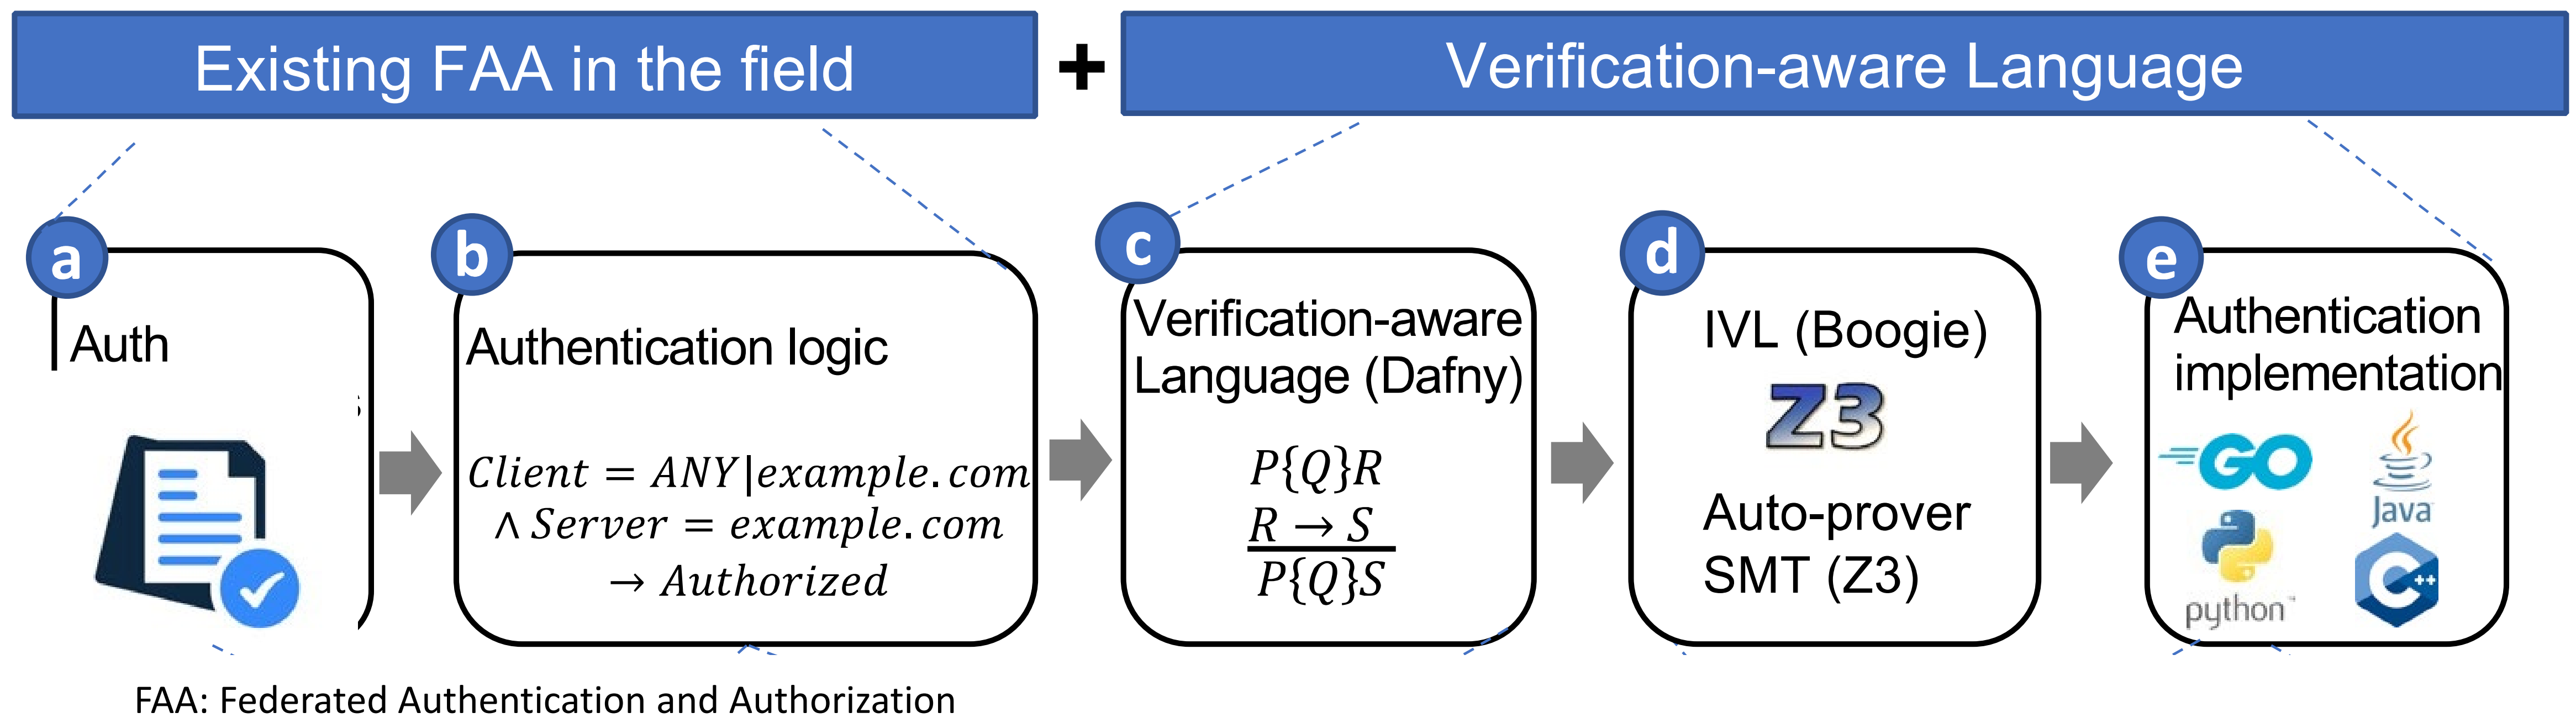


Bringing Verification-Aware Languages and Federated Authentication to Enable Secure Computing for Scientific Communities

Phuong Cao (PI), Jim Basney, Ravishankar Iyer, Anita Nikolich

<https://pmcao.github.io/projects/nsf-2319190>



FAA: Federated Authentication and Authorization

Challenges:

- Informal natural language specifications
- Complicated federated authentication logic
- Wide developer experience spectrum w.r.t security knowledge
- Operational, albeit unclear memory safety model in existing implementations
- Significant effort in validating downstream implementations upon specs revision

Scientific Impact

- Identified novel threat, failures, and uncertainty of security alerts affecting federated authentication infrastructure [IEEE QCE'24, IEEE DependSys'20, USENIX Security'24]
- Formalized a subset of critical authentication functions, using Dafny, in token-based authentication based on SciTokens and FABRIC Testbed Credential Manager [Github]
- Automated program synthesis of such implementation in Python, Java, and Rust [Github]
- Security testbed for reproducing authentication-related attacks [Secure-HPC Workshop @ Supercomputing'24]

Technical Approach Demonstration

Input:

- SciTokens specs
- Source repositories: SciTokens Python, FABRIC credential manager
- Knowledge of experts in SciTokens and FABRIC

As in RFC7519, the aud claim is not necessarily a URI. The service may accept several different possible audiences; the service endpoint at https://storage.example.com may accept an audience of either Site_Example or https://storage.example.com.

Client	Server	Result
ANY	ANY	Error
ANY	example.com	Success
example.com	ANY	Error
example.com	example.com	Success
network.com	example.com	Fail

```
method validateAudience(
  client_audience: set<string>,
  server_audience: set<string>)
  returns(c: bool)
  requires |client_audience| >= 1;
  requires |server_audience| >= 1;
{
  assert ("ANY" in server_audience) == false;
  if ("ANY" in client_audience) &&
    (|server_audience| >= 1){
    c := true;
  }
  return c;
}
var allowedAudience: set<string> := {};
allowedAudience := client_audience * server_audience;
```

```
def validateAudience(
  client_audience:
  server_audience:
  c: bool = False
  if (dafny.Seq("ANY") in (server_audience)):
    c = False
    c = c
    return c
  if ((dafny.Seq("ANY") in (client_audience)) and
    (len(server_audience) >= (1))):
    c = True
    c = c
    return c
  d_0_allowedAudience = _dafny.Set({})
  d_0_allowedAudience = _dafny.Set({})
  d_0_allowedAudience = (client_audience).intersection((server_audience))
```

Output

- Verified correct implementation of scoping, caching functionalities in SciTokens Python
- Synthesized implementations in Python, Rust, and Java

Feedback to Federated Authentication & Authorization (FAA) development team

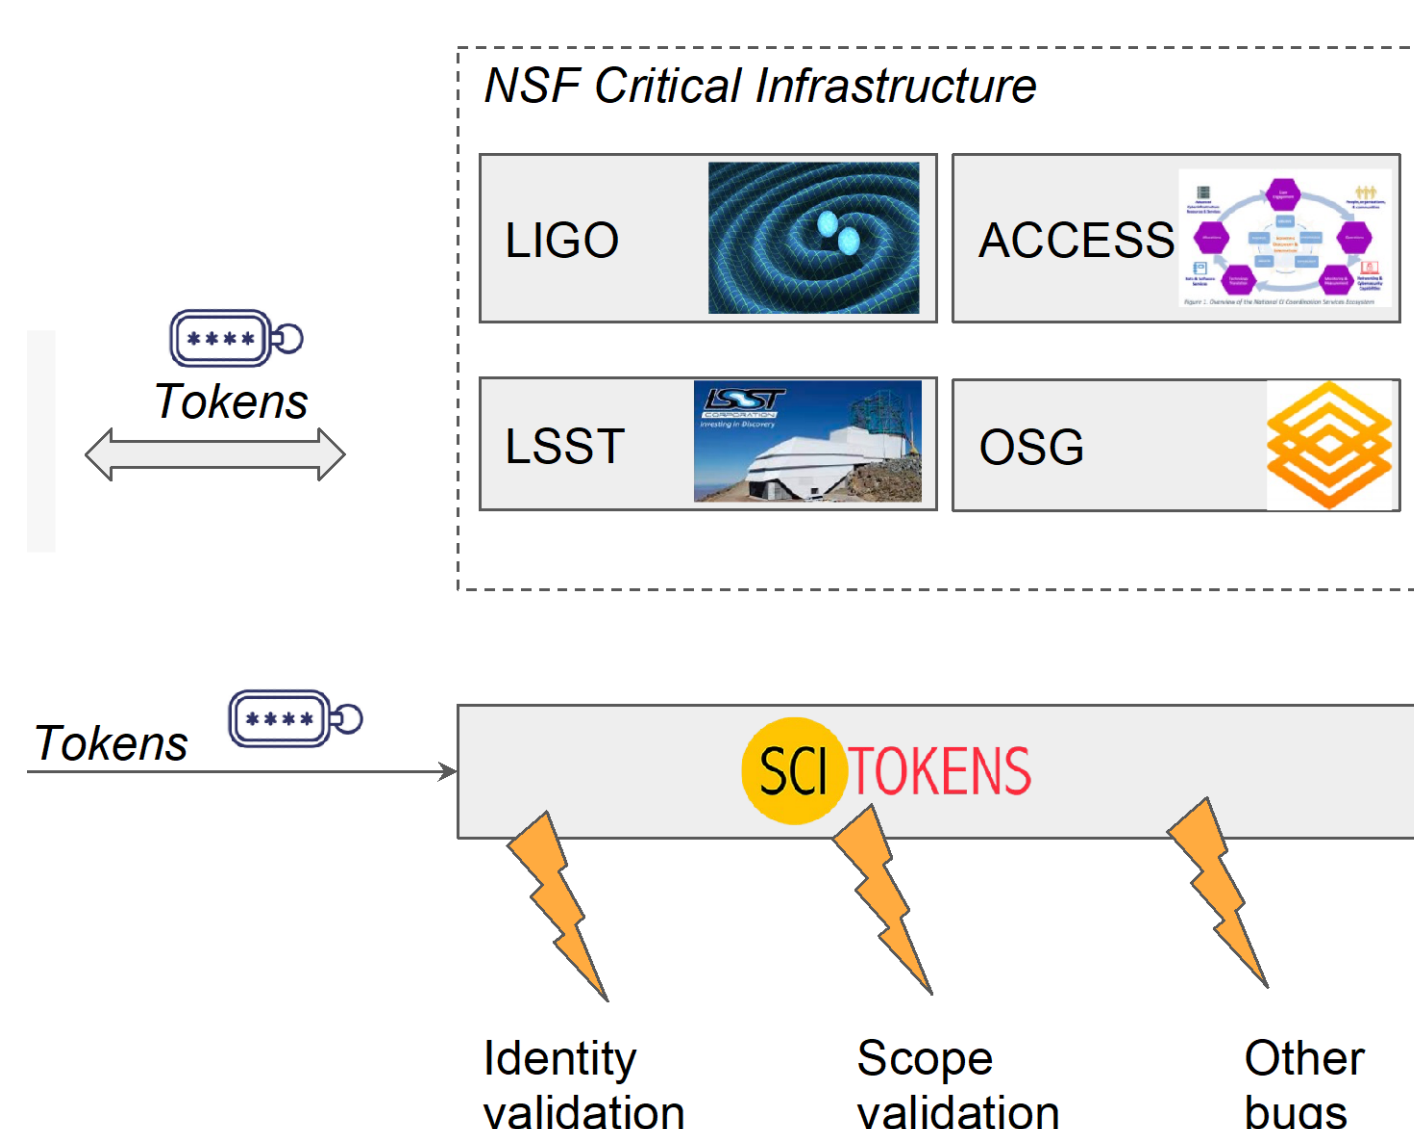
Report UNSAT (bugs)

Correct implementation

Impact on NSF Cyberinfrastructure



Impact on Security Operators

Broadening Participation in K-12



- Quantitative survey of security staff attending Supercomputing, NSF Cybersecurity Summit, and IEEE QCE
- Seminar on federated authentication and authorization in TrustedCI Webinar
- Hands-on tutorial on security log analyses with National Center for Supercomputing Applications (NCSA) Incident Response Team, and attendees located at Berkeley Lab and CMU

- Explained Federated Authentication for K-12 students through interactive game at Carrie Busey Elementary Scientific Night



References

Complete bibliography are available on project page: <https://pmcao.github.io/projects/nsf-2319190>

[NSF Research Infrastructure Workshop '24]

