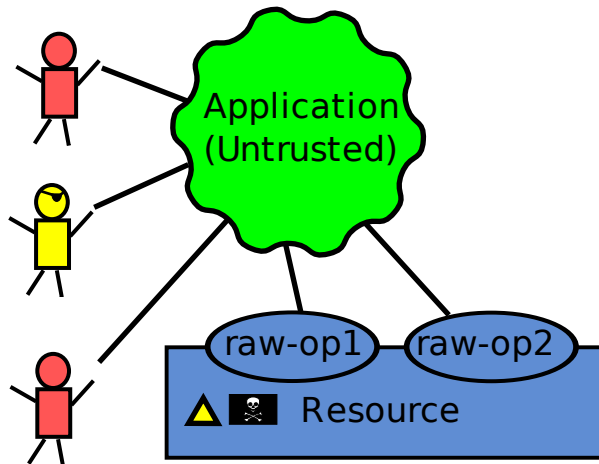
Access control, locally.



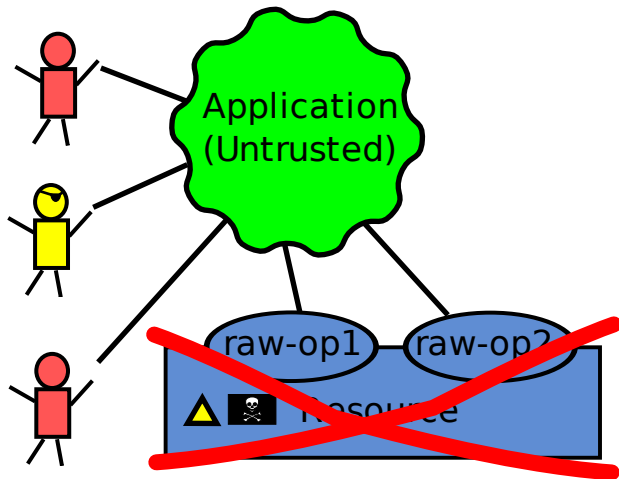
Access control, locally.



Access control, locally.



Access control, locally.



Jargon needed for this talk

Institutional Policy Human-level rules about principals, data values, and resources.

Bad Access A system state change violating institutional policy.

Formal Rules Machine-level encoding of policy.

Jargon needed for this talk

Institutional Policy Human-level rules about principals, data values, and resources.

Bad Access A system state change violating institutional policy.

Formal Rules Machine-level encoding of policy.

Foreshadowing

Users care about *institutional policy*, but technology tries to enforce *formal rules*.

Why do reference monitors allow bad accesses to occur?

Problem 1

The trusted computing base's implementation may be buggy.

Problem 2

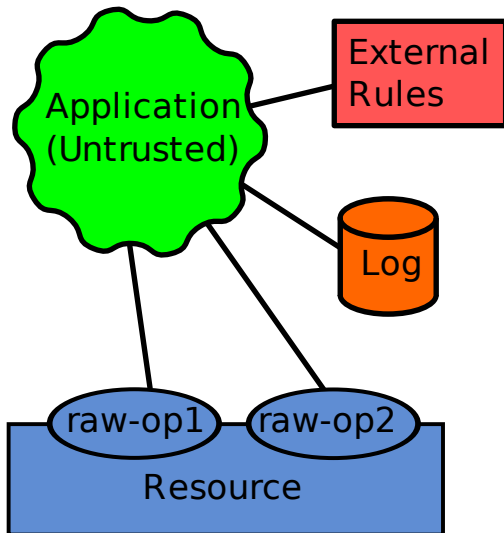
The formal rules language may be too impoverished to express institutional policy.

Problem 3

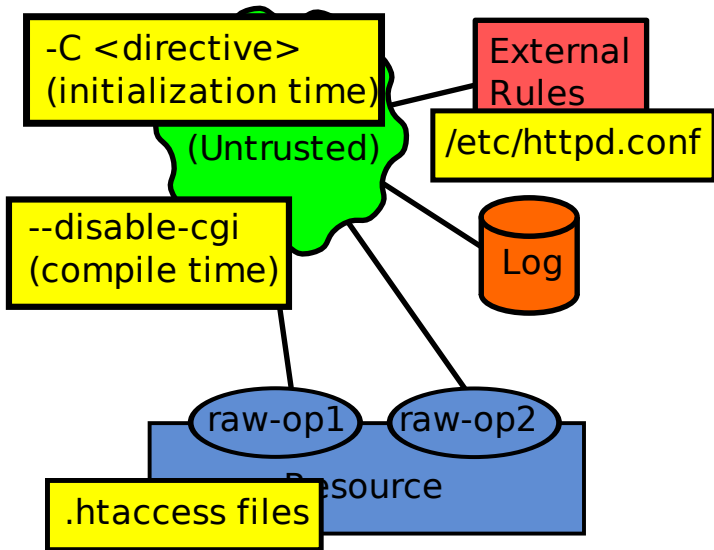
The system may be configured with incorrect formal rules.

And many other reasons not addressed here. . .

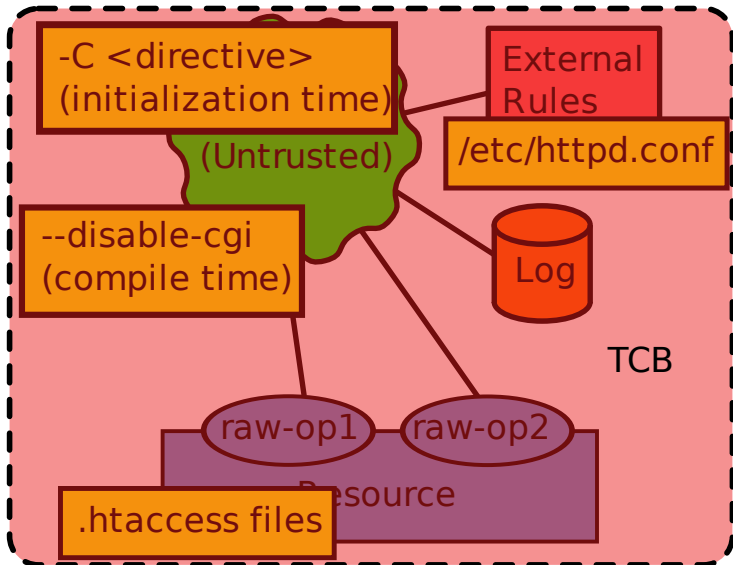
Common application design exacerbates these problems.



Common application design exacerbates these problems.



Common application design exacerbates these problems.




Problem 1


The trusted computing base's implementation may be buggy.

Aura Solution

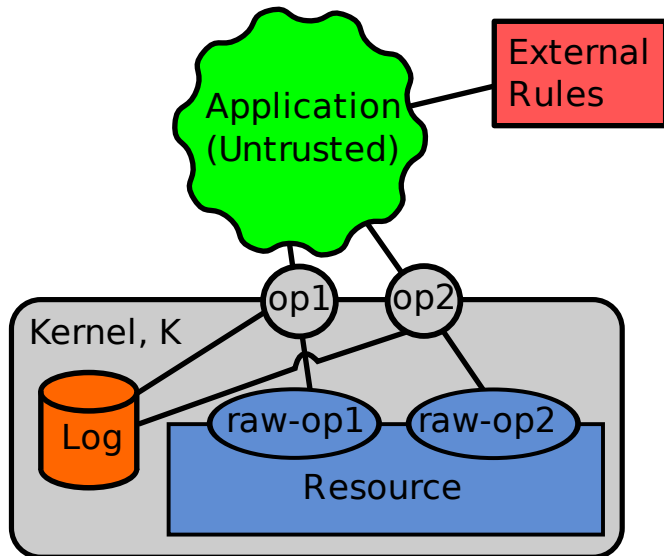
Trust only a small, generic *kernel* that has no application-specific functionality.

 [Saltzer+ 75], [Bauer+ 99], [Jia+ 08]


In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.



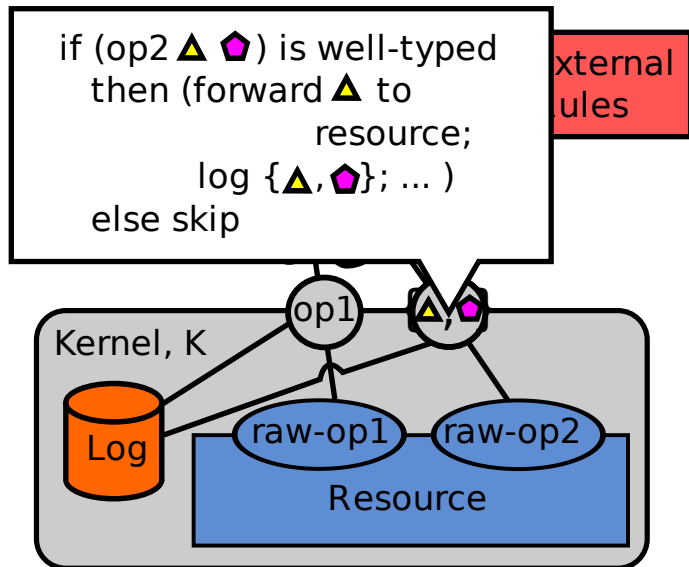
In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.



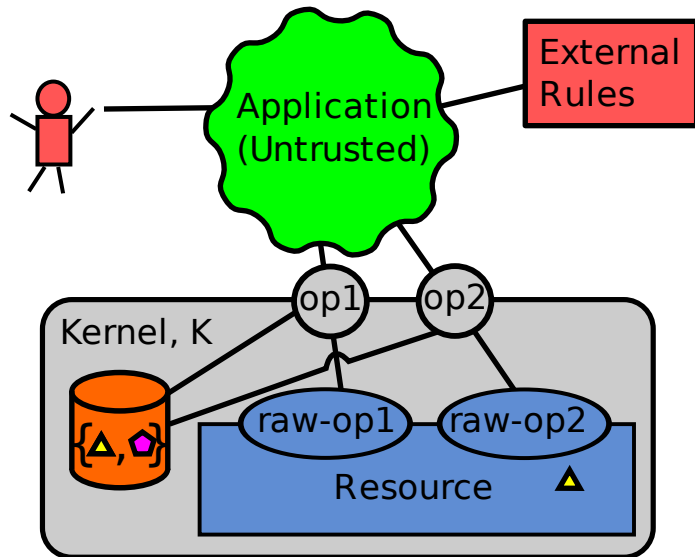
In Aura, a lightweight kernel protects resources.



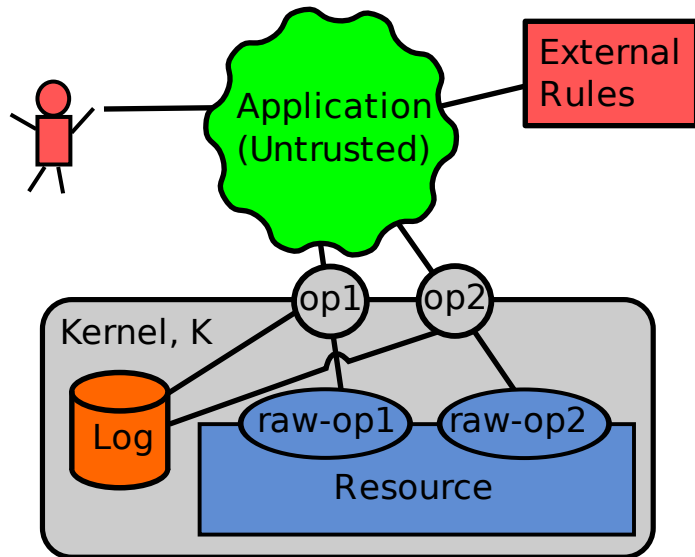
In Aura, a lightweight kernel protects resources.



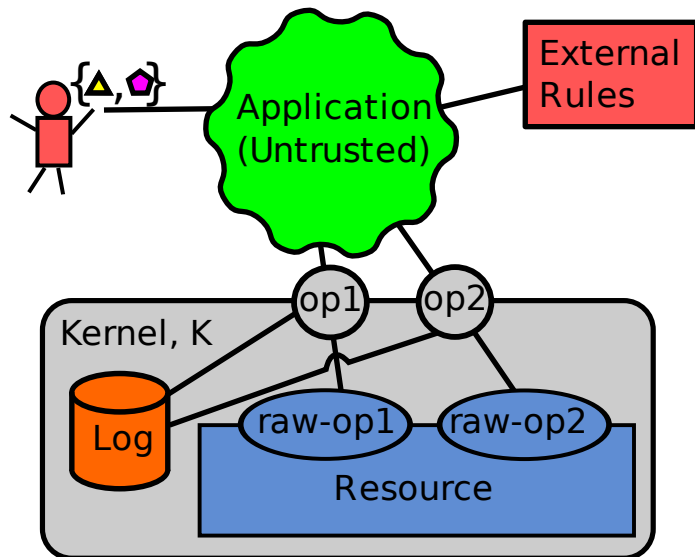
In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.



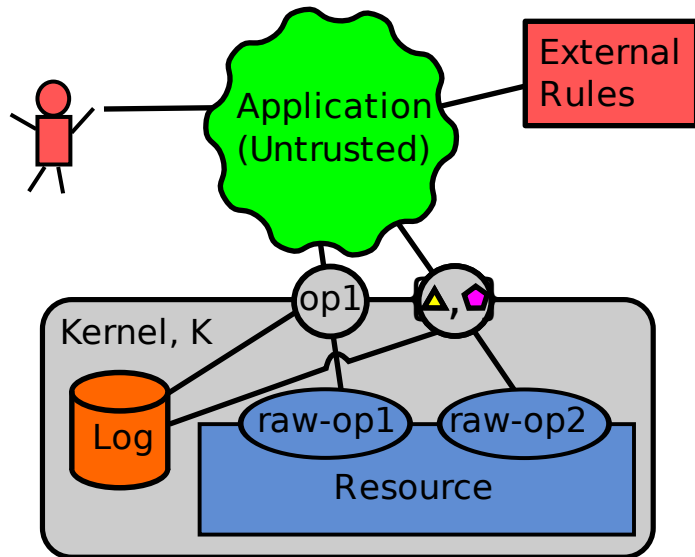
In Aura, a lightweight kernel protects resources.



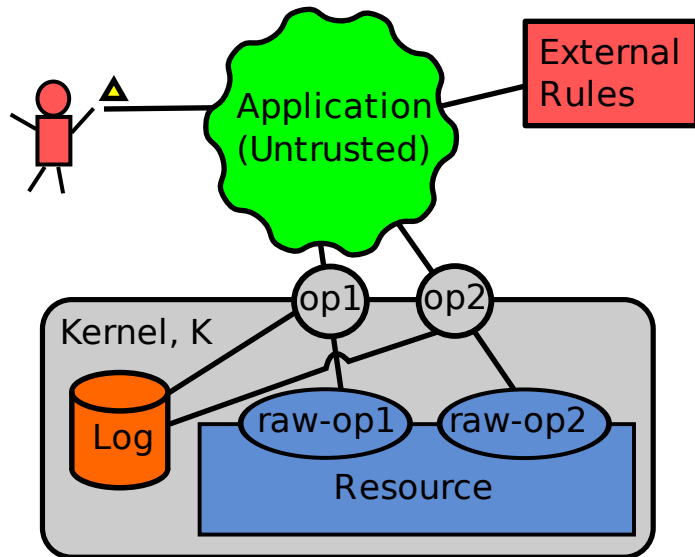
In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.



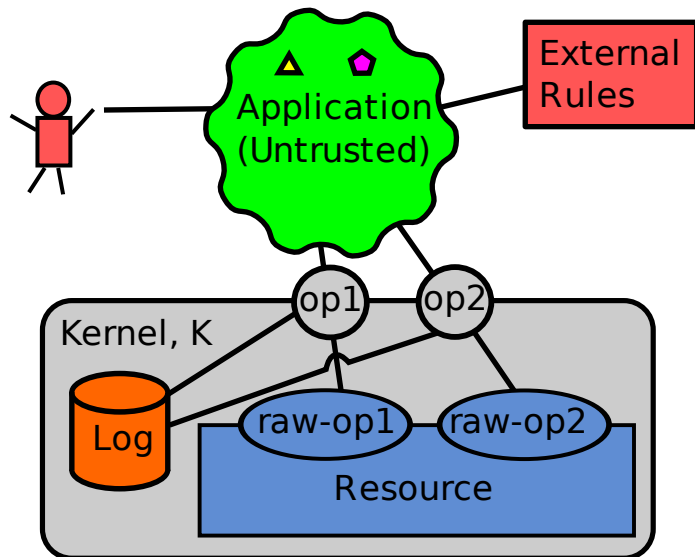
In Aura, a lightweight kernel protects resources.



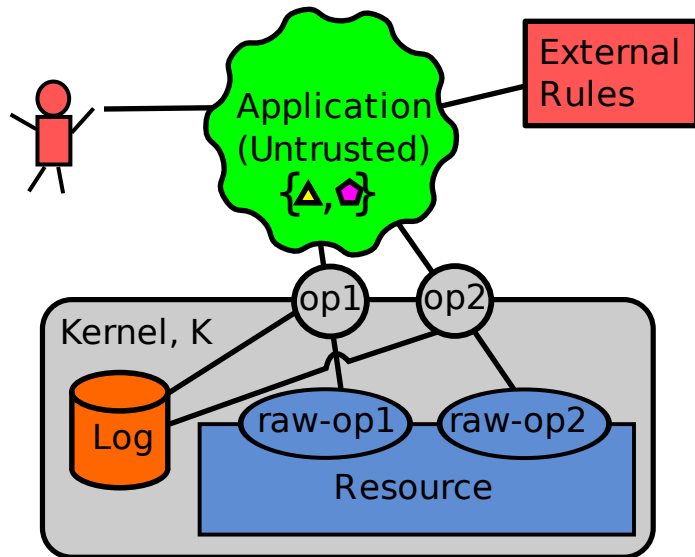
In Aura, a lightweight kernel protects resources.




In Aura, a lightweight kernel protects resources.



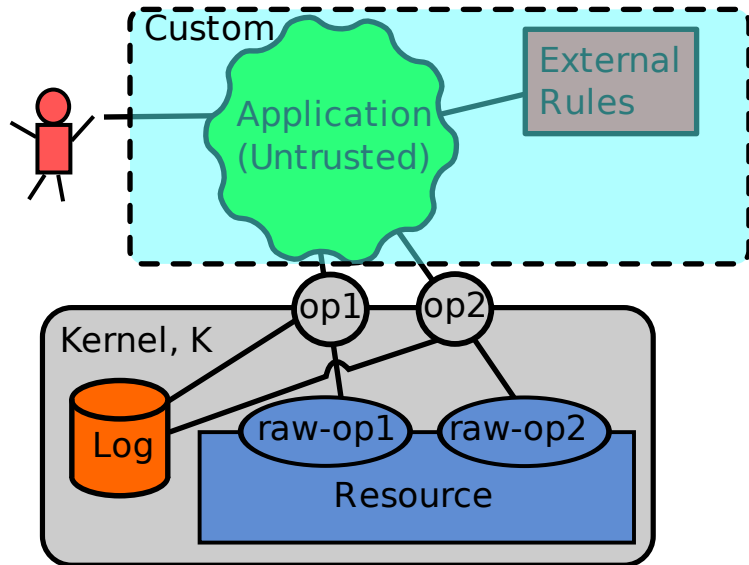
In Aura, a lightweight kernel protects resources.



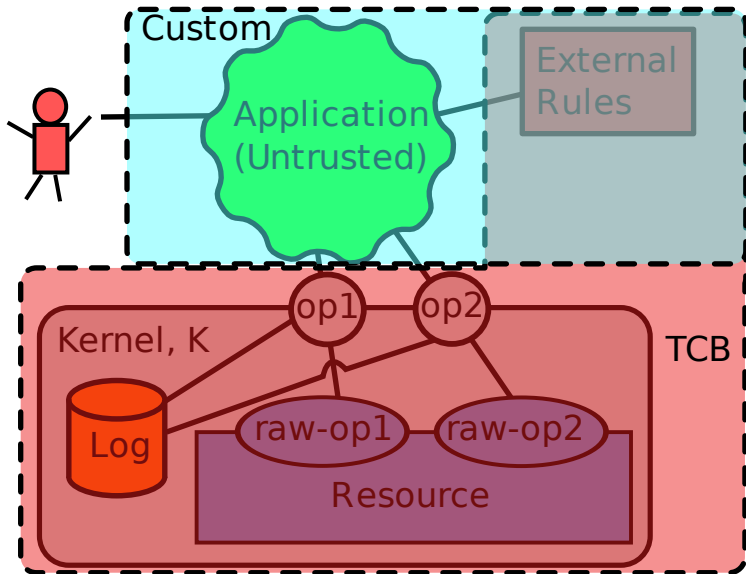
In Aura, a lightweight kernel protects resources.



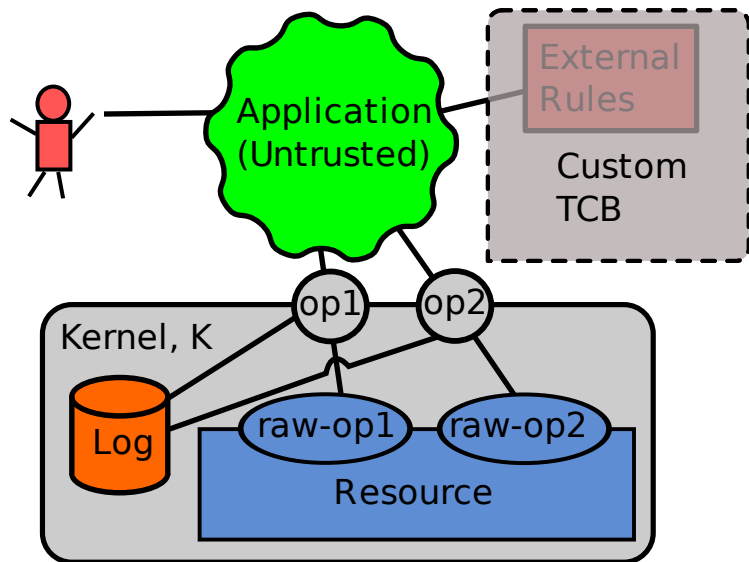
In Aura, a lightweight kernel protects resources.



In Aura, a lightweight kernel protects resources.



In Aura, a lightweight kernel protects resources.



Example: Types for a remote procedure call server.

The RPC server resource has a single raw operation

$$\text{raw-rpc} : \mathbf{string} \Rightarrow \mathbf{string}$$

The kernel exposes an *extended signature* of propositions.

$$\text{OkToRPC} : \mathbf{string} \rightarrow \mathbf{Prop}$$
$$\text{DidRPC} : \mathbf{string} \rightarrow \mathbf{string} \rightarrow \mathbf{Prop}$$

Applications are written against the *kernel interface*.

$$\begin{aligned} \text{rpc} : (x : \mathbf{string}) &\Rightarrow (\text{Kernel } \mathbf{says} \text{ OkToRPC } x) \\ &\Rightarrow \{y : \mathbf{string}; \text{Kernel } \mathbf{says} \text{ DidRPC } x \ y\} \end{aligned}$$

Example: Types for a remote procedure call server.

The RPC server resource has a single raw operation

$$\text{raw-rpc} : \mathbf{string} \Rightarrow \mathbf{string}$$

The kernel exposes an *extended signature* of propositions.

$$\text{OkToRPC} : \mathbf{string} \rightarrow \mathbf{Prop}$$
$$\text{DidRPC} : \mathbf{string} \rightarrow \mathbf{string} \rightarrow \mathbf{Prop}$$

Applications are written against the *kernel interface*.

$$\begin{aligned} \text{rpc} : (x : \mathbf{string}) &\Rightarrow (\text{Kernel } \mathbf{says} \text{ OkToRPC } x) \\ &\Rightarrow \{y : \mathbf{string}; \text{Kernel } \mathbf{says} \text{ DidRPC } x \ y\} \end{aligned}$$


The kernel interface and extended signature
may be mechanically generated.

Problem 2

The formal rules language may be too impoverished to express institutional policy.

Aura Solution

Use dependent DCC with signature objects to specify rules.

 [Abadi+ 06], [Fournet+ 07], [Bengtson+ 08] ...

Aura's **says** modality represents affirmation. (1/2)

$$\frac{\Gamma \vdash P : \mathbf{Prop}}{\Gamma \vdash A \mathbf{says} P : \mathbf{Prop}}$$

The proposition “Principal A affirms proposition P .”

$$\frac{\Gamma \vdash P : \mathbf{Prop}}{\Gamma \vdash \mathbf{sign}(A, P) : A \mathbf{says} P}$$

A 's signature on P . P might be unprovable.

Aura's **says** modality represents affirmation. (1/2)

$$\frac{\Gamma \vdash P : \mathbf{Prop}}{\Gamma \vdash A \mathbf{says} P : \mathbf{Prop}}$$

The proposition “Principal A affirms proposition P .”

$$\frac{\Gamma \vdash P : \mathbf{Prop}}{\Gamma \vdash \mathbf{sign}(A, P) : A \mathbf{says} P}$$

A 's signature on P . P might be unprovable.



Aura's **says** modality represents affirmation. (2/2)

$$\frac{\Gamma \vdash p : P}{\Gamma \vdash \mathbf{return}@[A]p : A \mathbf{says} P}$$

A affirms proven propositions.

$$\frac{\begin{array}{l} \Gamma \vdash p : A \mathbf{says} P \\ \Gamma, x : P \vdash q : A \mathbf{says} Q \end{array}}{\Gamma \vdash \mathbf{bind} x = p \mathbf{in} q : A \mathbf{says} Q}$$

A affirms the result of hypothetical reasoning. We can “reason from A 's point of view.”

(These are standard monad rules.)

Aura₀'s syntax supports restricted dependent types.

Partial syntax

$t ::=$	string prin	Base types
	x a	Variables and constants
	t says t	Says modality
	$(x: t) \rightarrow t$	Logical implication/quantification
	$(x: t) \Rightarrow t$	Computational arrows
	$\{x: t; t\}$	Dependent pair type
	$t t$	Application
	\vdots	

Syntactic separation of computation and logical arrows stop arbitrary resource-effects from polluting the logic.

Dependent types allow for expressive rules.

Example (Bob acts for Alice)

Alice **says** $((P: \mathbf{Prop}) \rightarrow \text{Bob says } P \rightarrow P)$

Dependent types allow for expressive rules.

Example (Bob acts for Alice)

Alice **says** $((P: \mathbf{Prop}) \rightarrow \text{Bob says } P \rightarrow P)$

Example (Bob acts for Alice only regarding validity)

Alice **says** $((x: \mathbf{string}) \rightarrow \text{Bob says valid } x \rightarrow \text{valid } x)$

Dependent types allow for expressive rules.

Example (Bob acts for Alice)

Alice **says** $((P: \mathbf{Prop}) \rightarrow \text{Bob says } P \rightarrow P)$

Example (Bob acts for Alice only regarding validity)

Alice **says** $((x: \mathbf{string}) \rightarrow \text{Bob says valid } x \rightarrow \text{valid } x)$

Example (Kernel allows RPC calls on strings endorsed by Alice)

Kernel **says** $((x: \mathbf{string}) \rightarrow \text{Alice says valid } x \rightarrow \text{OkToRPC } x)$

Problem 3

The system may be configured with incorrect formal rules.

Aura Solution

Consistently log runtime proof objects for later analysis.



[Wee 95], [Cederquist+ 05]

When something unexpected happens: look at the log.

- System design guarantees a one-to-one correspondence between log entries and resource state changes.
- If a Alice's signature does not appear in a log entry, she could not have caused the associated action.
- Proofs get convoluted, but *proof reduction* can restore clarity.

Example (A convoluted proof)

$(\lambda x. \lambda y. y) (\mathbf{sign}(\text{Alice}, P)) (\mathbf{sign}(\text{Bob}, Q))$

Here Alice's signature is "irrelevant."

Reduction relation includes special cases for **bind**.

R-Bind-Specious

$$\frac{x \notin \text{fv}(t_2)}{\mathbf{bind} \ x = t_1 \ \mathbf{in} \ t_2 \longrightarrow t_2}$$

Drops unused hypotheses. Most interesting when t_1 is a signature.

R-Bind-Commute

$$\frac{y \notin \text{fv}(t_3)}{\mathbf{bind} \ x = (\mathbf{bind} \ y = t_1 \ \mathbf{in} \ t_2) \ \mathbf{in} \ t_3 \longrightarrow \mathbf{bind} \ y = t_1 \ \mathbf{in} \ \mathbf{bind} \ x = t_2 \ \mathbf{in} \ t_3}$$

Commutation rule that can enable further reductions.

... plus standard β and structural rules.

Example (The same convoluted proof)

$$(\lambda x. \lambda y. y) (\mathbf{sign}(\text{Alice}, P)) (\mathbf{sign}(\text{Bob}, Q)) \longrightarrow^* \mathbf{sign}(\text{Bob}, Q)$$

Example (Reducing by special **bind** rules)

$$\begin{aligned} & \mathbf{bind} \ x = (\mathbf{bind} \ y = (f \ \mathbf{sign}(\text{Bob}, Q)) \\ & \quad \mathbf{in} \ \mathbf{sign}(\text{Alice}, P)) \\ & \quad \mathbf{in} \ x \\ & \longrightarrow^* \mathbf{sign}(\text{Alice}, P) \end{aligned}$$

Theorem (Subject Reduction)

If $p \longrightarrow p'$ and $\Gamma \vdash p : s$ then $\Gamma \vdash p' : s$.

Theorem (Confluence)

If $p \longrightarrow^ p_1$, and $p \longrightarrow^* p_2$, then there exists p_3 such that $p_1 \longrightarrow^* p_3$ and $p_2 \longrightarrow^* p_3$.*


Theorem (Strong Normalization)

If $\Gamma \vdash p : s$, then p is strongly normalizing (SN). That is, all reduction sequences starting with p halt.

Aura₀ is strongly normalizing.


Fact

The Calculus of Constructions with dependent pairs (CoC) is SN.

 [Geuvers 95]

Proof Idea

Show Aura reductions can be simulated in terminating system based on CoC.



Aura is a framework for proof carrying authorization and audit.

Aura includes

- a small and generic trusted computing base,
- an expressive authorization logic, and
- a principled audit methodology.

Aura is a framework for proof carrying authorization and audit.

Aura includes

- a small and generic trusted computing base,
- an expressive authorization logic, and
- a principled audit methodology.

Code and technical reports available from
<http://www.cis.upenn.edu/~stevez/sol/aura.html>

Or ask us for a demo!

- Typing Rules for Pair and Arrow
- Strong Normalization Proof

Typing rules for arrow and pair.

$$\frac{\begin{array}{l} \Sigma; \Gamma \vdash t_1 : k_1 \quad \Sigma; \Gamma, x : t_1 \vdash t_2 : k_2 \\ k_1 \in \{\mathbf{Kind}^P, \mathbf{Prop}, \mathbf{Type}\} \quad k_2 \in \{\mathbf{Prop}, \mathbf{Type}\} \end{array}}{\Sigma; \Gamma \vdash (x : t_1) \rightarrow t_2 : k_2}$$

$$\frac{\begin{array}{l} \Sigma; \Gamma \vdash t_1 : k_1 \quad \Sigma; \Gamma, x : t_1 \vdash t_2 : k_2 \quad k_1, k_2 \in \{\mathbf{Prop}, \mathbf{Type}\} \end{array}}{\Sigma; \Gamma \vdash \{x : t_1; t_2\} : k_1}$$

Translation from Aura to CoC erases **says**

Definition

$$\llbracket A \text{ says } P \rrbracket \approx \llbracket P \rrbracket$$

$$\llbracket \text{return}@[A]p \rrbracket \approx \llbracket p \rrbracket$$

$$\llbracket \text{bind } x = p: P \text{ in } q \rrbracket \approx (\lambda x: \llbracket P \rrbracket . \llbracket q \rrbracket) \llbracket p \rrbracket$$

$$\llbracket \text{sign}(A, P) \rrbracket \approx x \text{ fresh}$$

Translation from Aura to CoC erases **says**

Definition

$$\llbracket A \text{ says } P \rrbracket_{\Delta} \approx \llbracket P \rrbracket_{\Delta}$$

$$\llbracket \text{return}@[A]p \rrbracket_{\Delta} \approx \llbracket p \rrbracket_{\Delta}$$

$$\llbracket \text{bind } x = p: P \text{ in } q \rrbracket_{\Delta} \approx (\lambda x: \llbracket P \rrbracket_{\Delta}. \llbracket q \rrbracket_{\Delta}) \llbracket p \rrbracket_{\Delta}$$

$$\llbracket \text{sign}(A, P) \rrbracket_{\Delta} \approx \Delta(\text{sign}(A, P))$$

Where Δ maps signatures to unique fresh variables.
(We'll treat the Δ 's implicitly from now on.)

Translation from Aura to CoC erases **says**

Definition

$$\llbracket A \text{ says } P \rrbracket_{\Delta} \approx \llbracket P \rrbracket_{\Delta}$$

$$\llbracket \text{return}@[A]p \rrbracket_{\Delta} \approx \llbracket p \rrbracket_{\Delta}$$

$$\llbracket \text{bind } x = p : P \text{ in } q \rrbracket_{\Delta} \approx (\lambda x : \llbracket P \rrbracket_{\Delta} . \llbracket q \rrbracket_{\Delta}) \llbracket p \rrbracket_{\Delta}$$

$$\llbracket \text{sign}(A, P) \rrbracket_{\Delta} \approx \Delta(\text{sign}(A, P))$$

Where Δ maps signatures to unique fresh variables.
(We'll treat the Δ 's implicitly from now on.)

Lemma

If p is a well typed term in Aura_0 , then—for an appropriate context— $\llbracket p \rrbracket$ is well typed in CoC.

Constructions needs a new reduction to simulate **bind**.

Definition (CC' reduction)

The CC' reduction relation augments the standard Calculus of Construction reduction relation with

$$\frac{}{(\lambda x:t. t_1)((\lambda y:s. t_2)u) \longrightarrow (\lambda y:s. ((\lambda x:t. t_1)t_2))u} \beta'$$


β' reductions simulate R-Bind-Commute reductions:

$$\frac{y \notin fv(t_3)}{\mathbf{bind\ } x = (\mathbf{bind\ } y = t_1 \mathbf{ in\ } t_2) \mathbf{ in\ } t_3 \longrightarrow \mathbf{bind\ } y = t_1 \mathbf{ in\ } \mathbf{bind\ } x = t_2 \mathbf{ in\ } t_3}$$

Aura₀ is strongly normalizing.

Lemma

Well typed CoC terms are SN under CC' reduction.

 [Lindley 05]

Lemma

If p is a well typed Aura₀ term and $p \longrightarrow p'$. Then $\llbracket p \rrbracket \longrightarrow_{CC'}^+ \llbracket p' \rrbracket$ using one or more steps.

Proof of strong normalization.

Imagine p is an looping Aura₀ term. Then by the second lemma we can build a term which loops according to CC'. This contradicts the first lemma. □