Alastor: Reconstructing the Provenance of Serverless Intrusions

Pubali Datta, Isaac Polinsky, Muhammad Adil Inam, Adam Bates, William Enck

31st Usenix Security Symposium 2022

ILLINOIS

NC STATE UNIVERSITY

Serverless Cloud Computing

- Serverless Programming Model
  - Modular, ephemeral, isolated functions
  - Application logic expressed as set of workflows

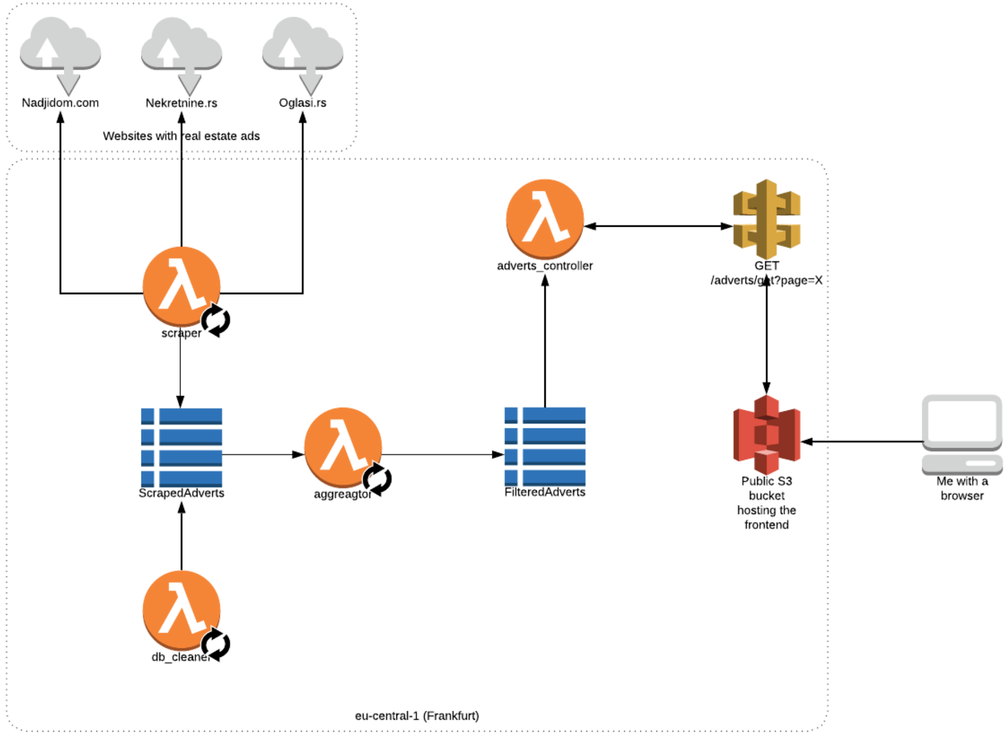

https://blog.codecentric.de/en/2018/04/serverless-application-for-scraping-and-filtering/

Serverless Cloud Computing

- Serverless Programming Model
  - Modular, ephemeral, isolated functions
  - Application logic expressed as set of workflows
- Why Popular?
  - Autoscaling
  - Less management
  - Pay per use model

https://blog.codecentric.de/en/2018/04/serverless-application-for-scraping-and-filtering/

Retail Serverless Application

Retail Website Portal

Product Purchase

Auth CC

Credit card Database

Publish Status

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

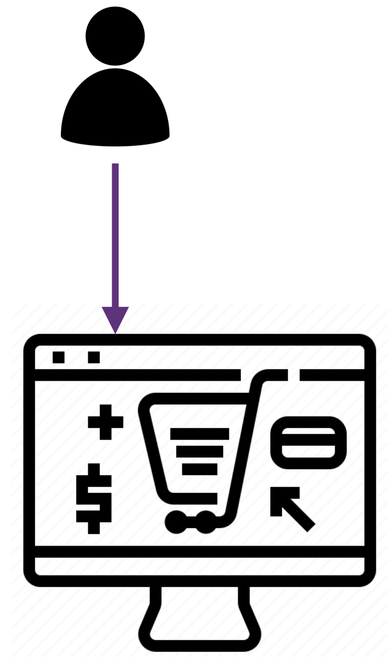

Retail Website Portal

Product Purchase

Auth CC

Credit card Database

Publish Status

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

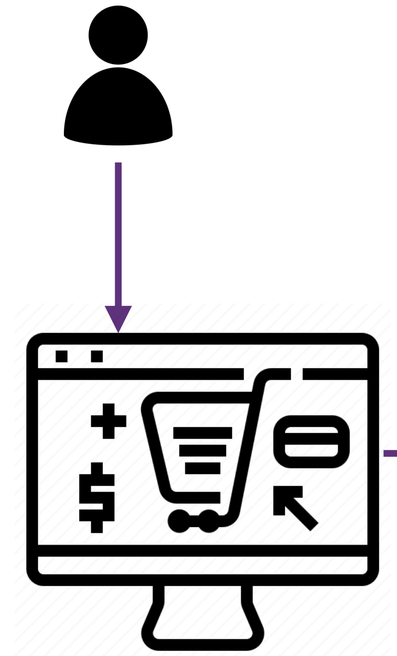

Retail Website Portal

Product Purchase

Auth CC

Credit card Database

Publish Status

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

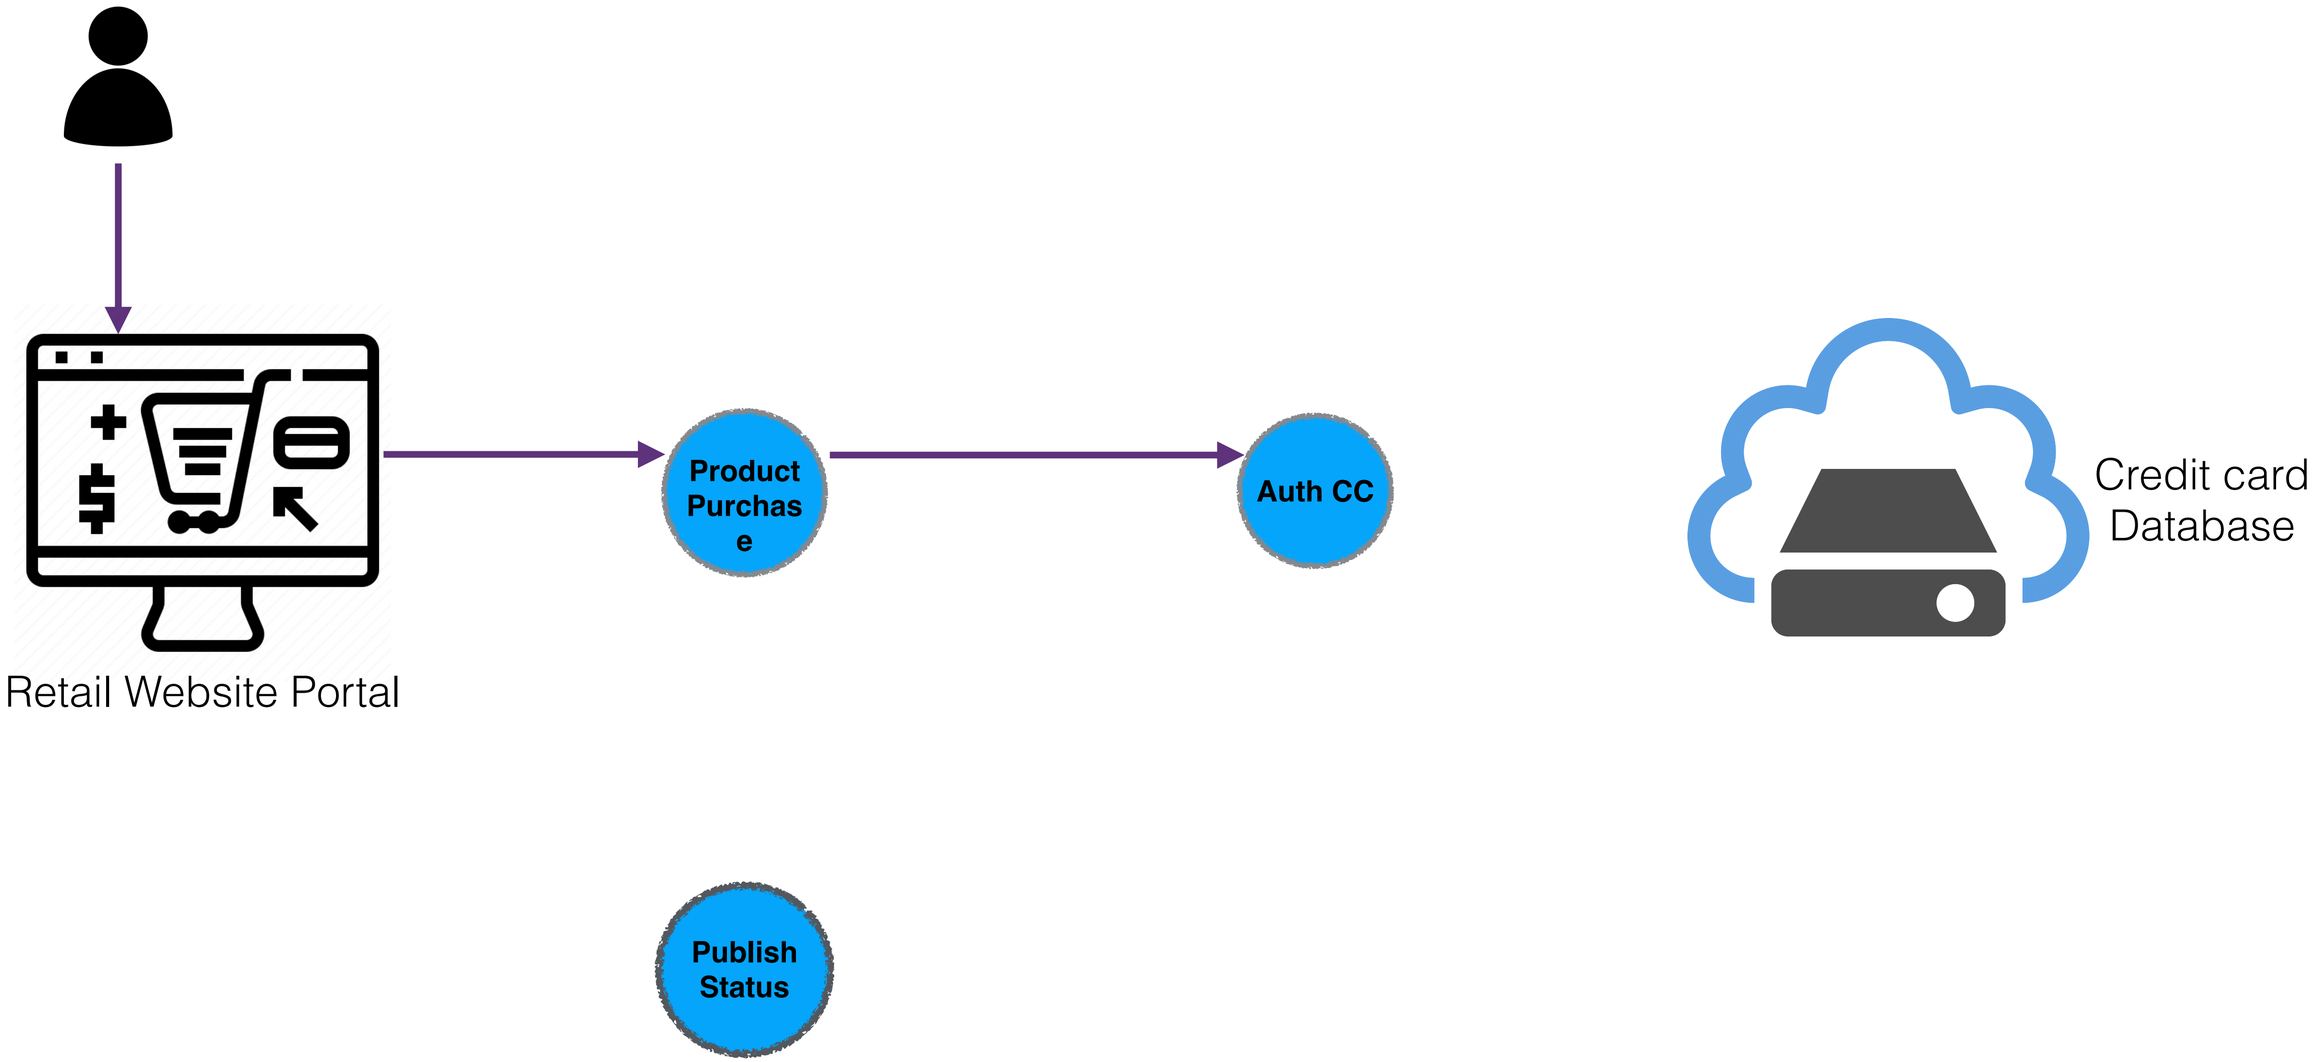

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

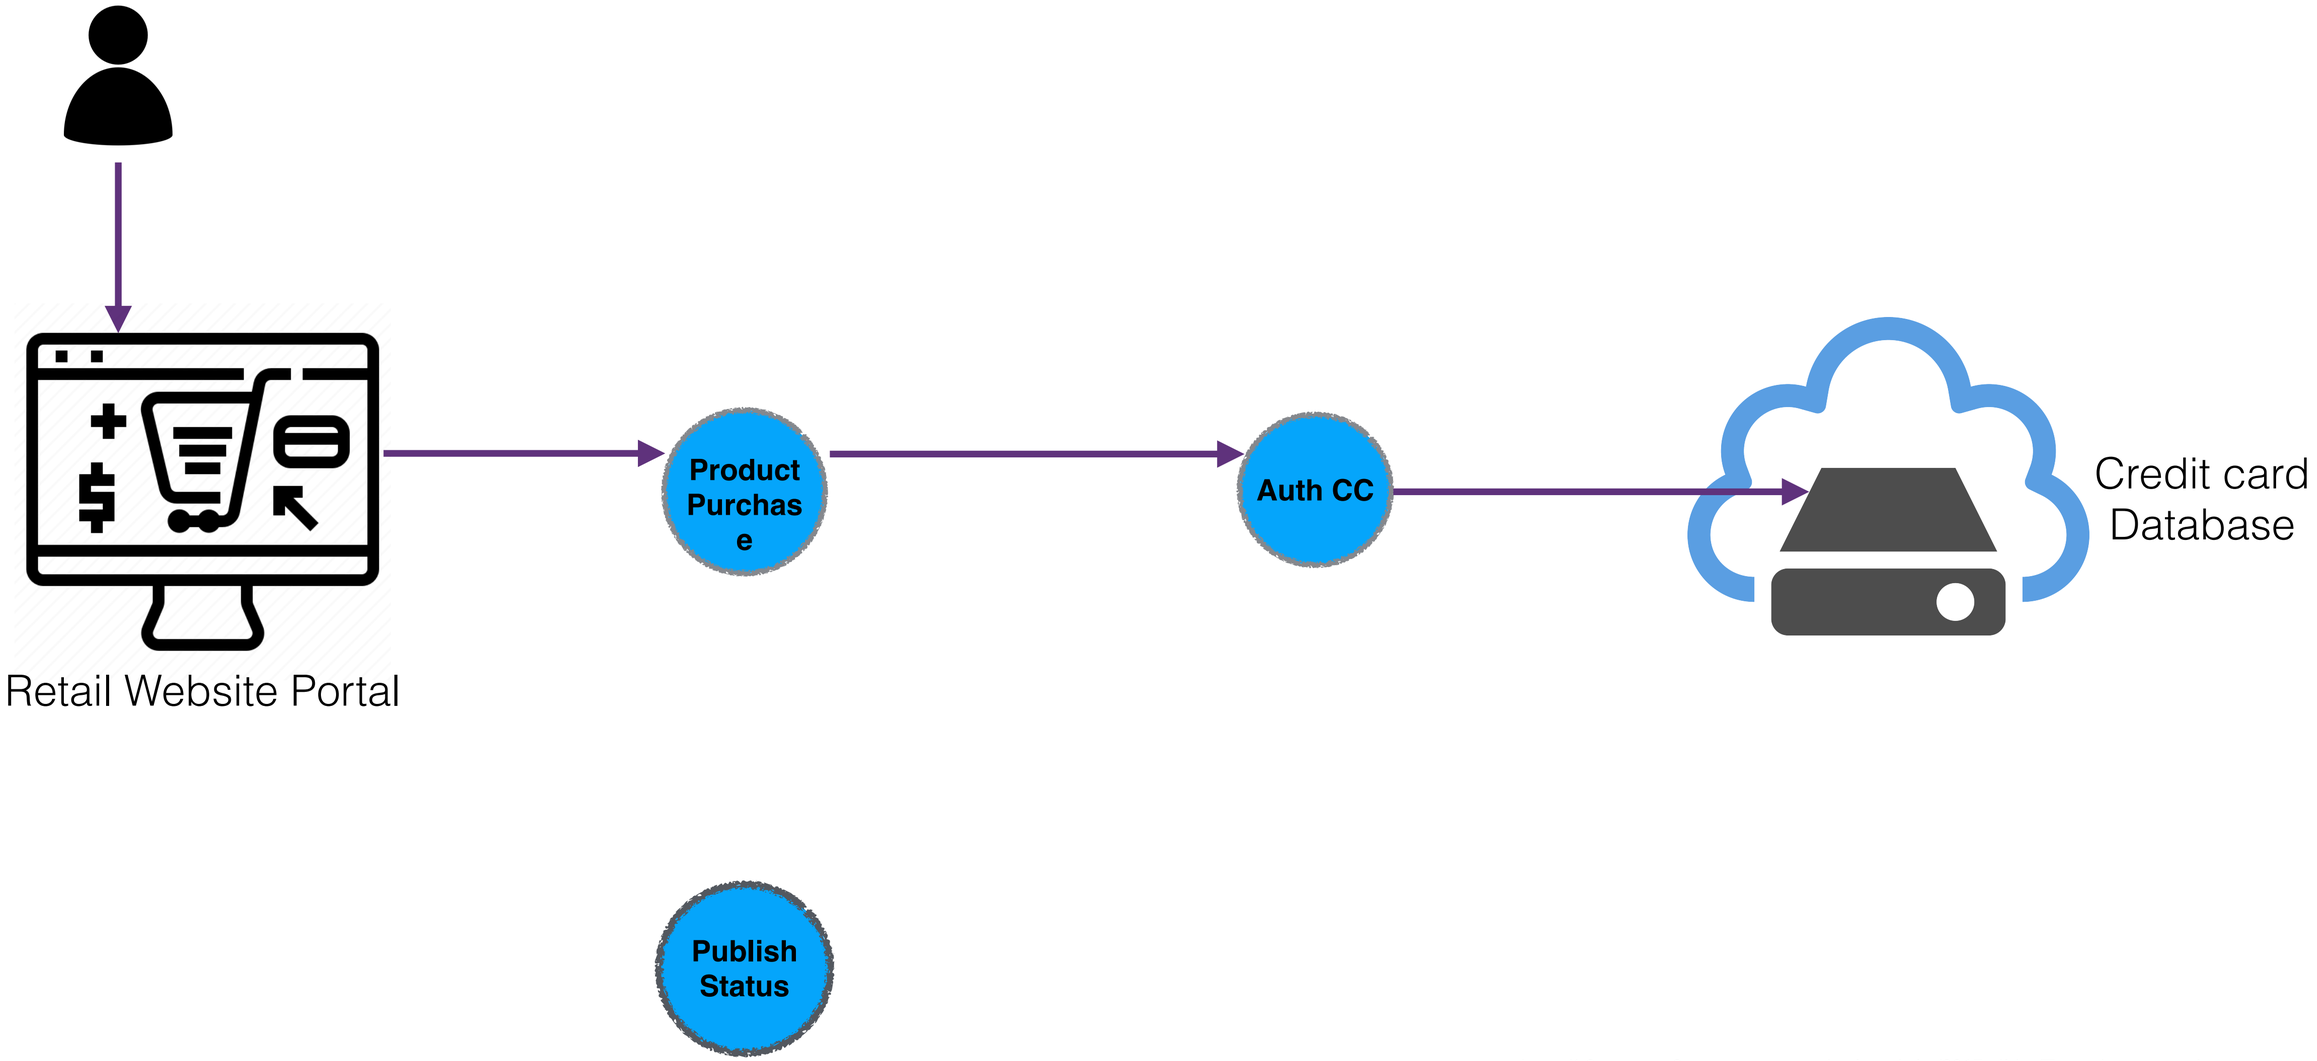

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

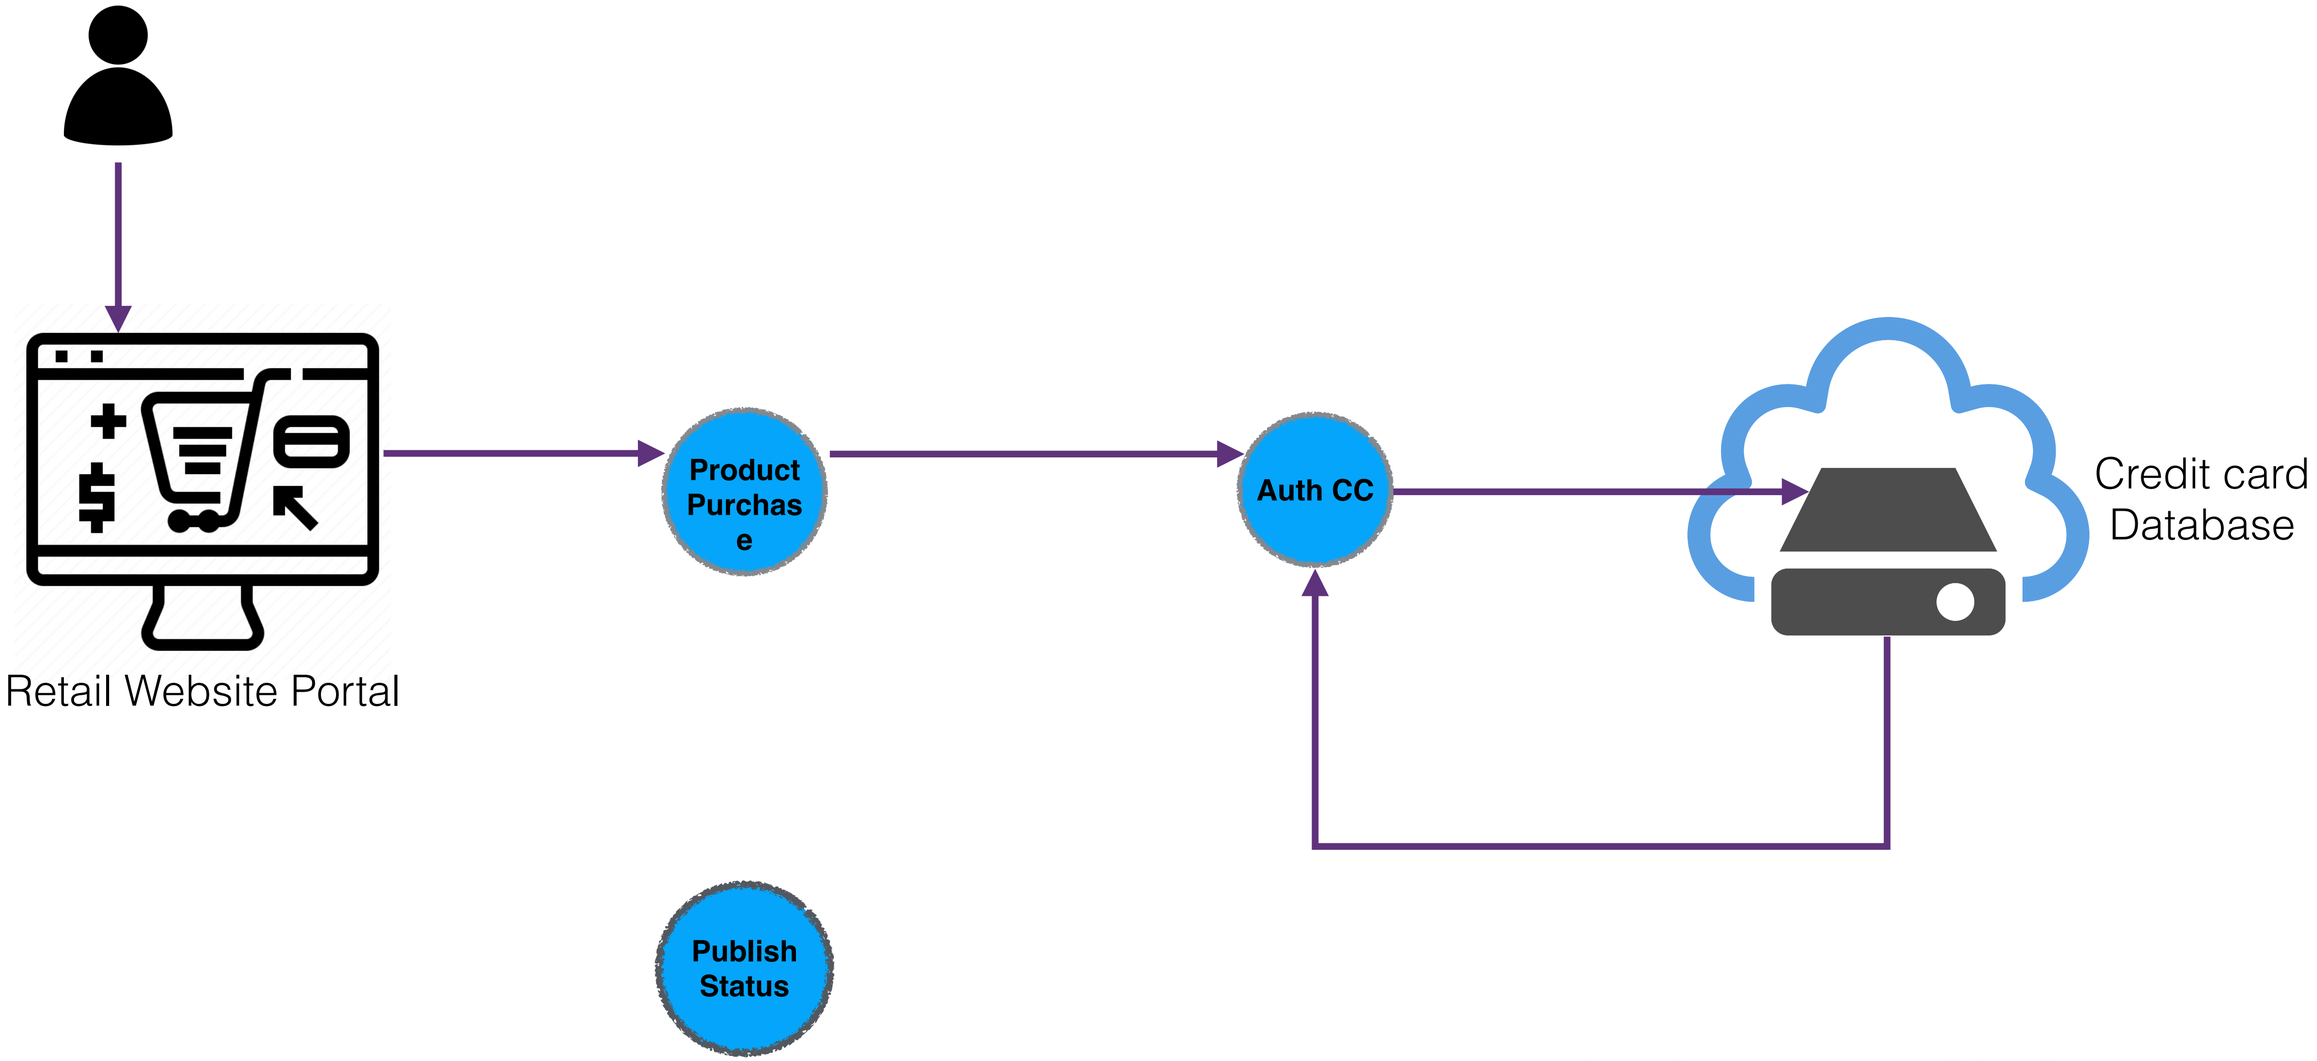

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

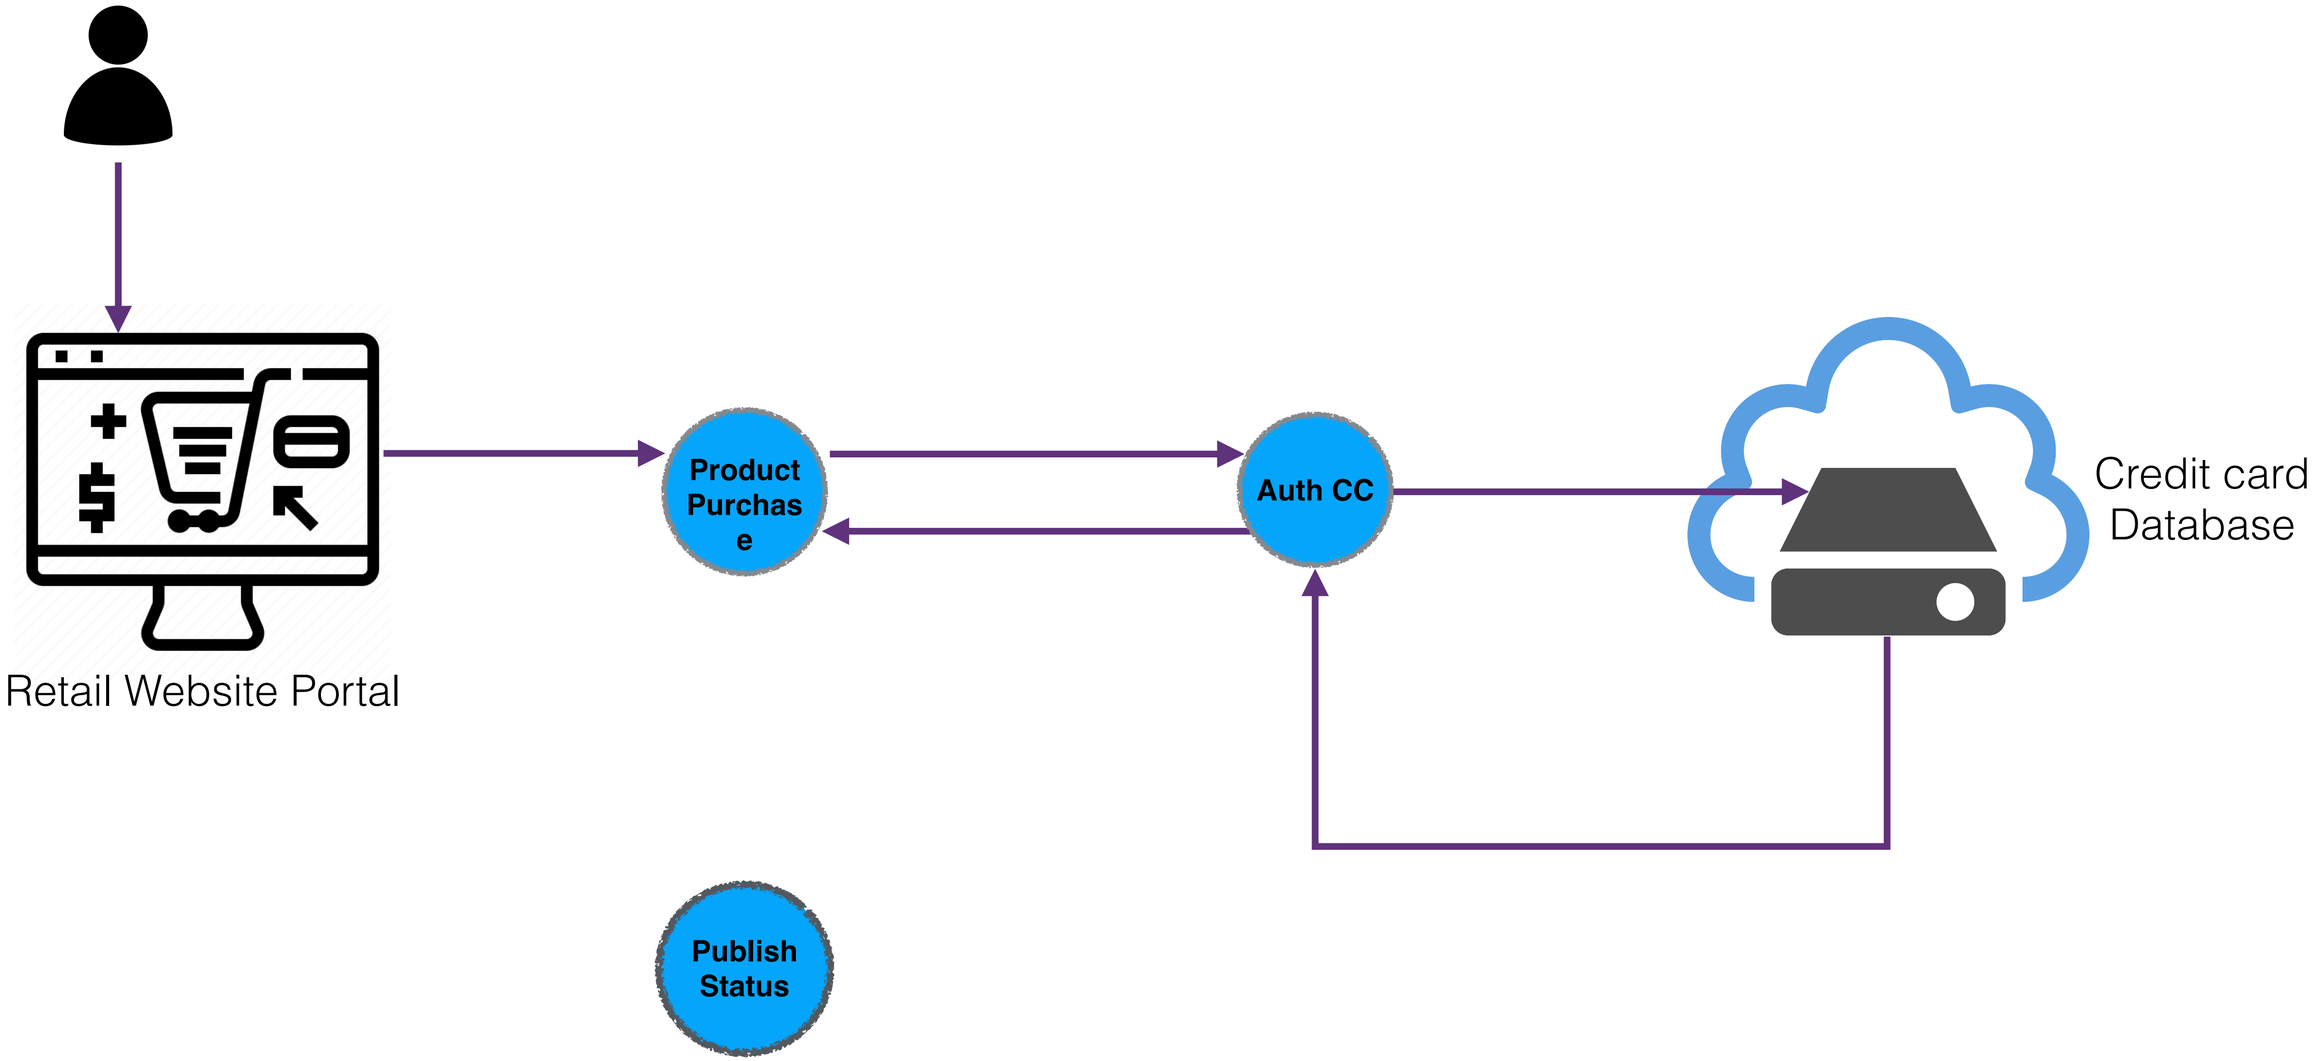

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

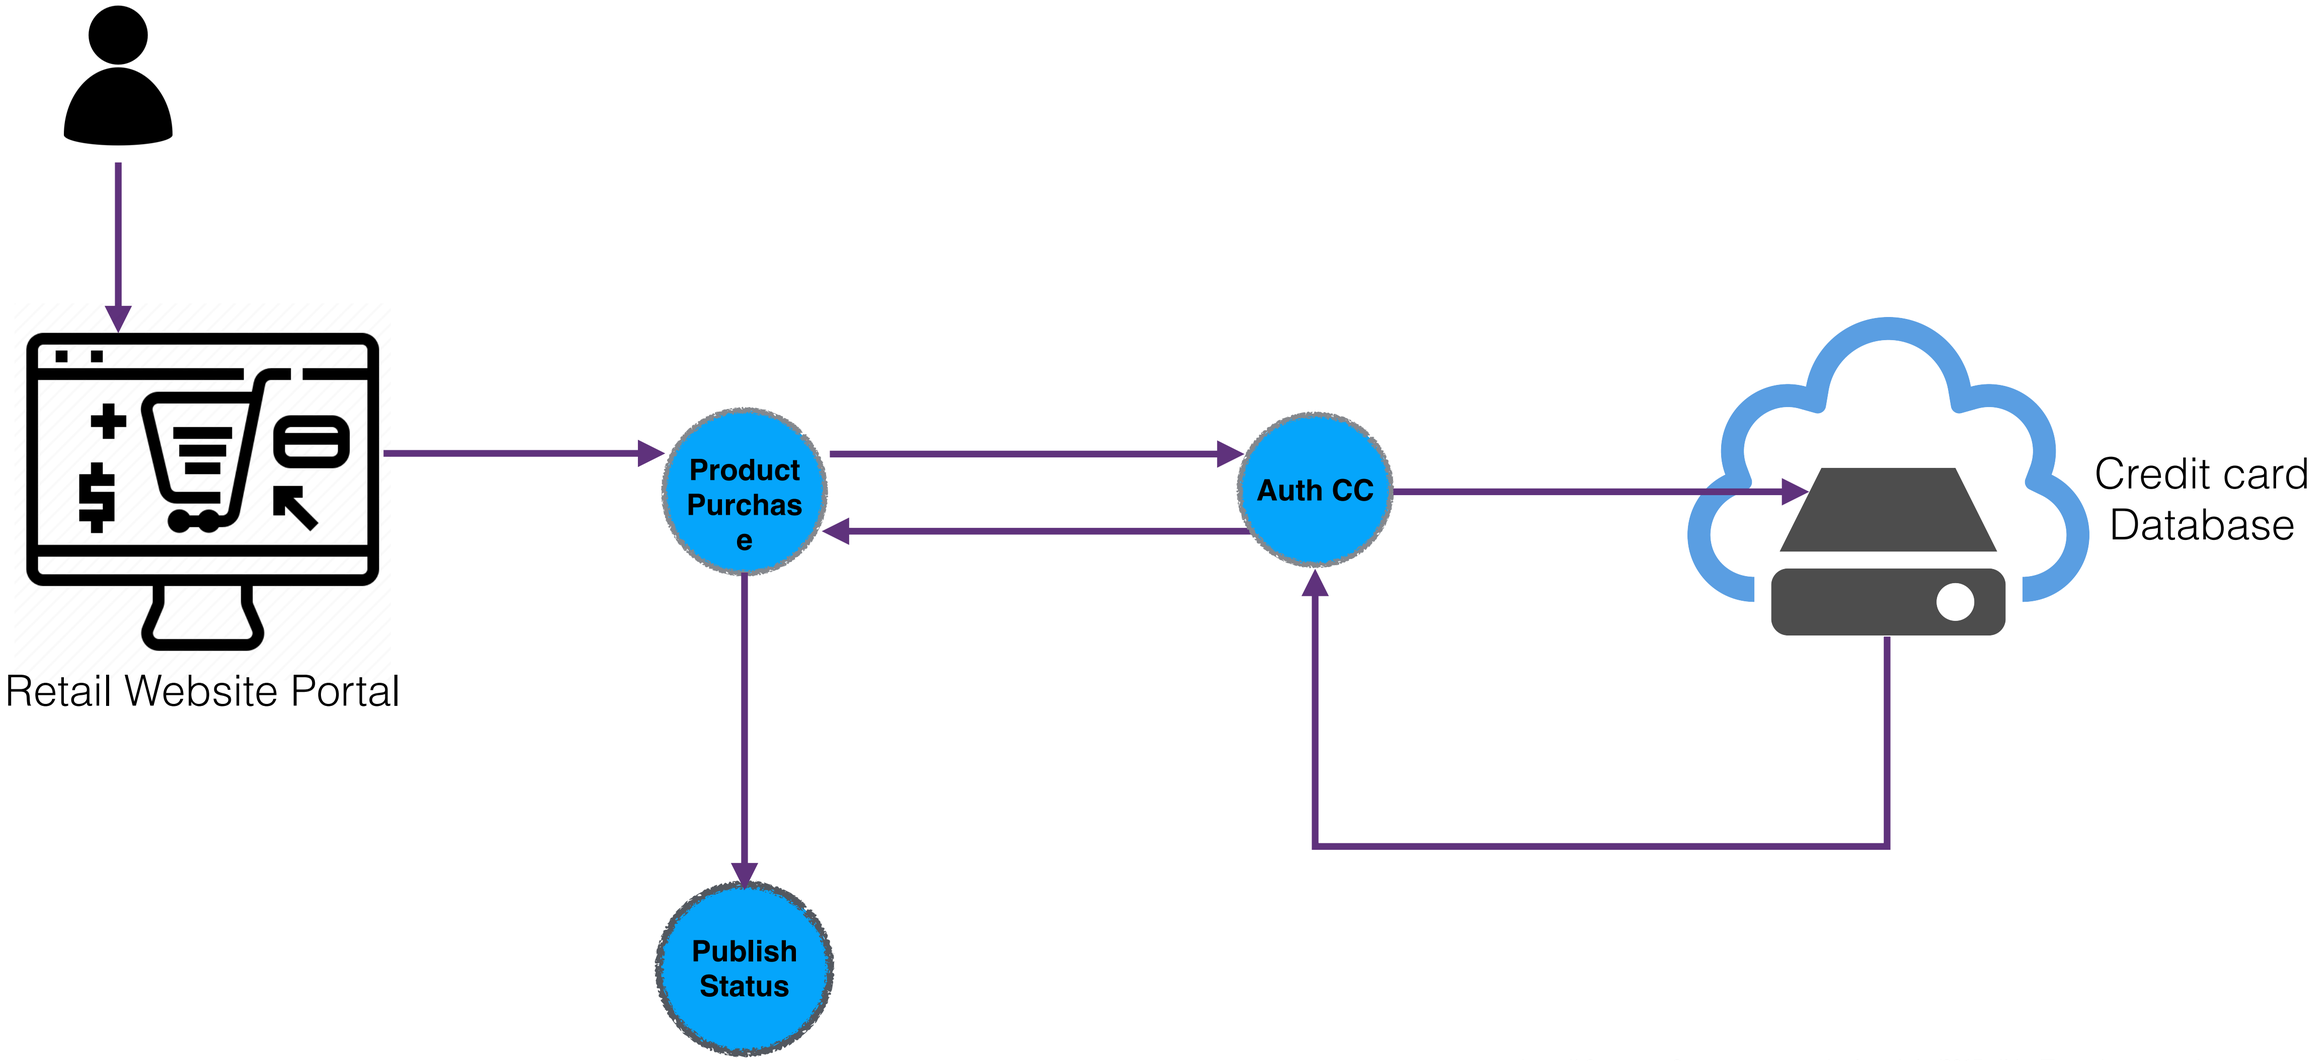

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

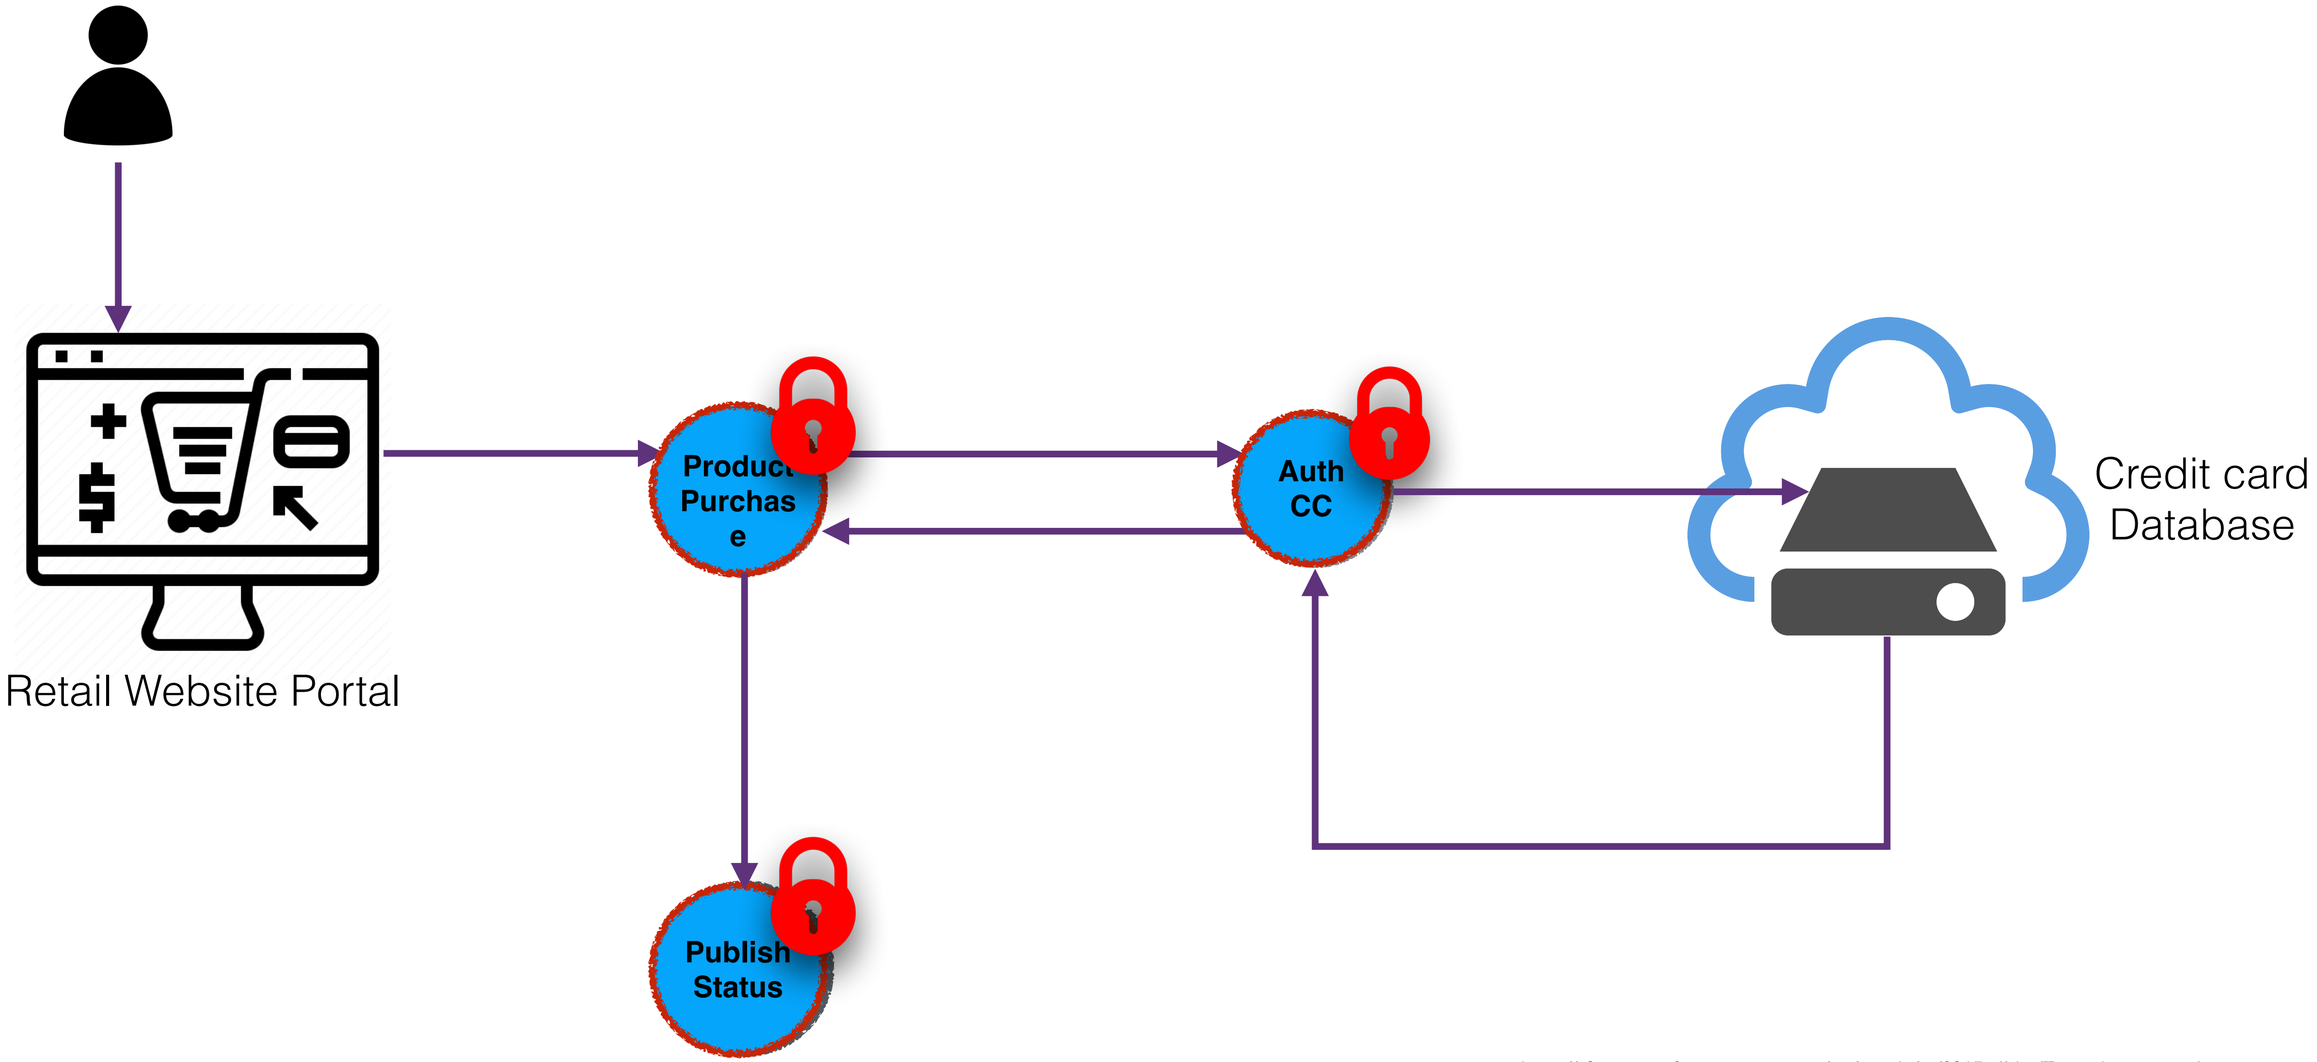

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

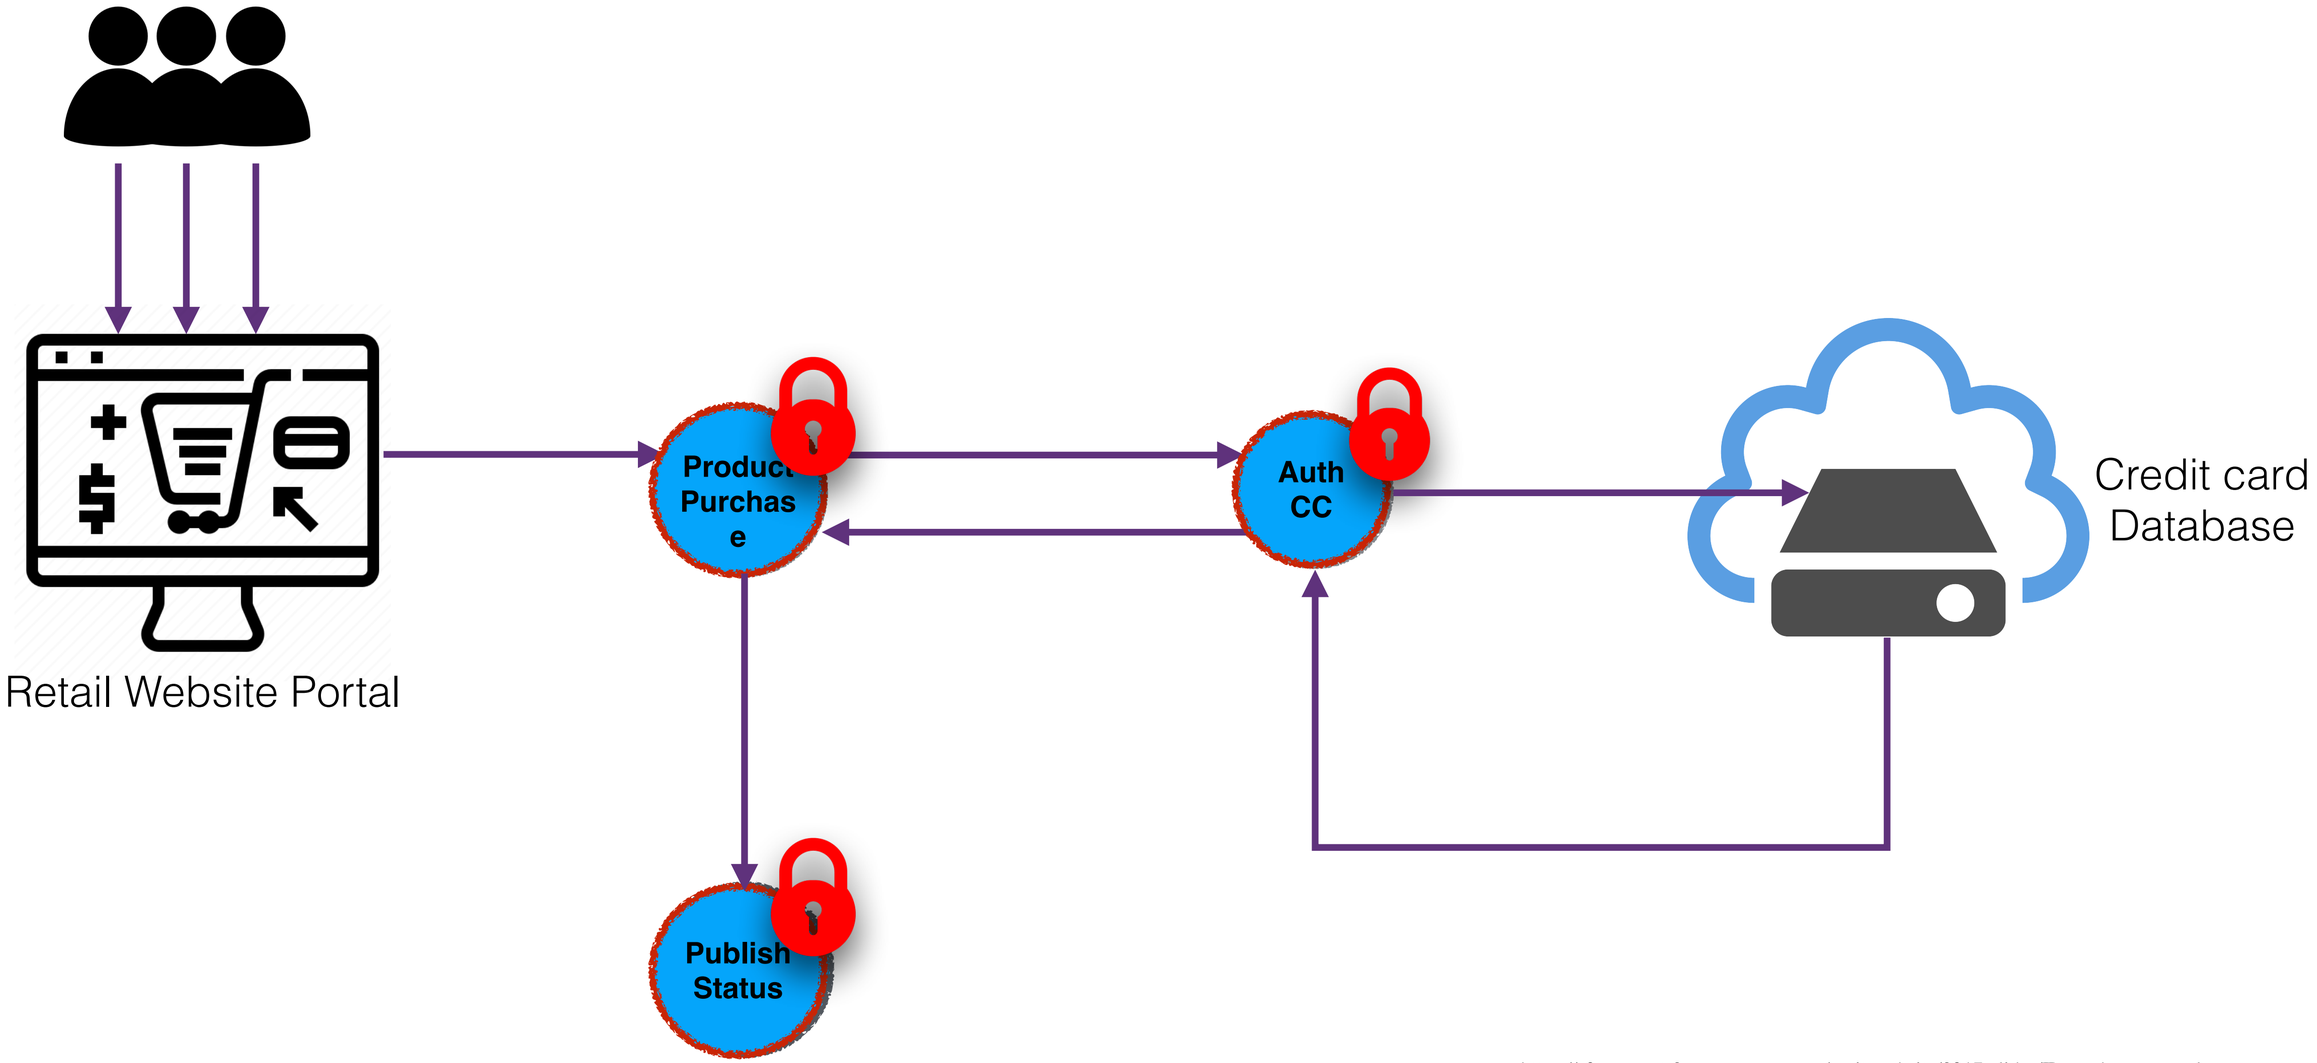

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

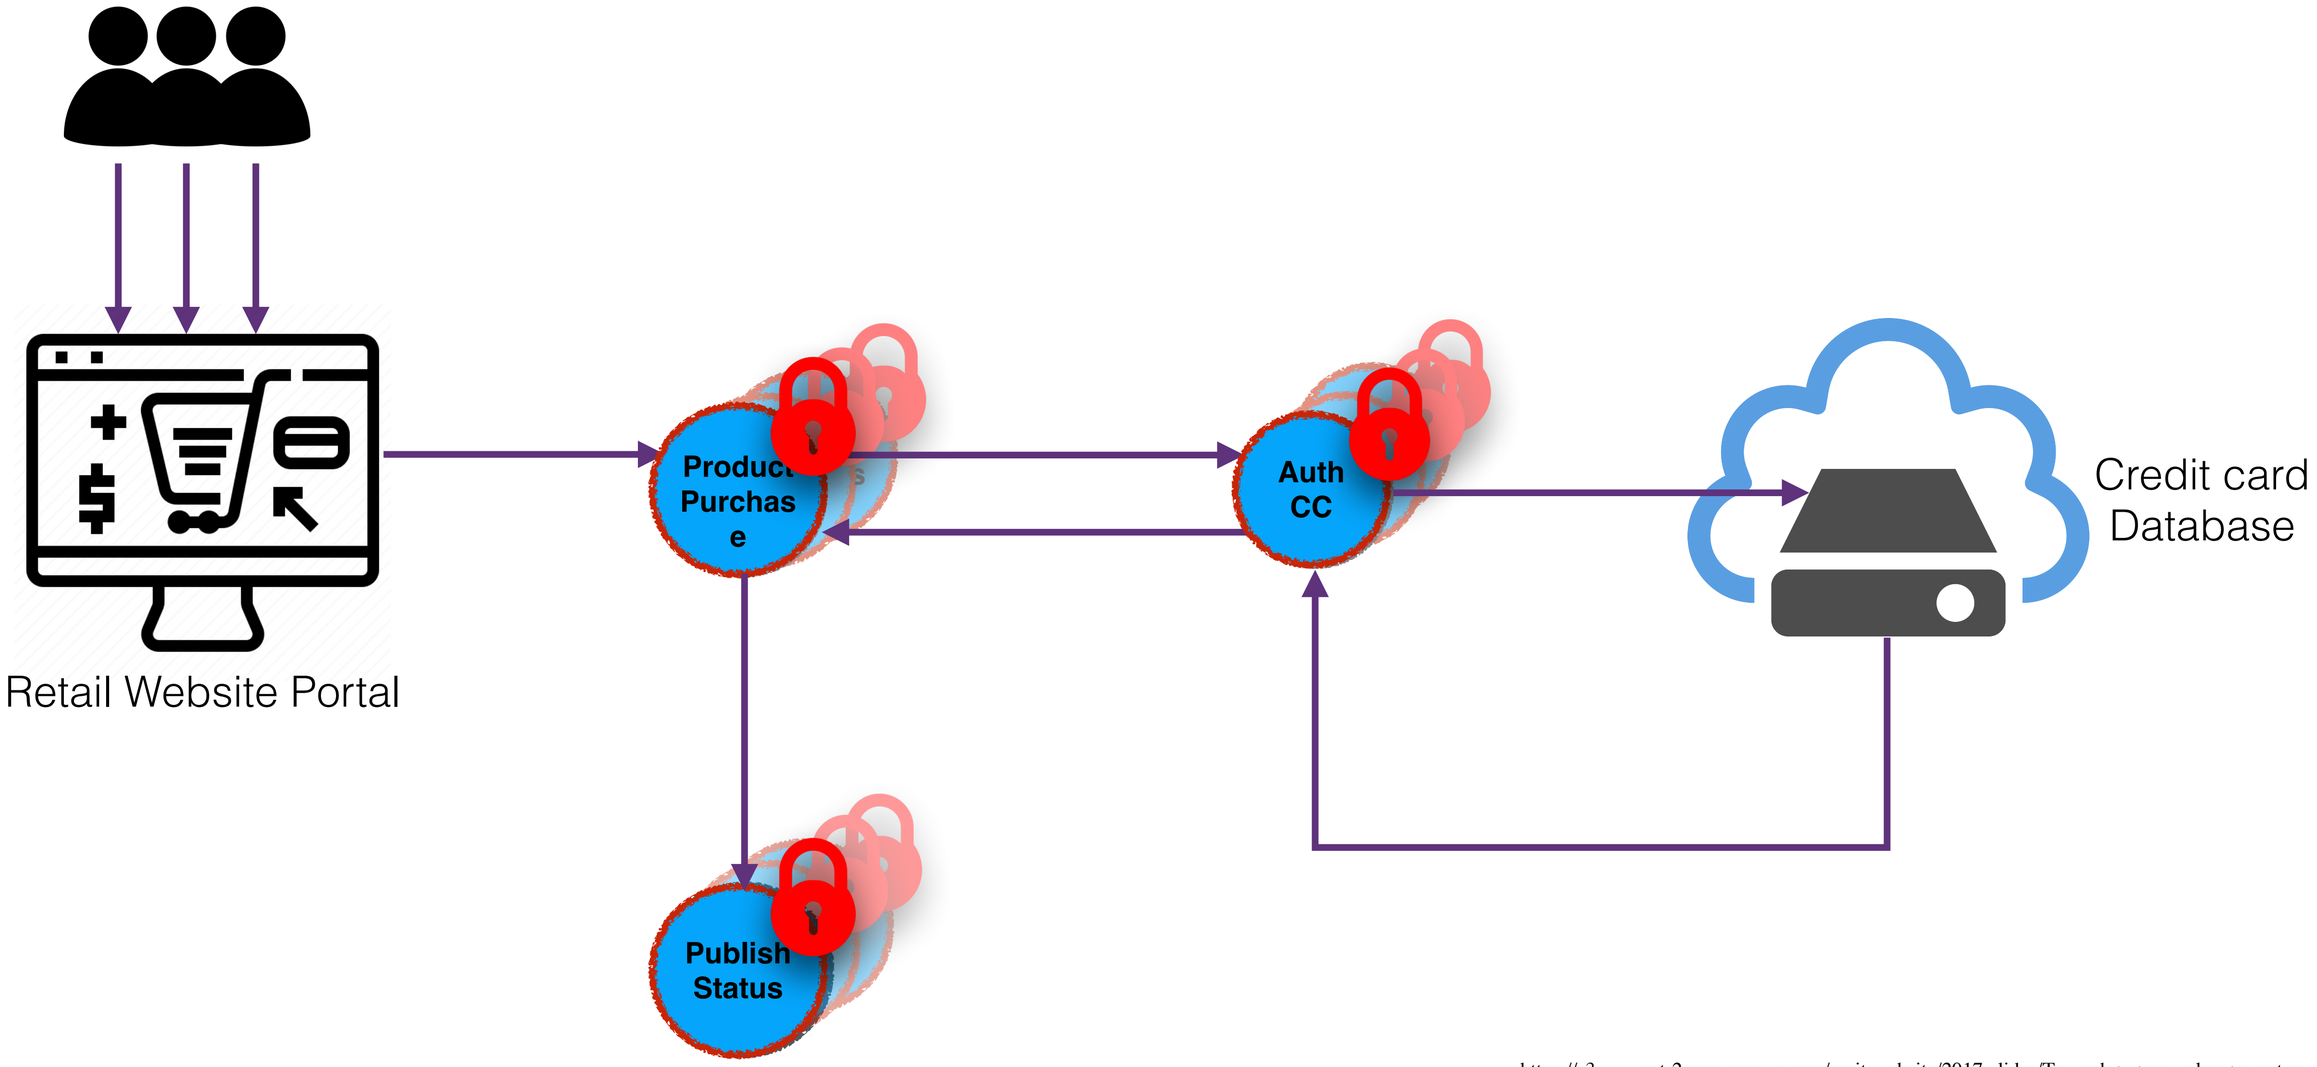

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

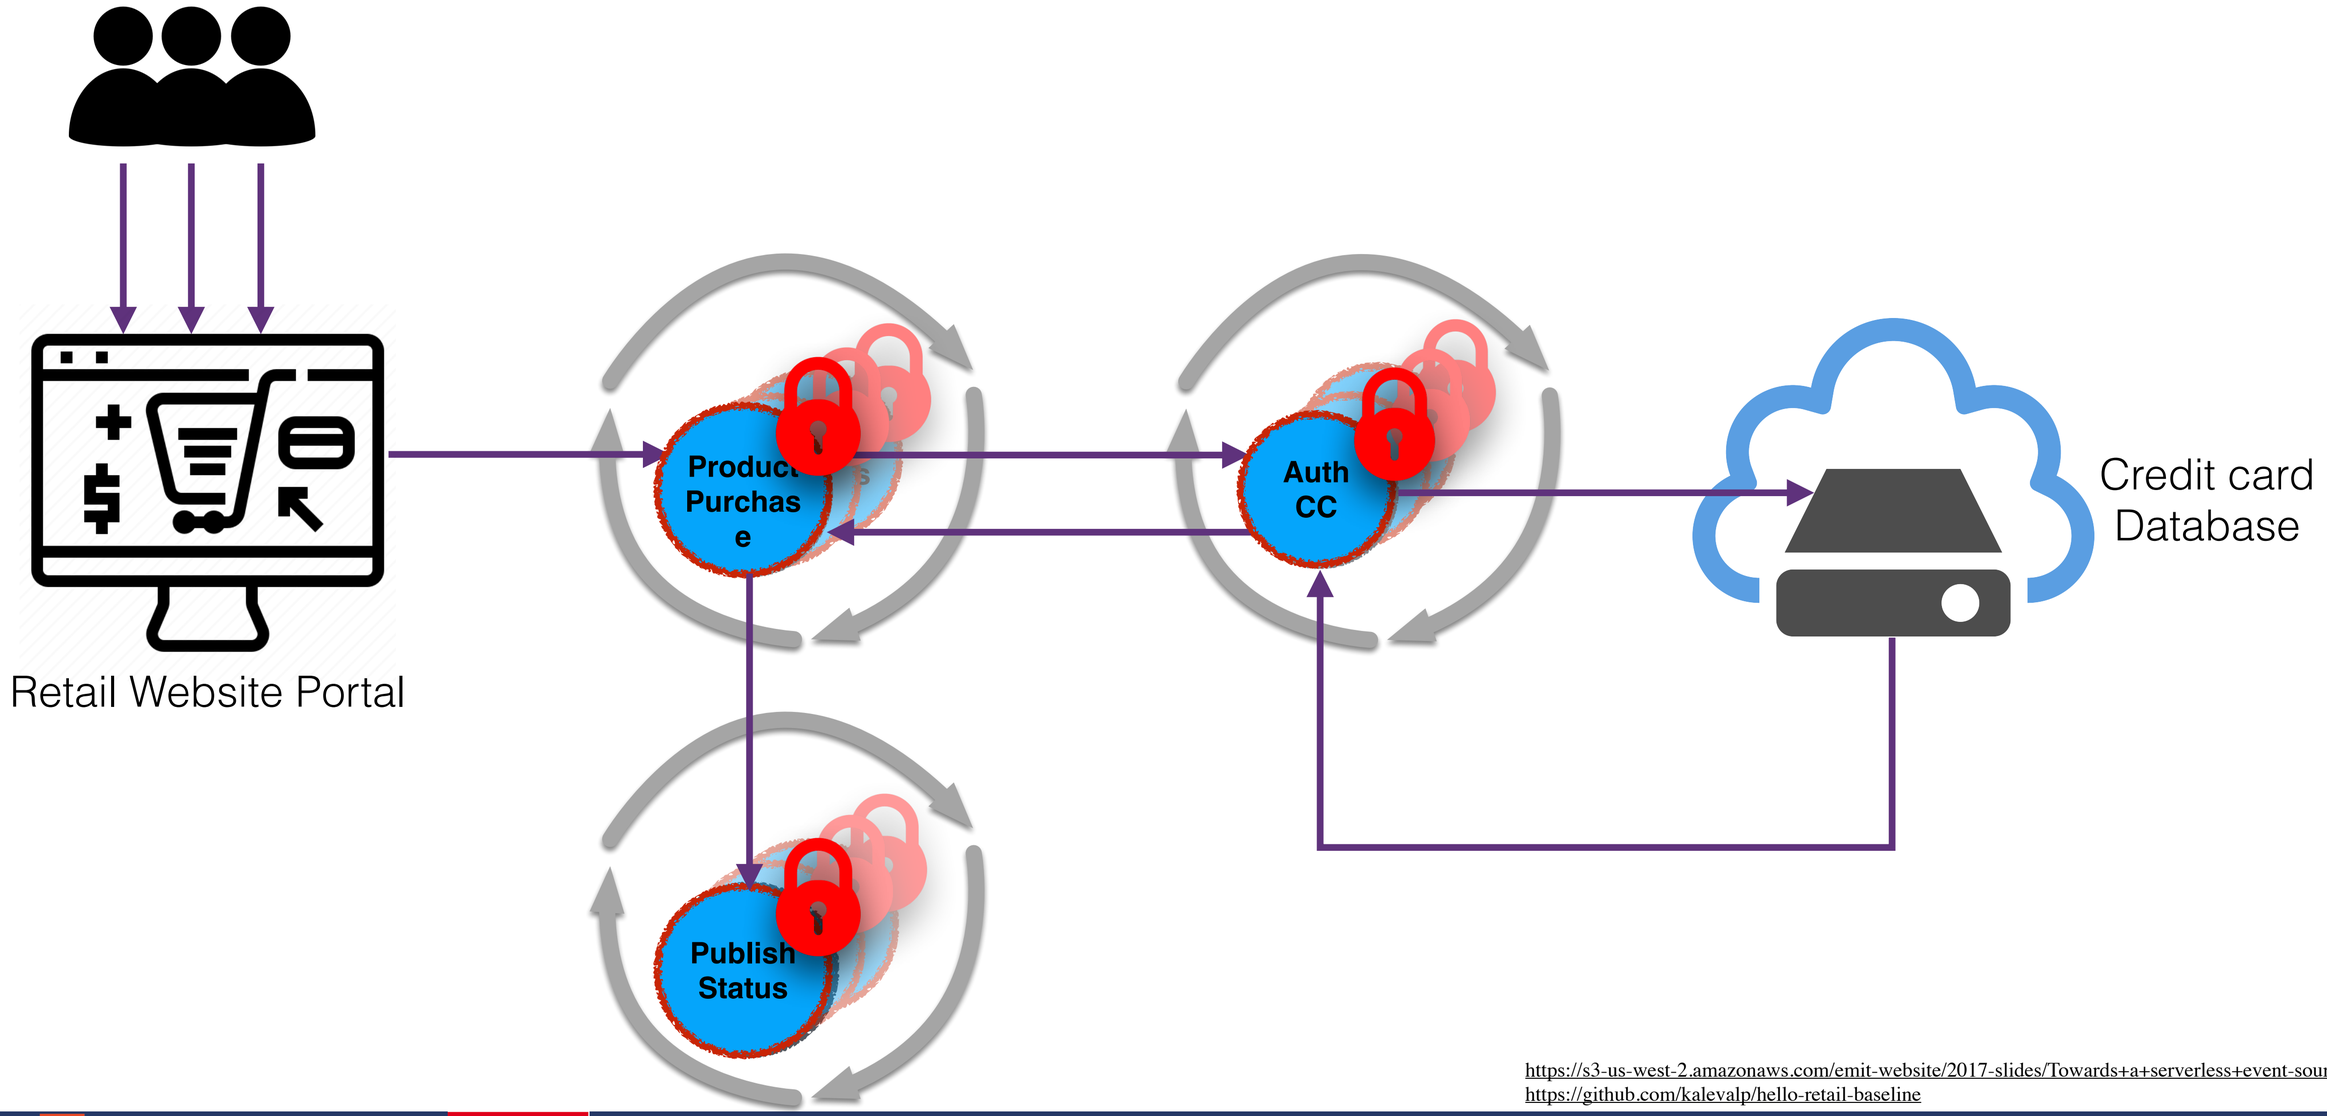

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

- IAM is the first line of defense in most cloud environments
- 99% of cloud identities are overly permissive
- 65% of the known cloud security incidents were due to misconfigurations

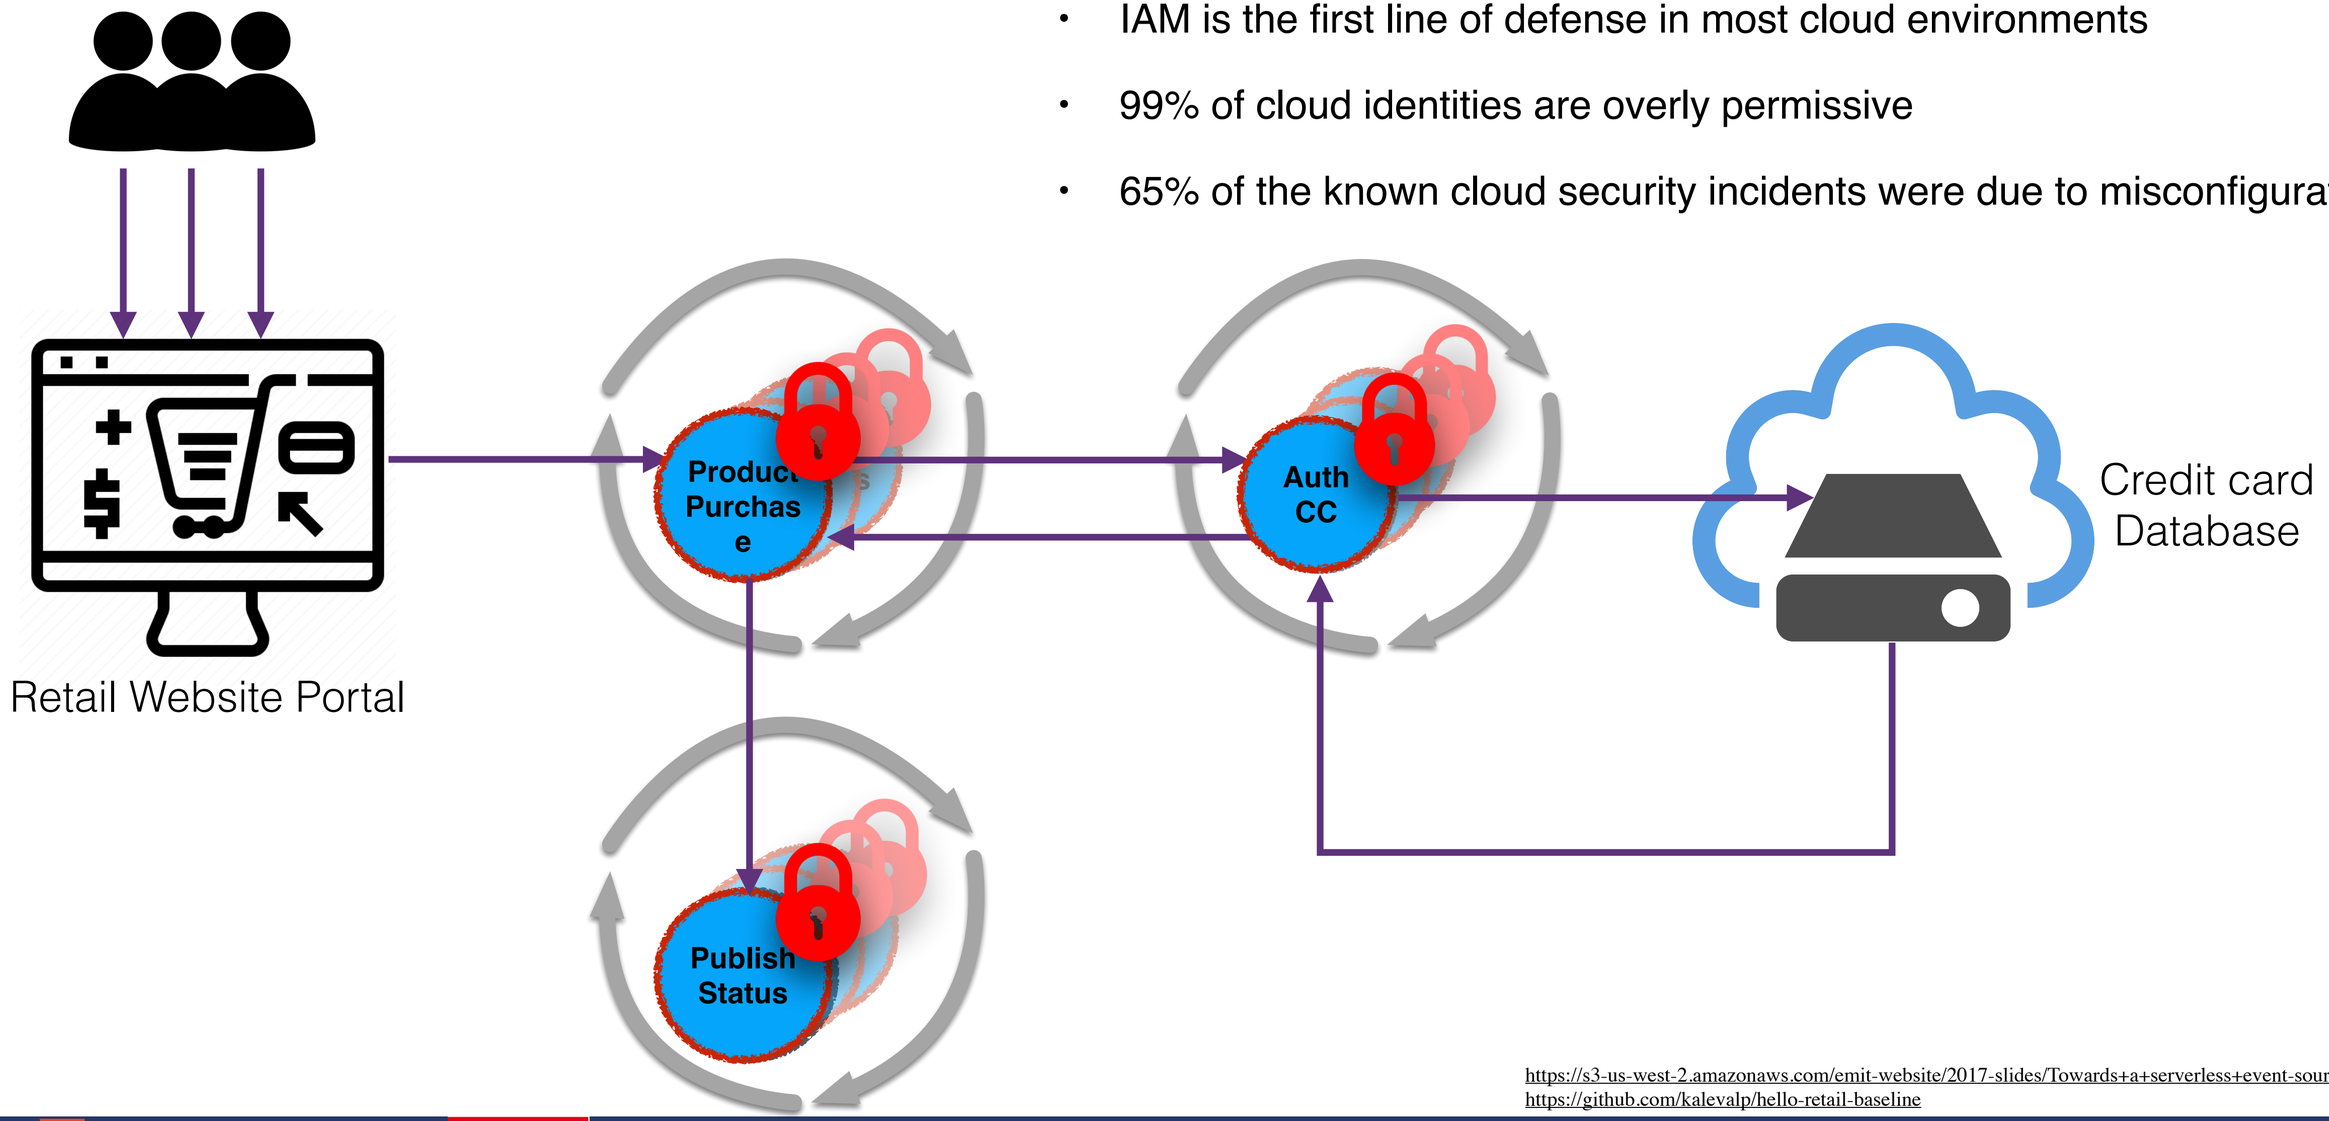

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

- IAM is the first line of defense in most cloud environments
- 99% of cloud identities are overly permissive
- 65% of the known cloud security incidents were due to misconfigurations

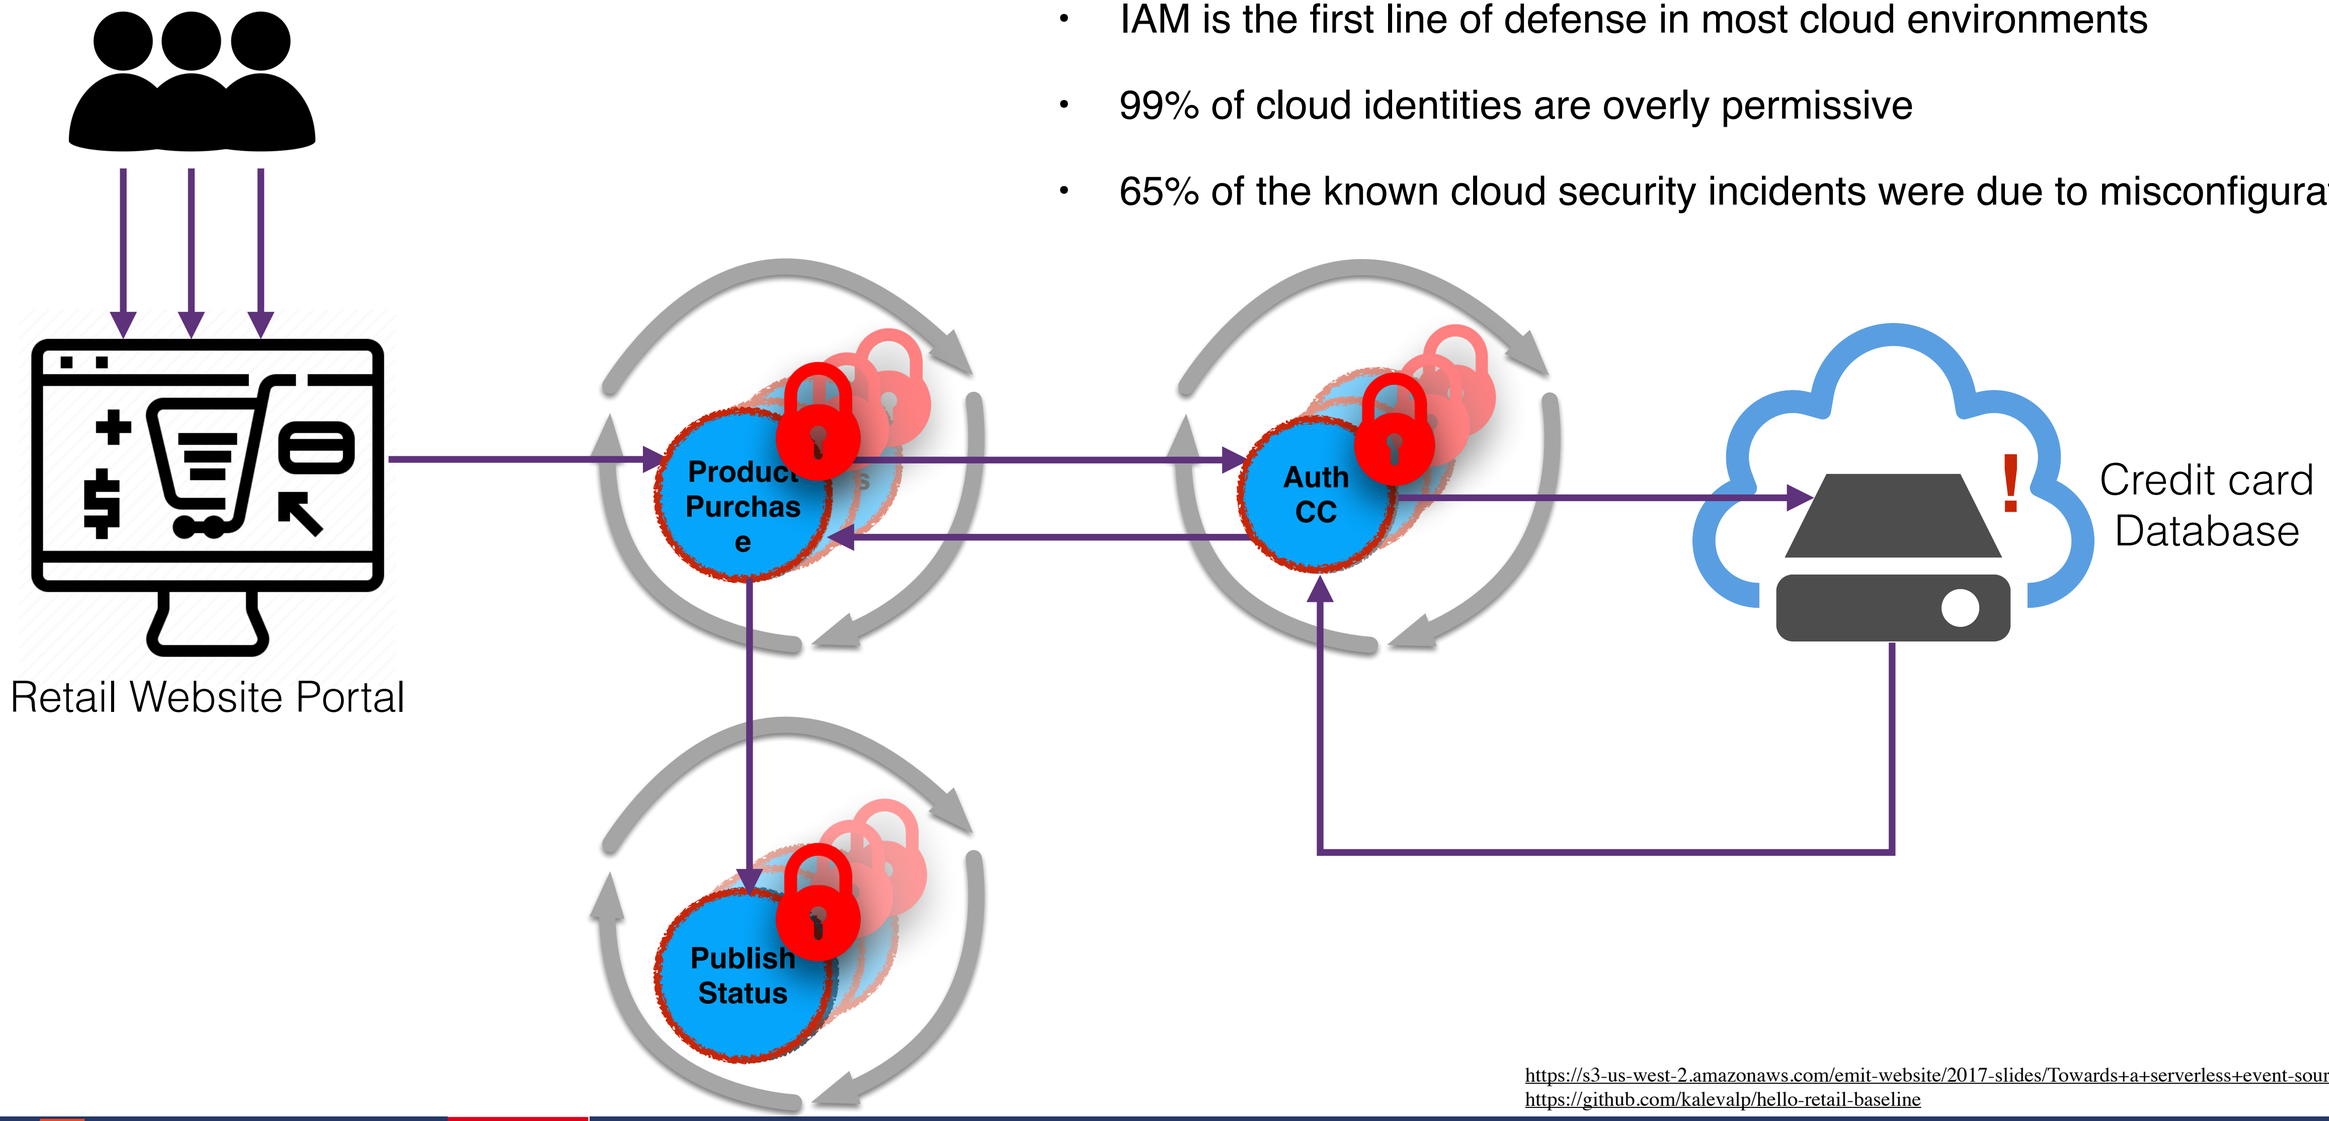

https://s3-us-west-2.amazonaws.com/emit-website/2017-slides/Towards+a+serverless+event-sourced+Nordstrom.pdf
https://github.com/kalevalp/hello-retail-baseline

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

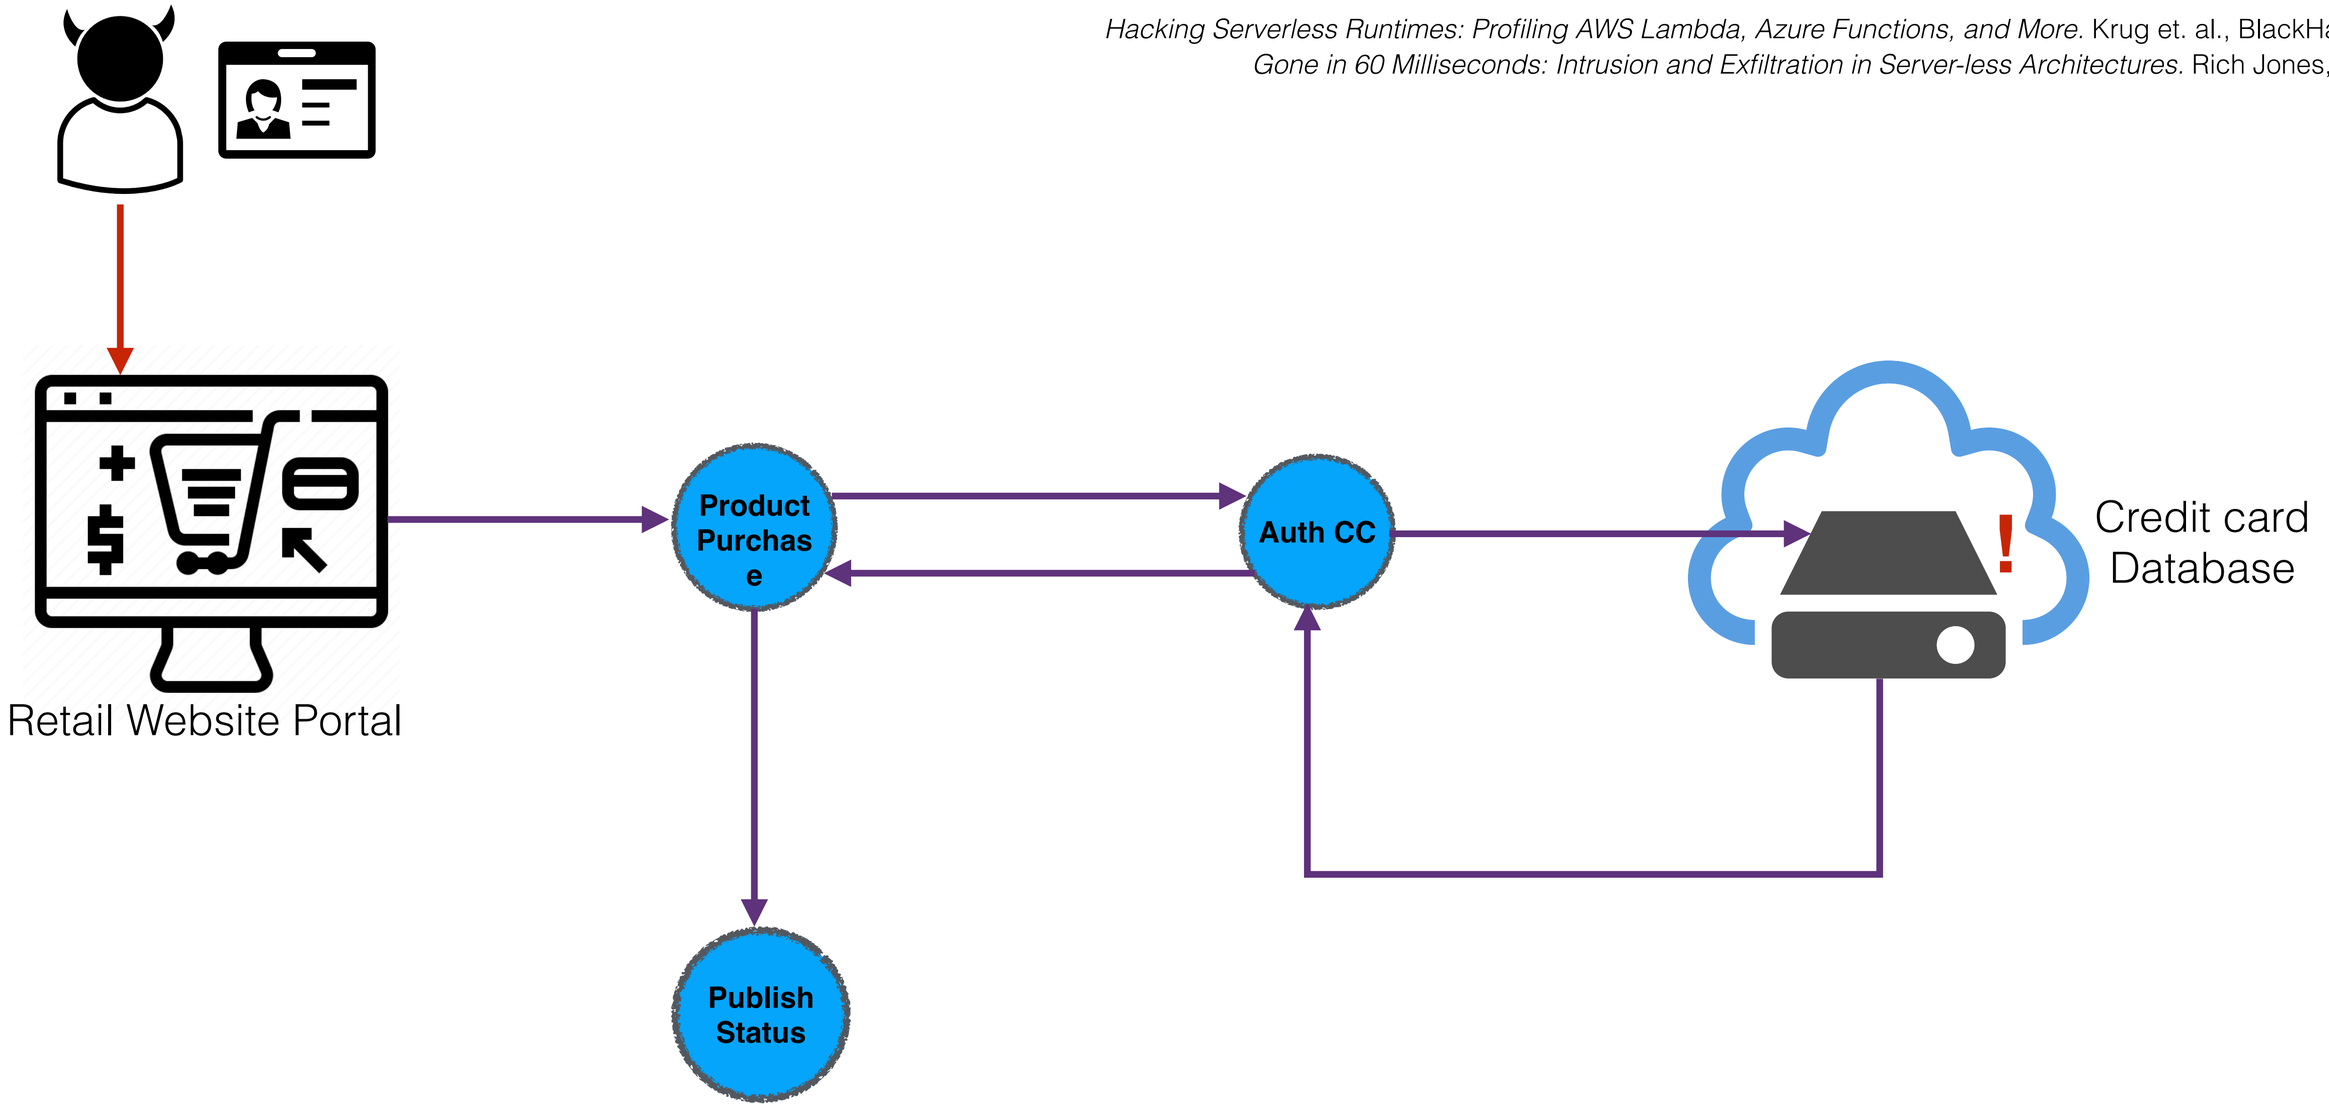

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

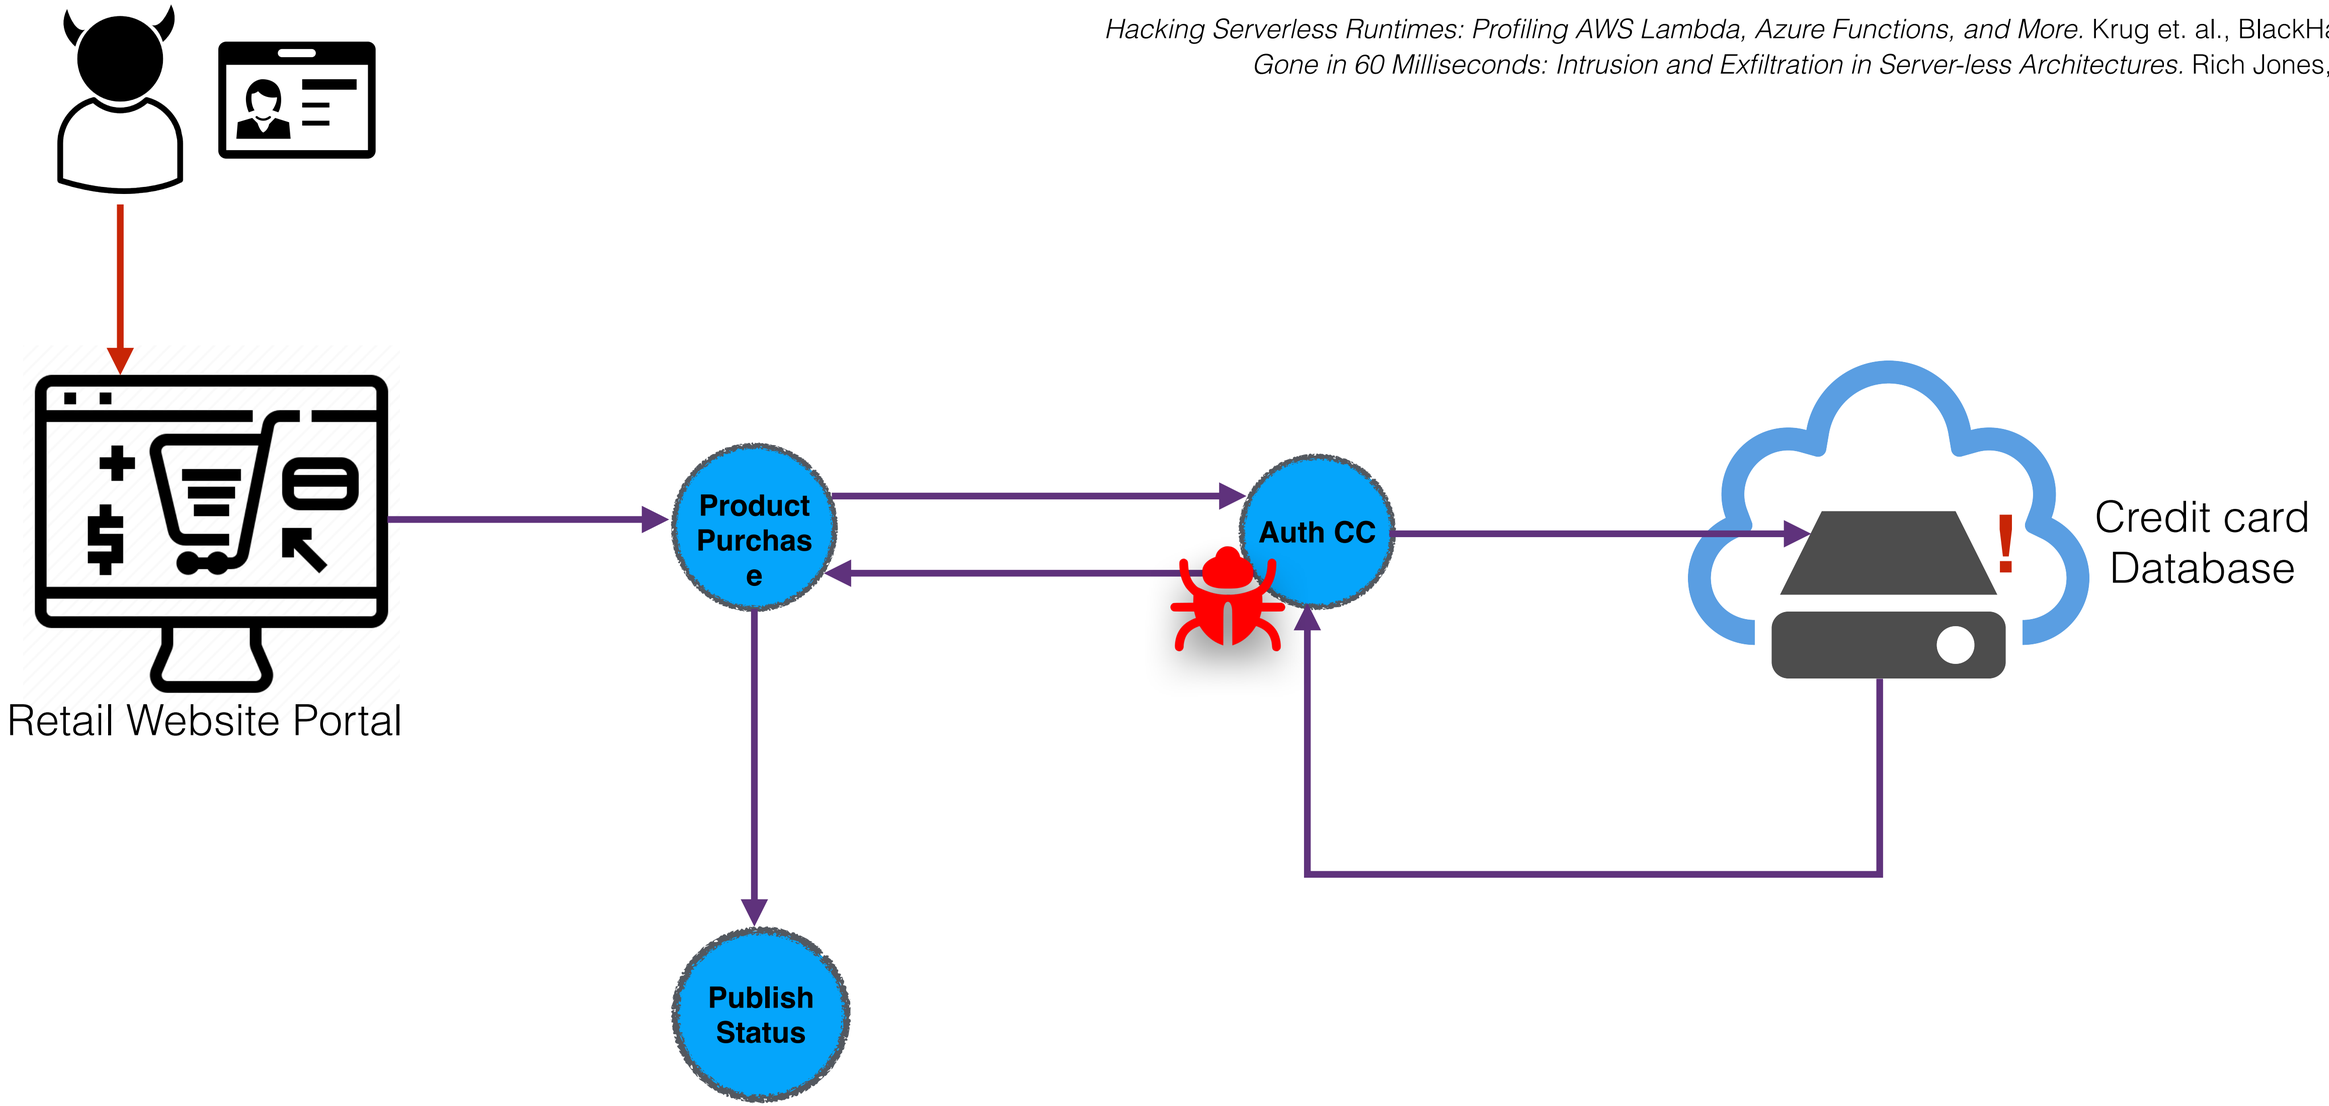

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

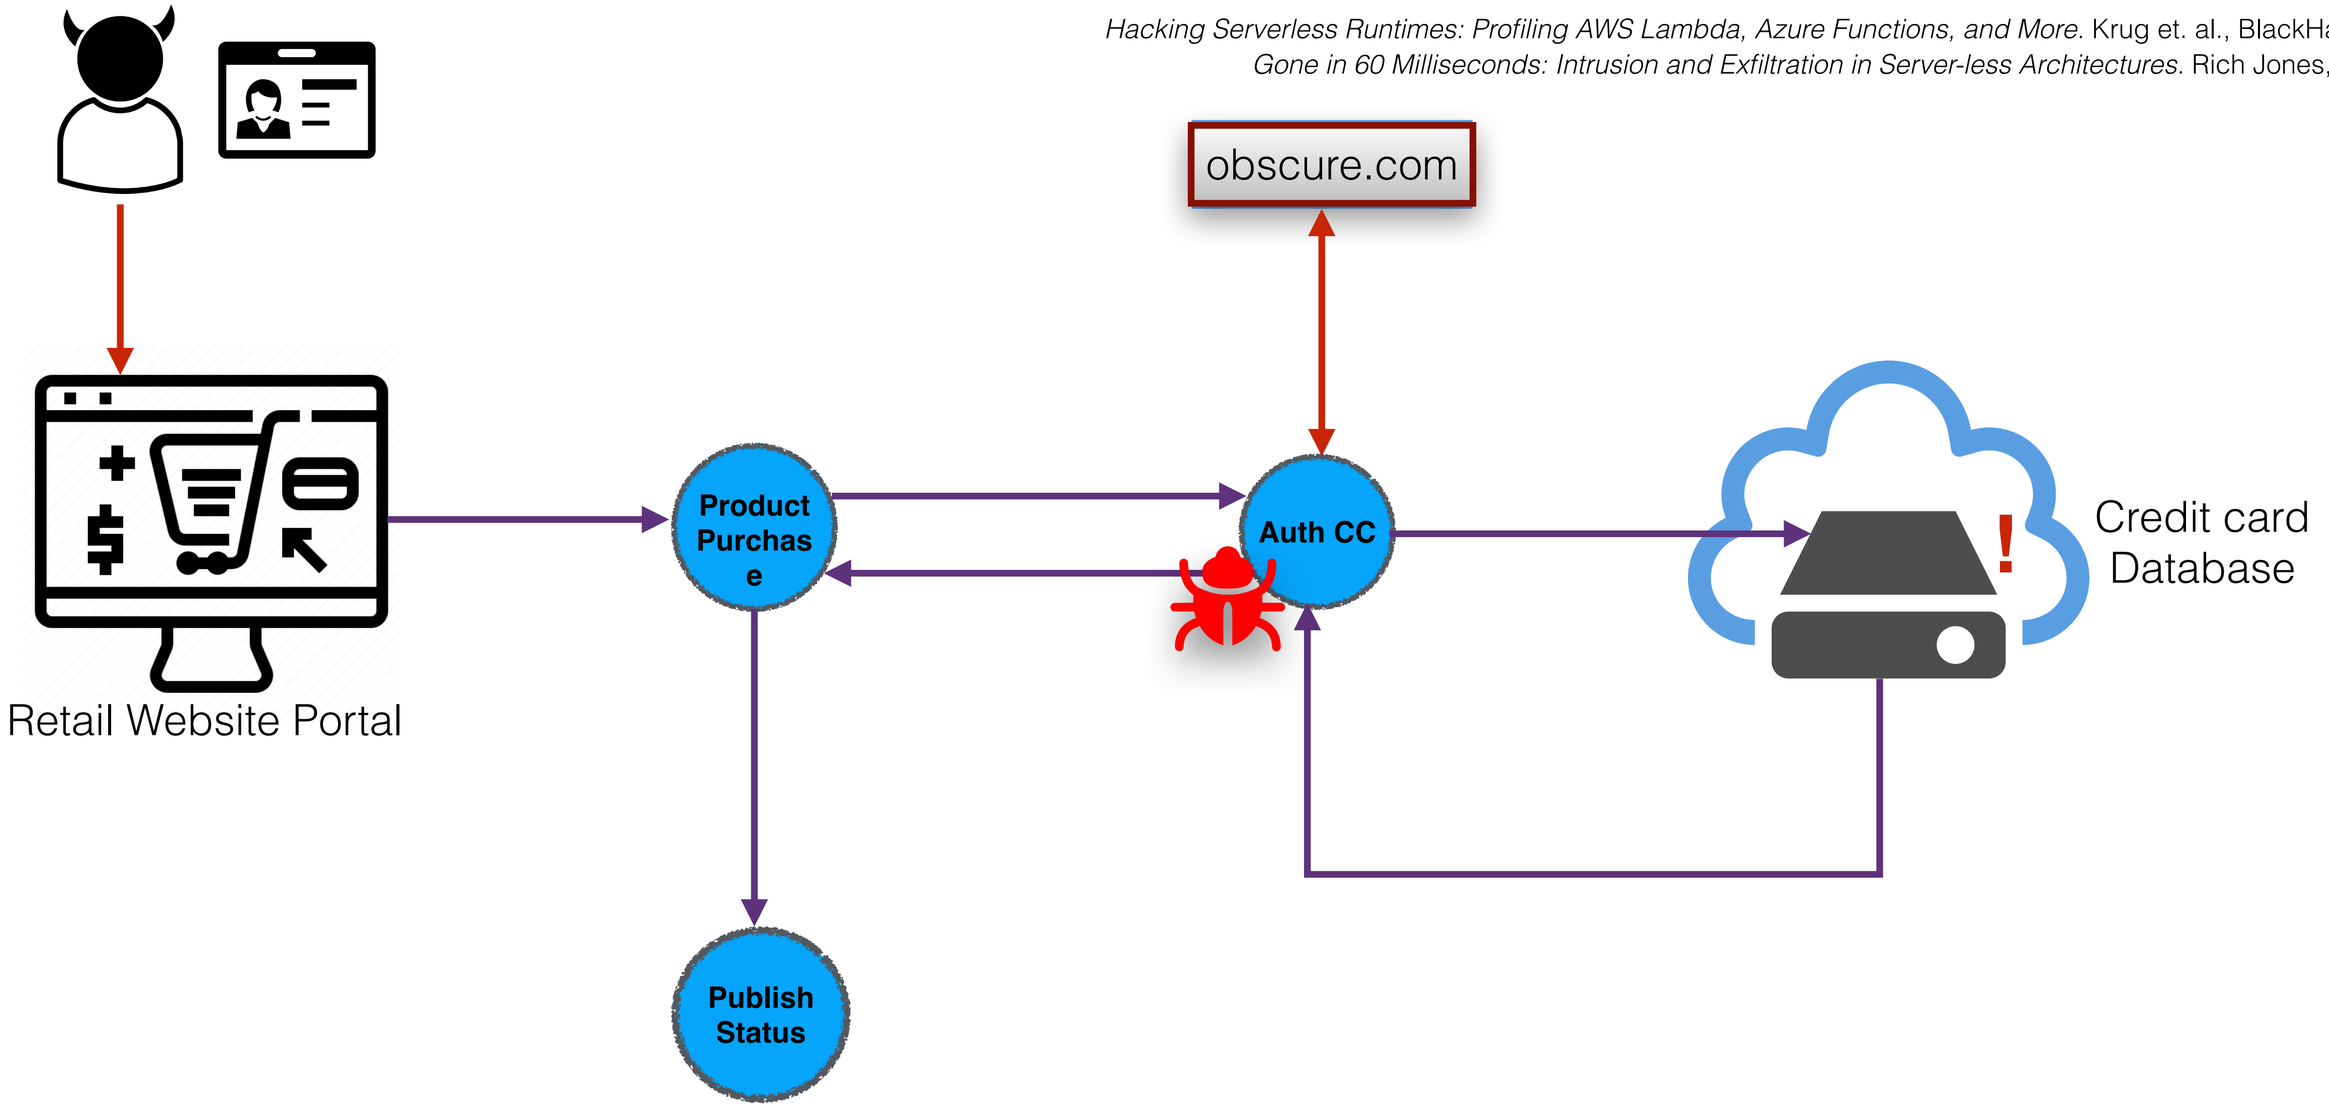

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

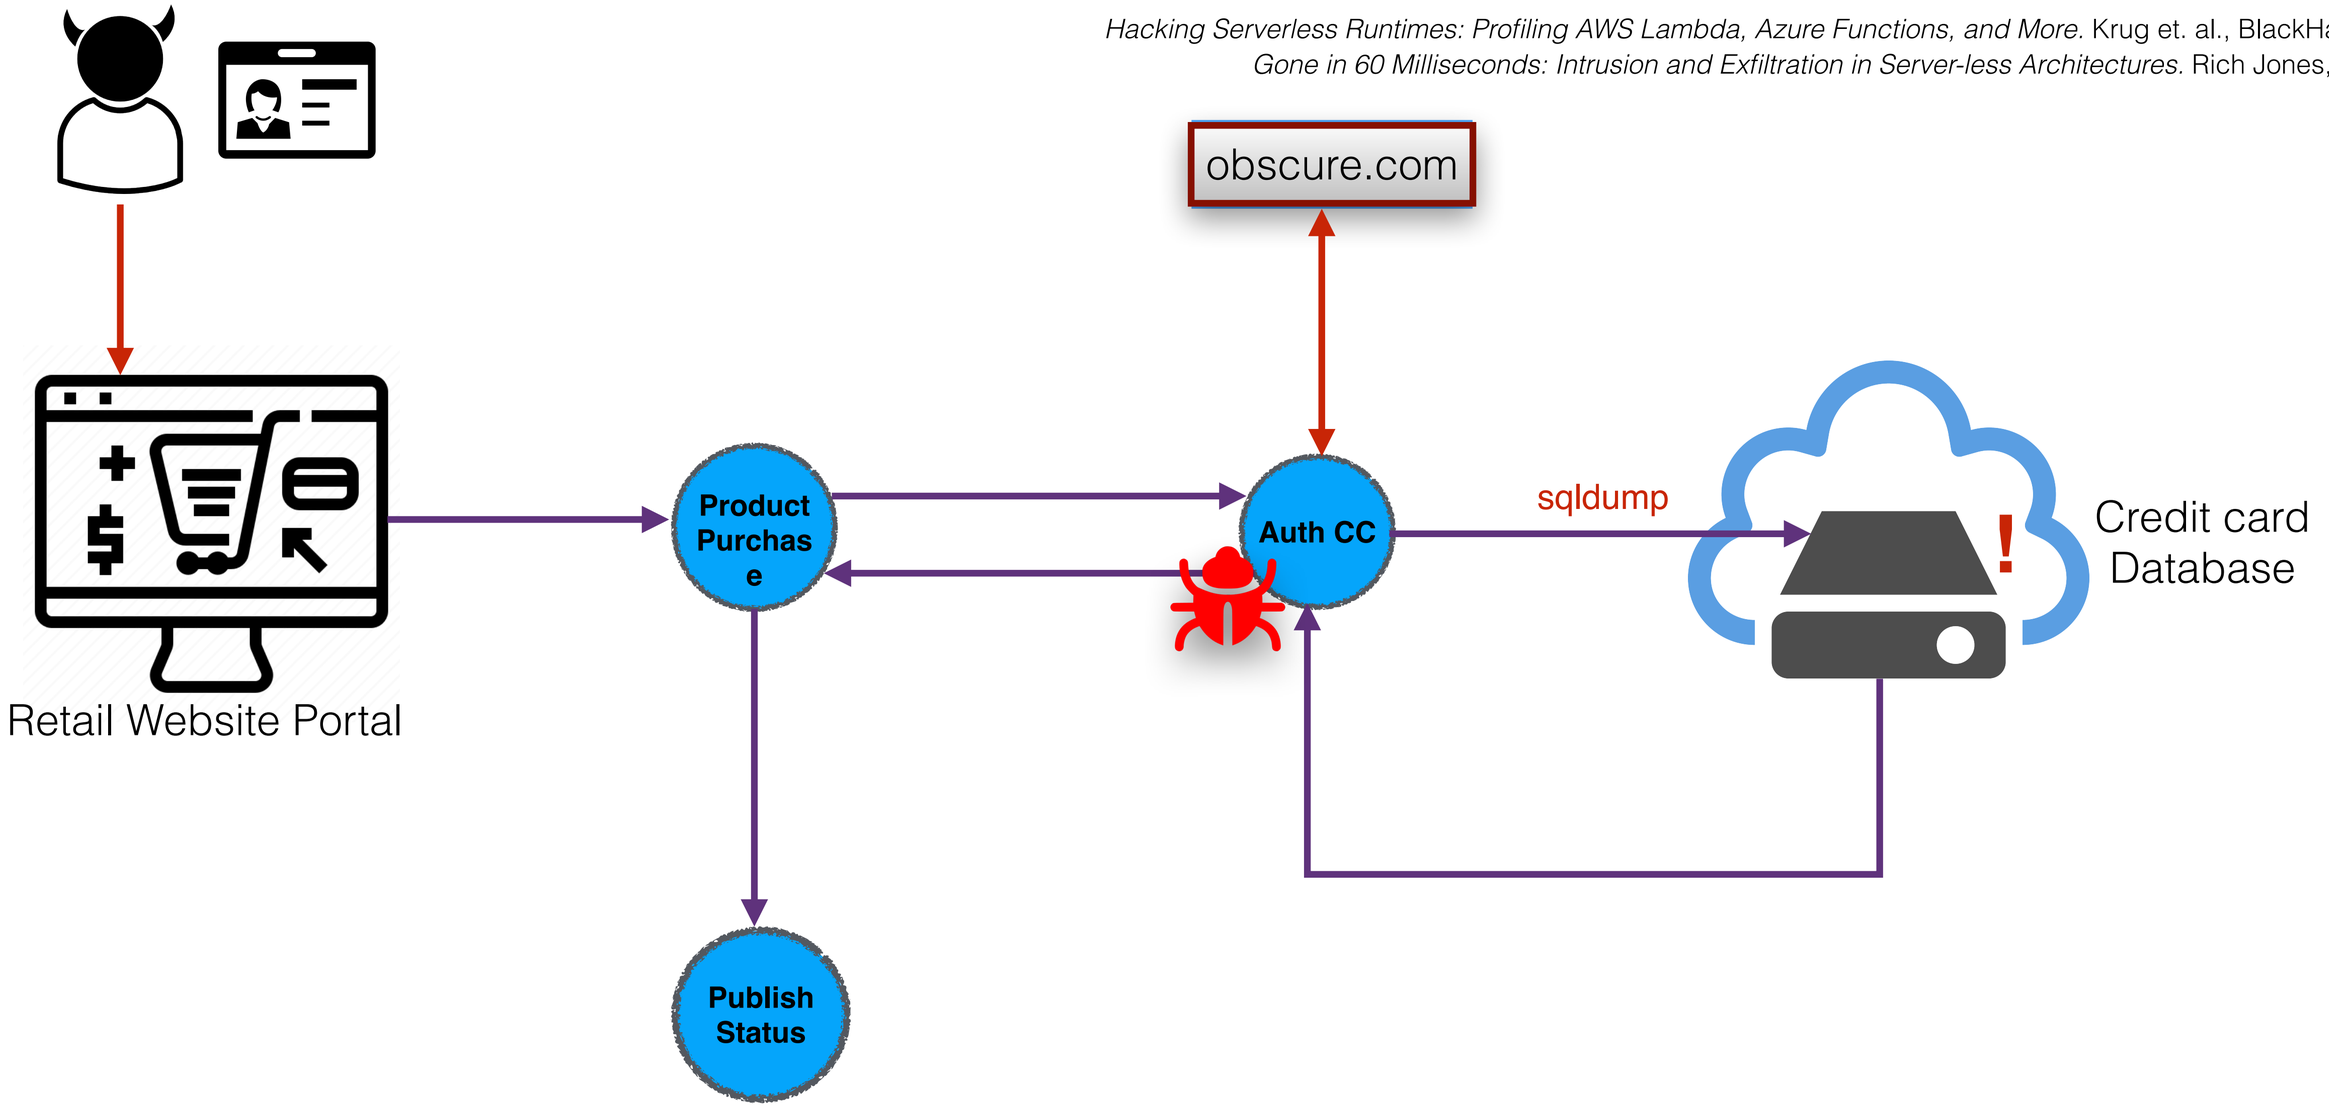

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

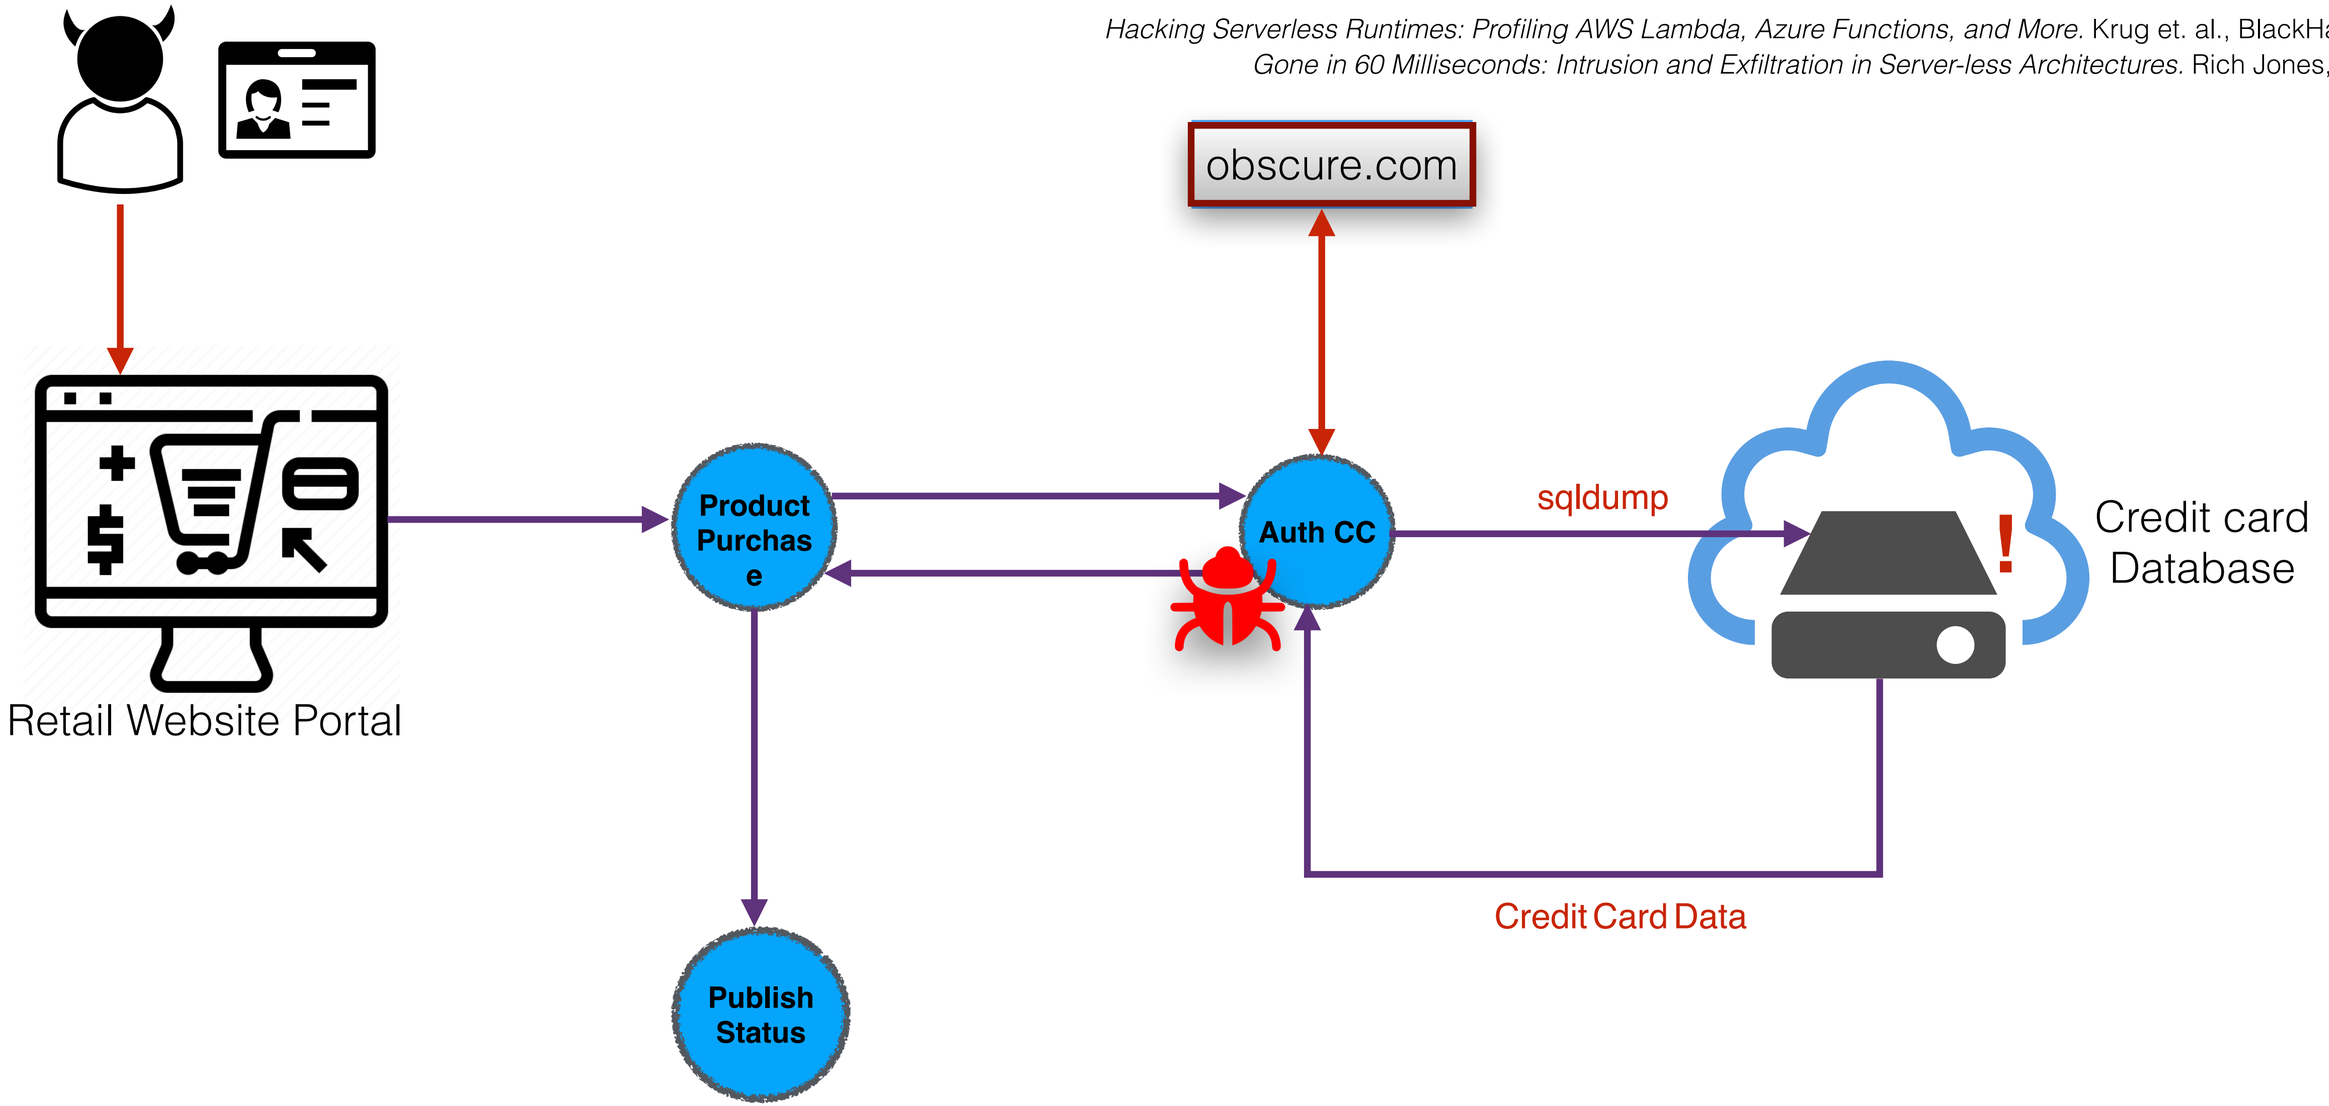

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

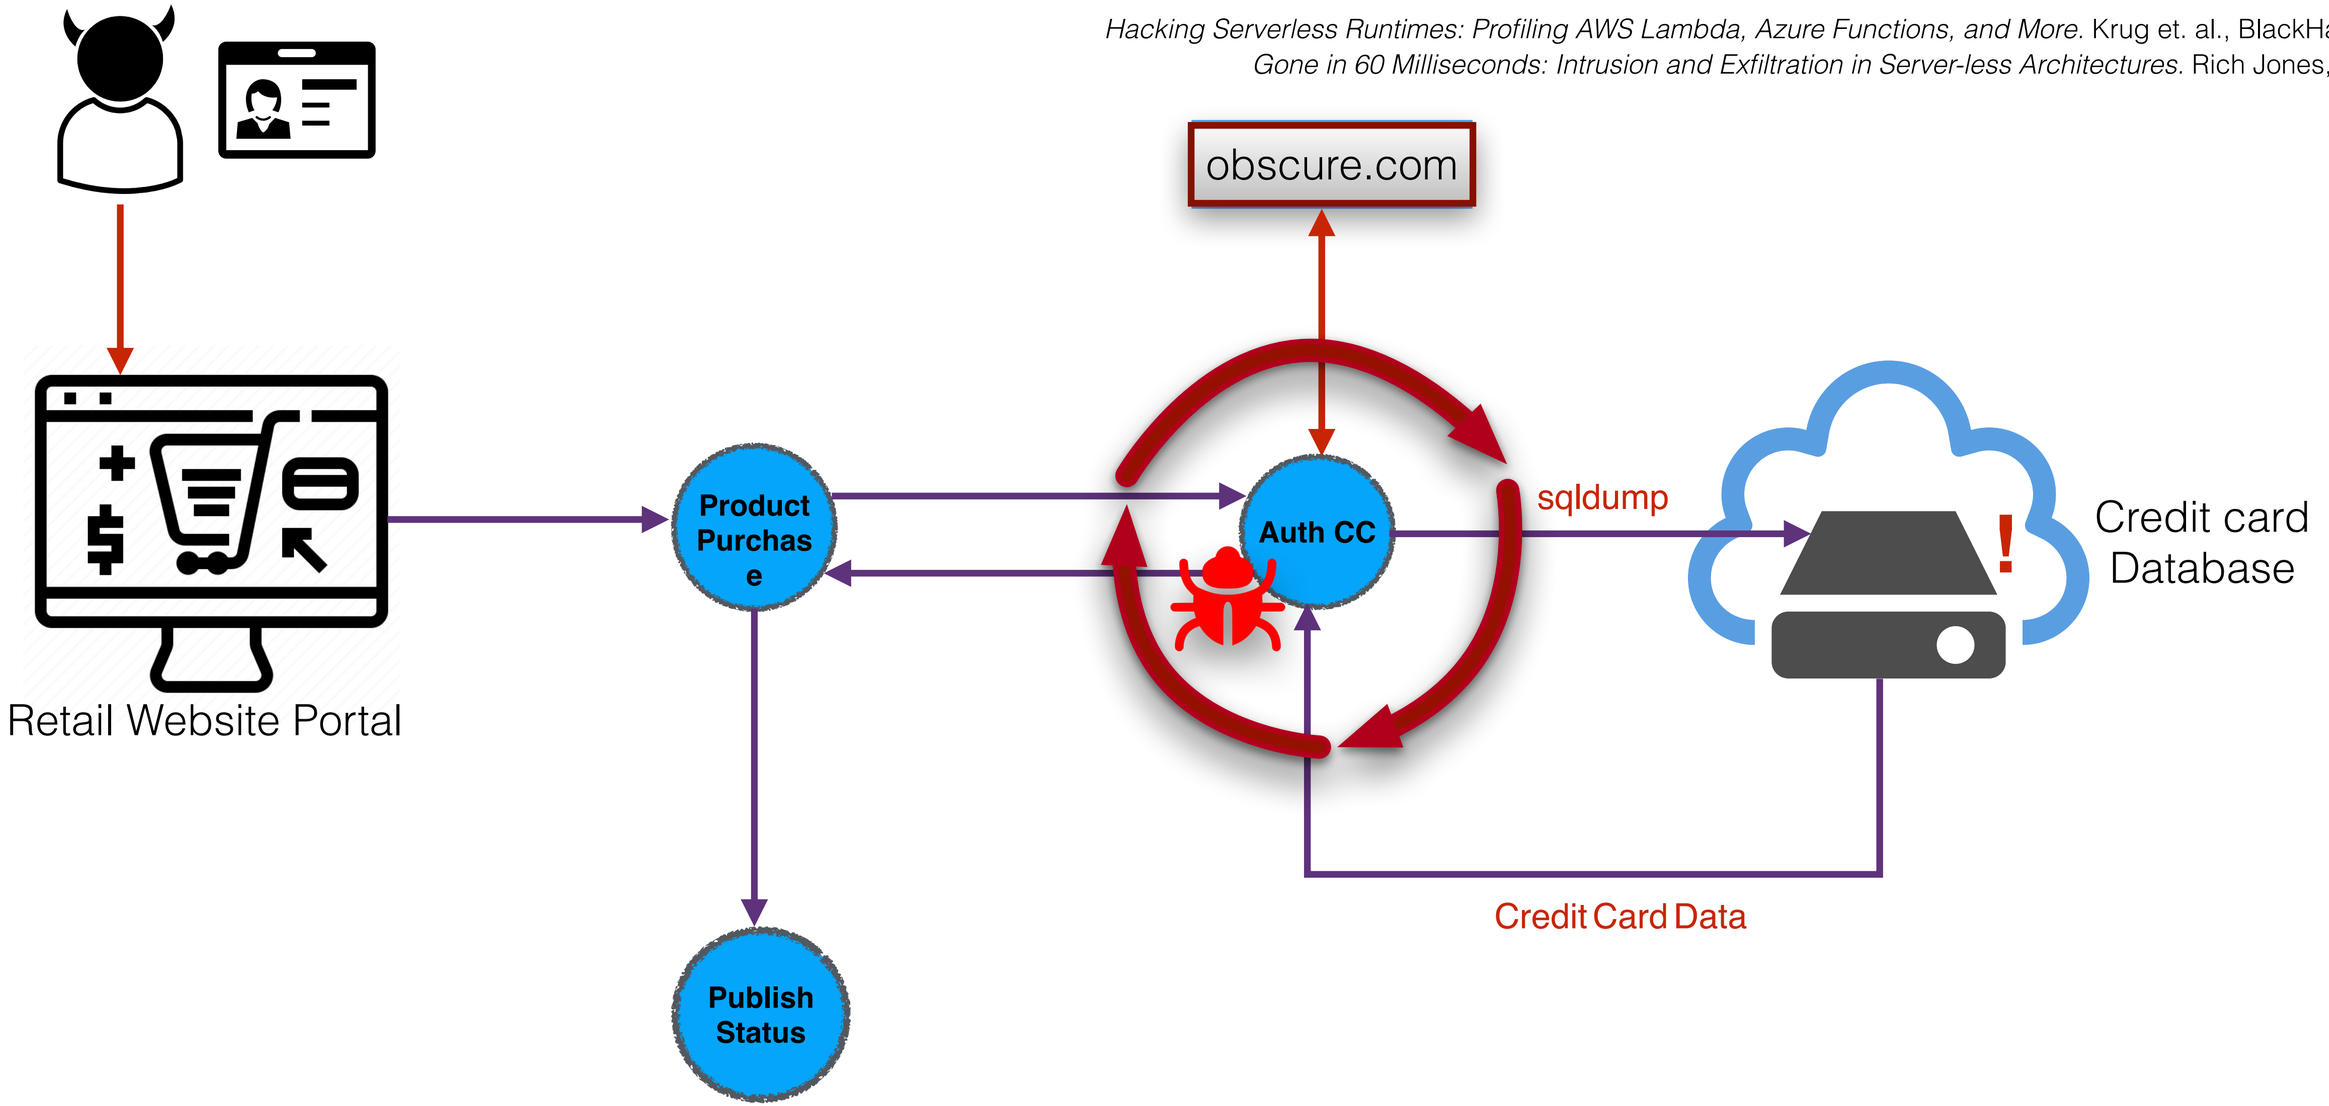

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

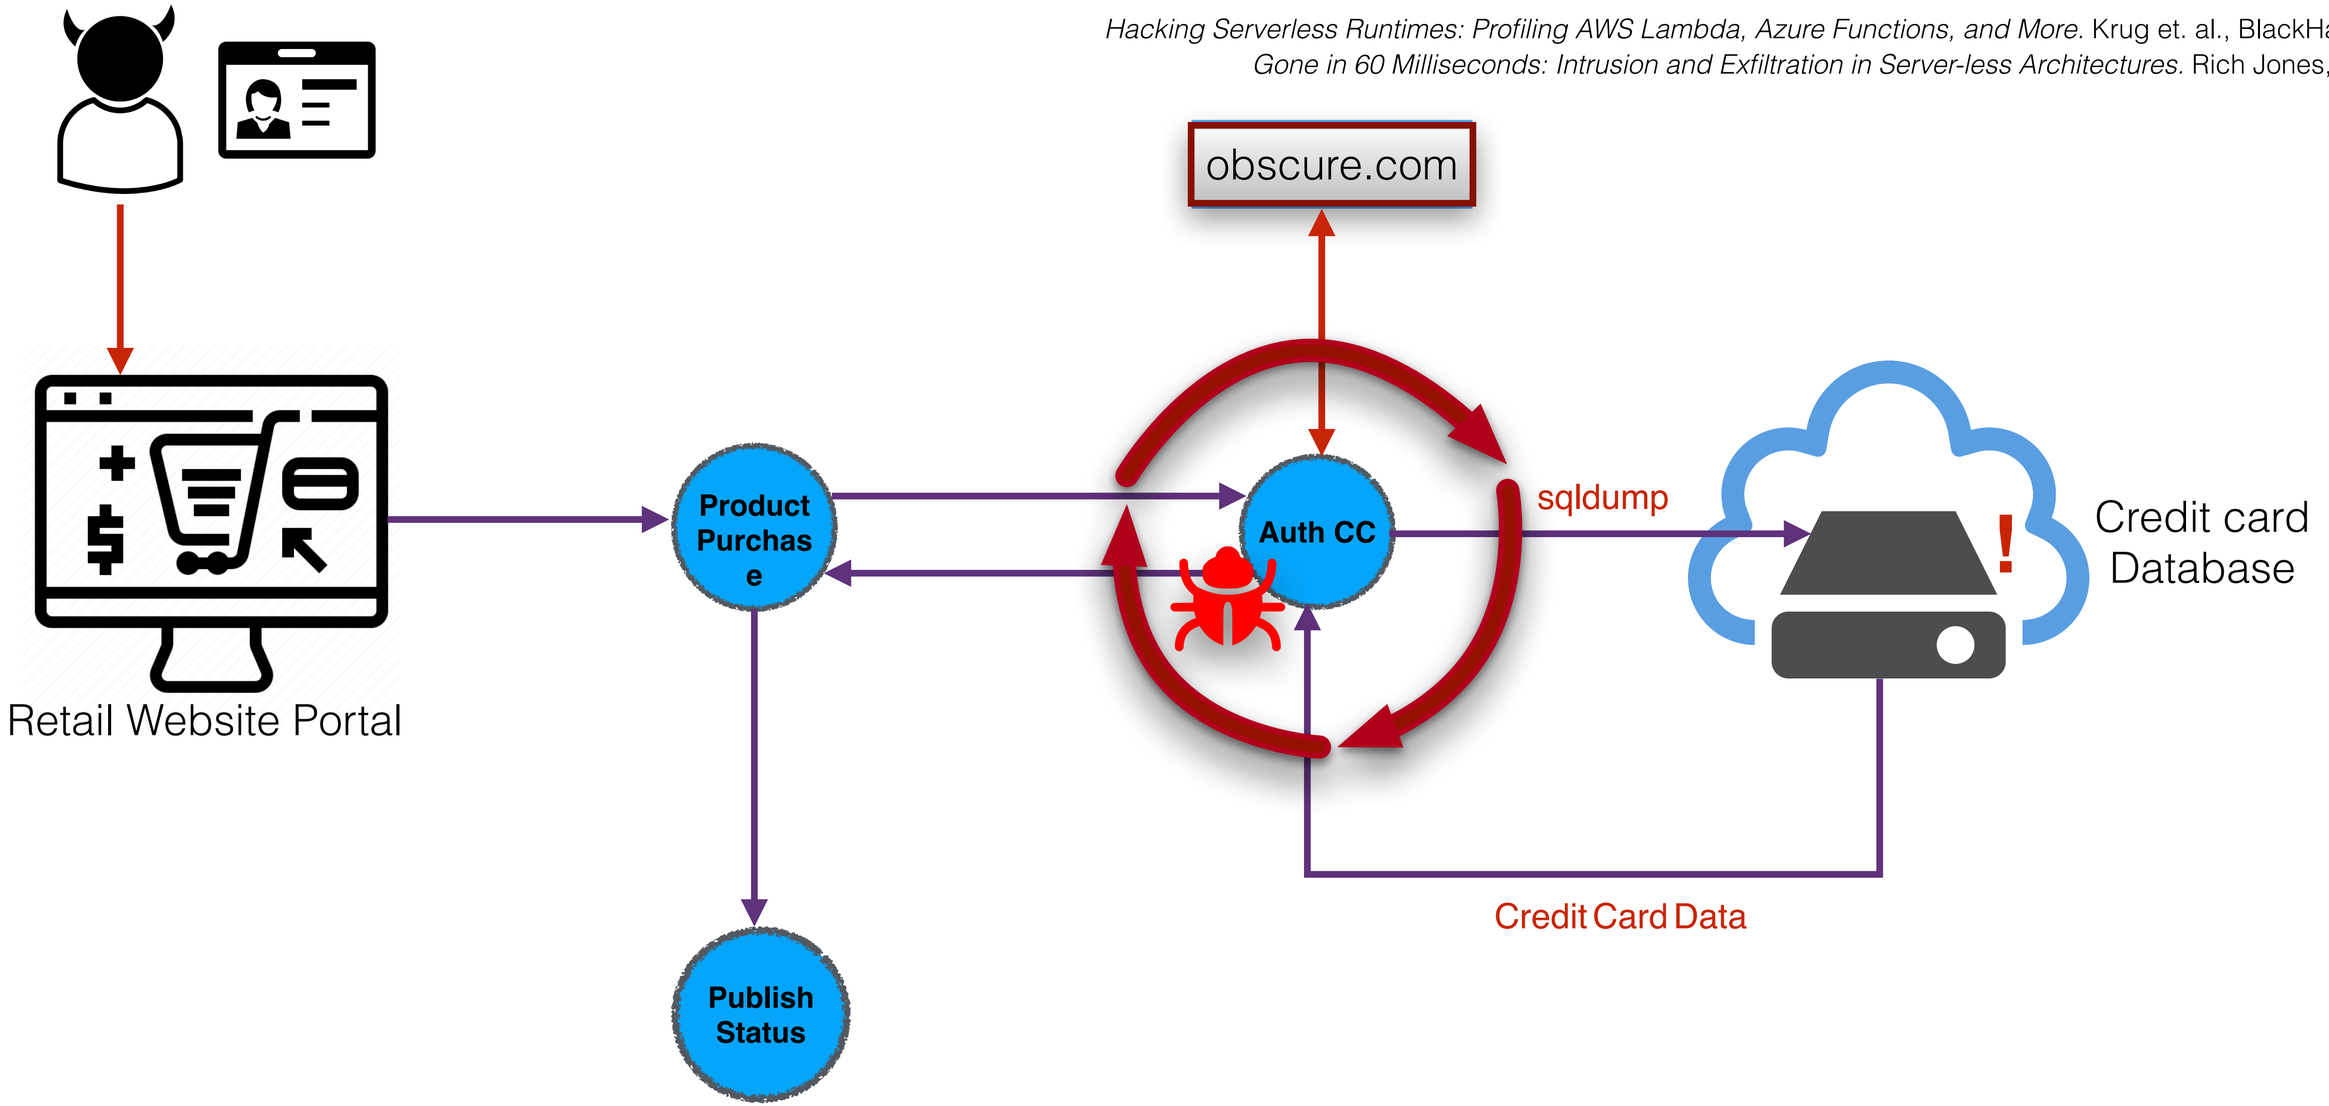

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

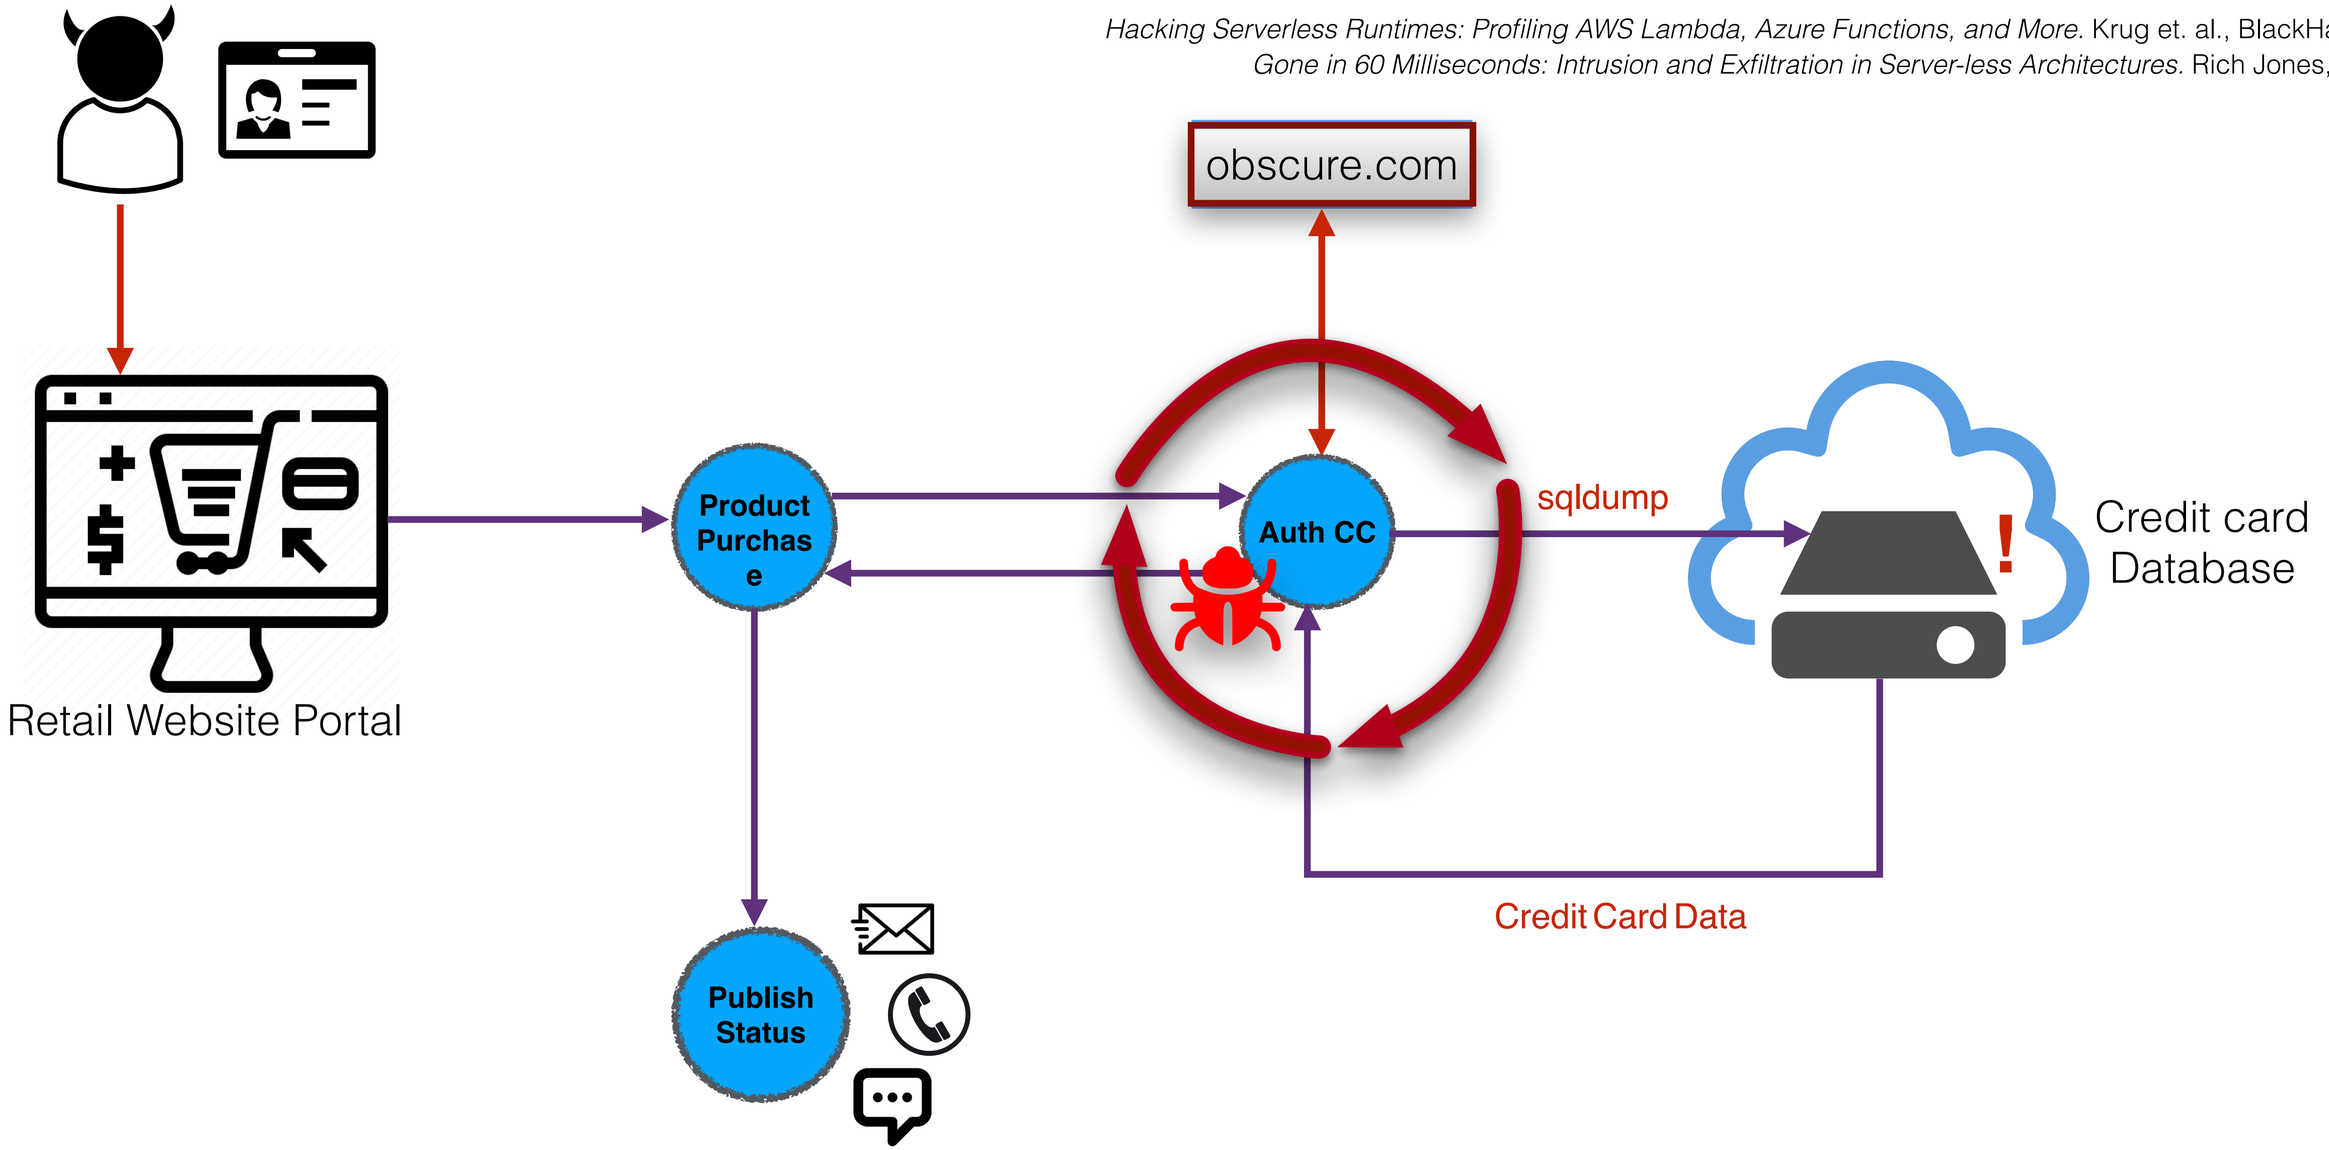

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

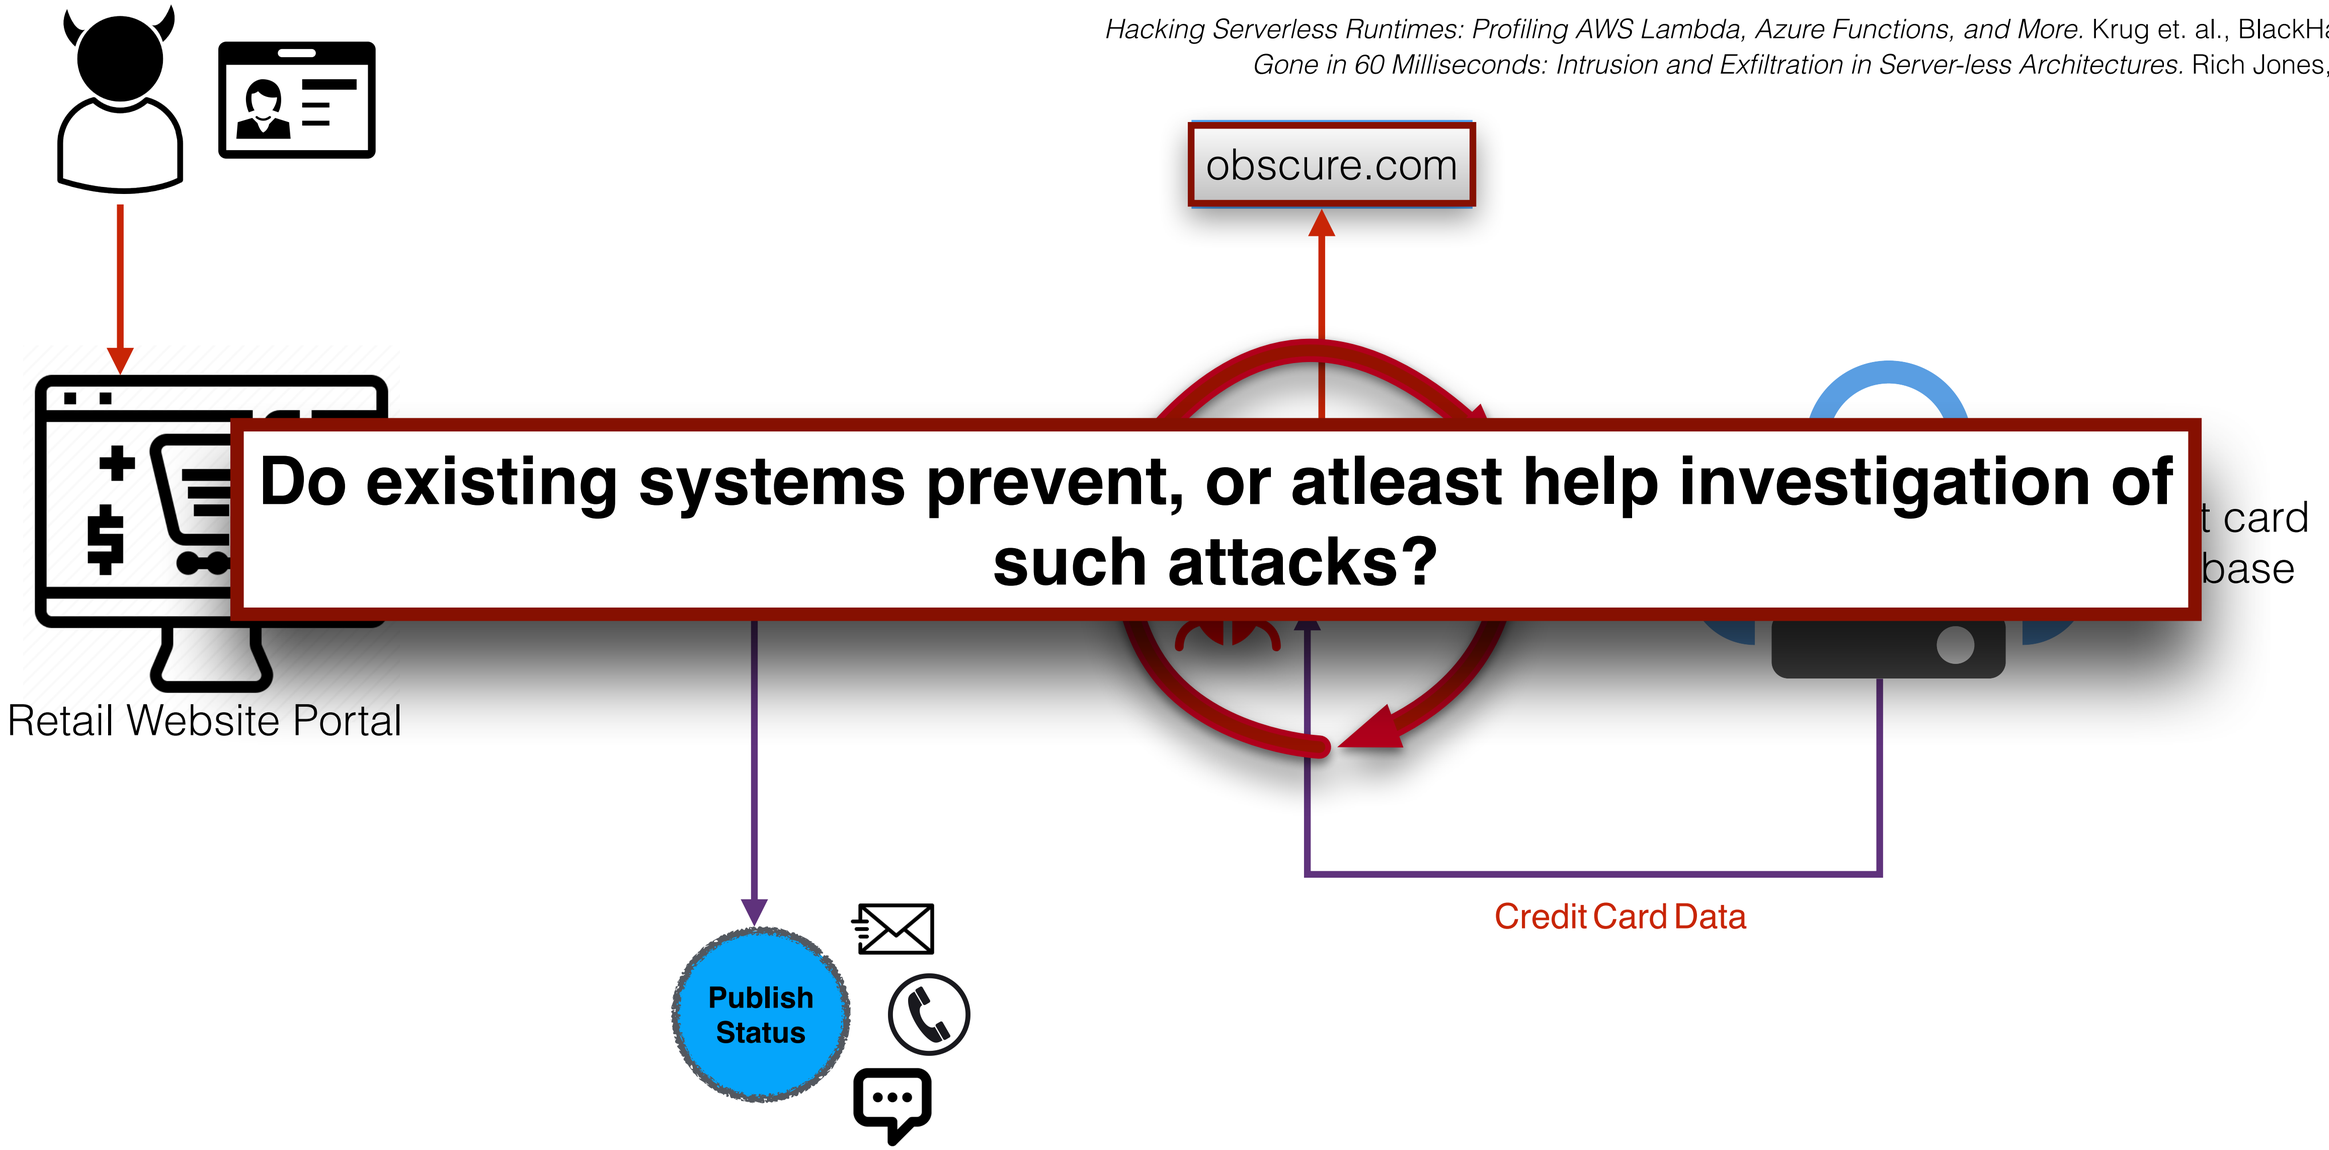

Limitations of existing solutions

- Traditional Approaches (e.g., Linux Audit)
  - Oblivious to serverless semantics
  - Pruning of events not useful for serverless forensics
- Observability tools
  - Usage limits, runtime and platform dependent

Limitations of existing solutions

- Traditional Approaches (e.g., Linux Audit)
  - Oblivious to serverless semantics
  - Pruning of events not useful for serverless forensics
- Observability tools
  - Usage limits, runtime and platform dependent

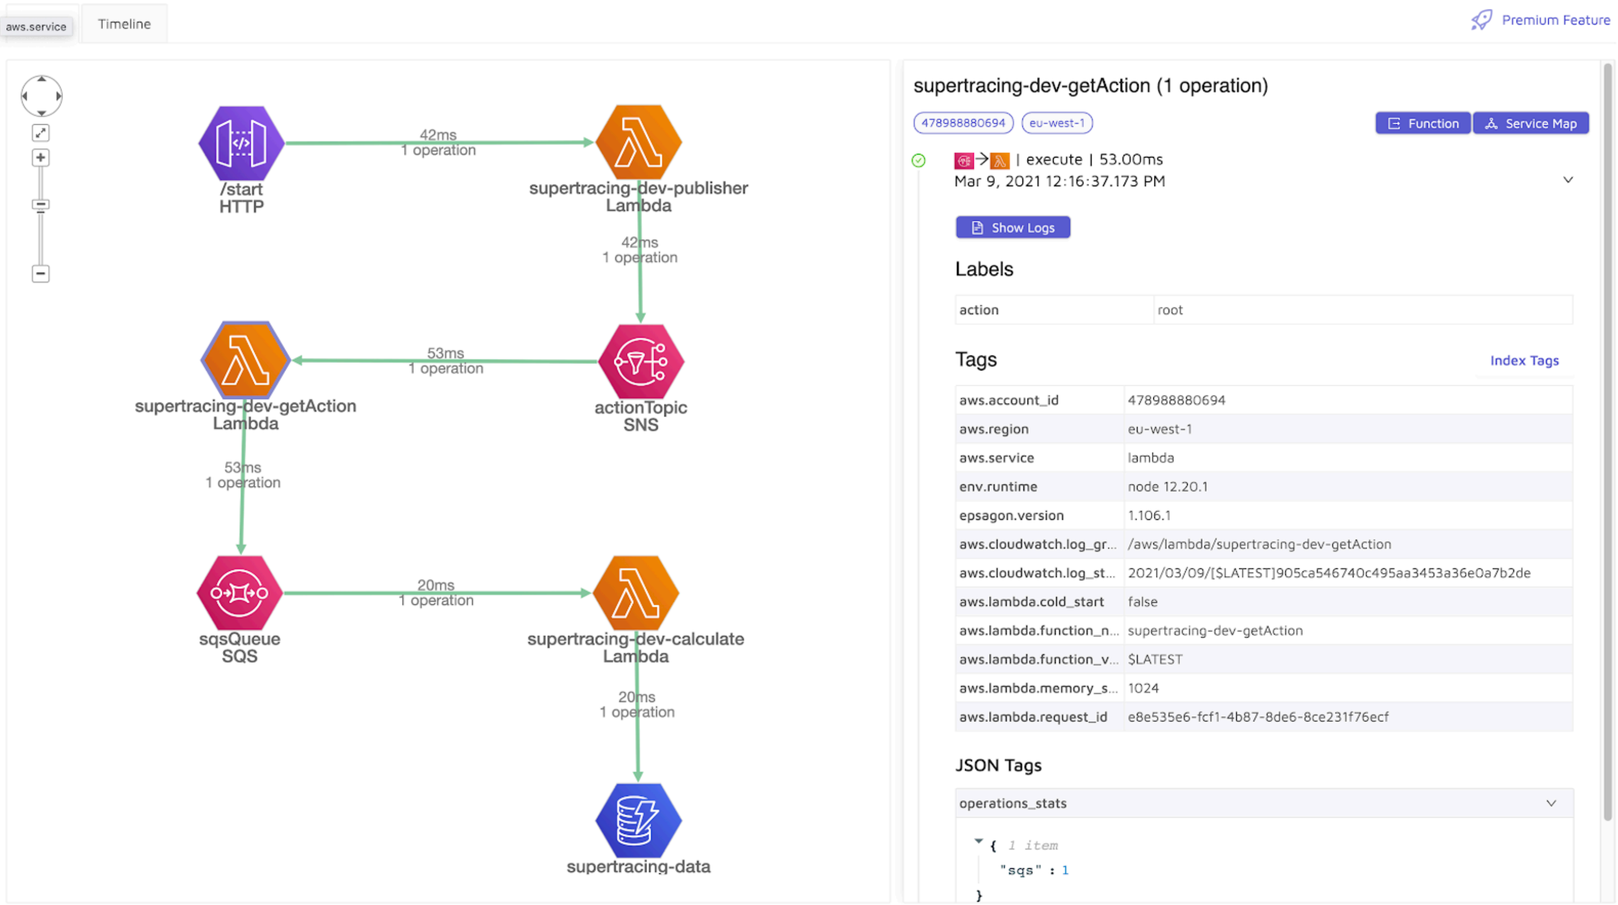

Limitations of existing solutions

- Traditional Approaches (e.g., Linux Audit)
  - Oblivious to serverless semantics
  - Pruning of events not useful for serverless forensics
- Observability tools
  - Usage limits, runtime and platform dependent

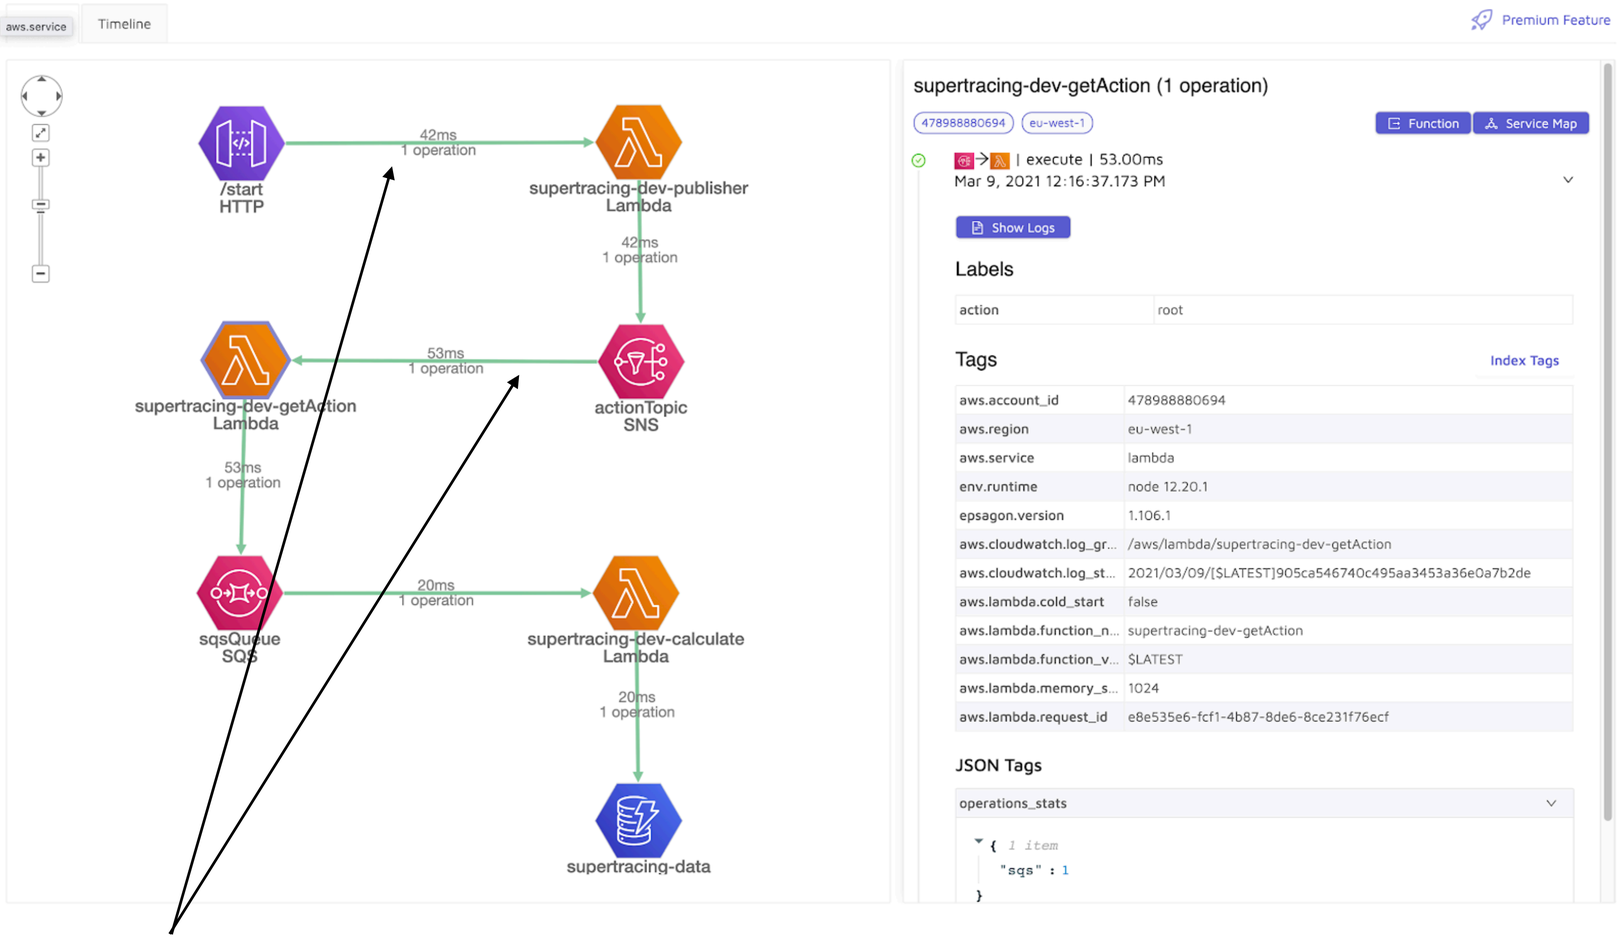

Coarse grained function communications

Limitations of existing solutions

- Traditional Approaches (e.g., Linux Audit)
  - Oblivious to serverless semantics
  - Pruning of events not useful for serverless forensics
- Observability tools
  - Usage limits, runtime and platform dependent

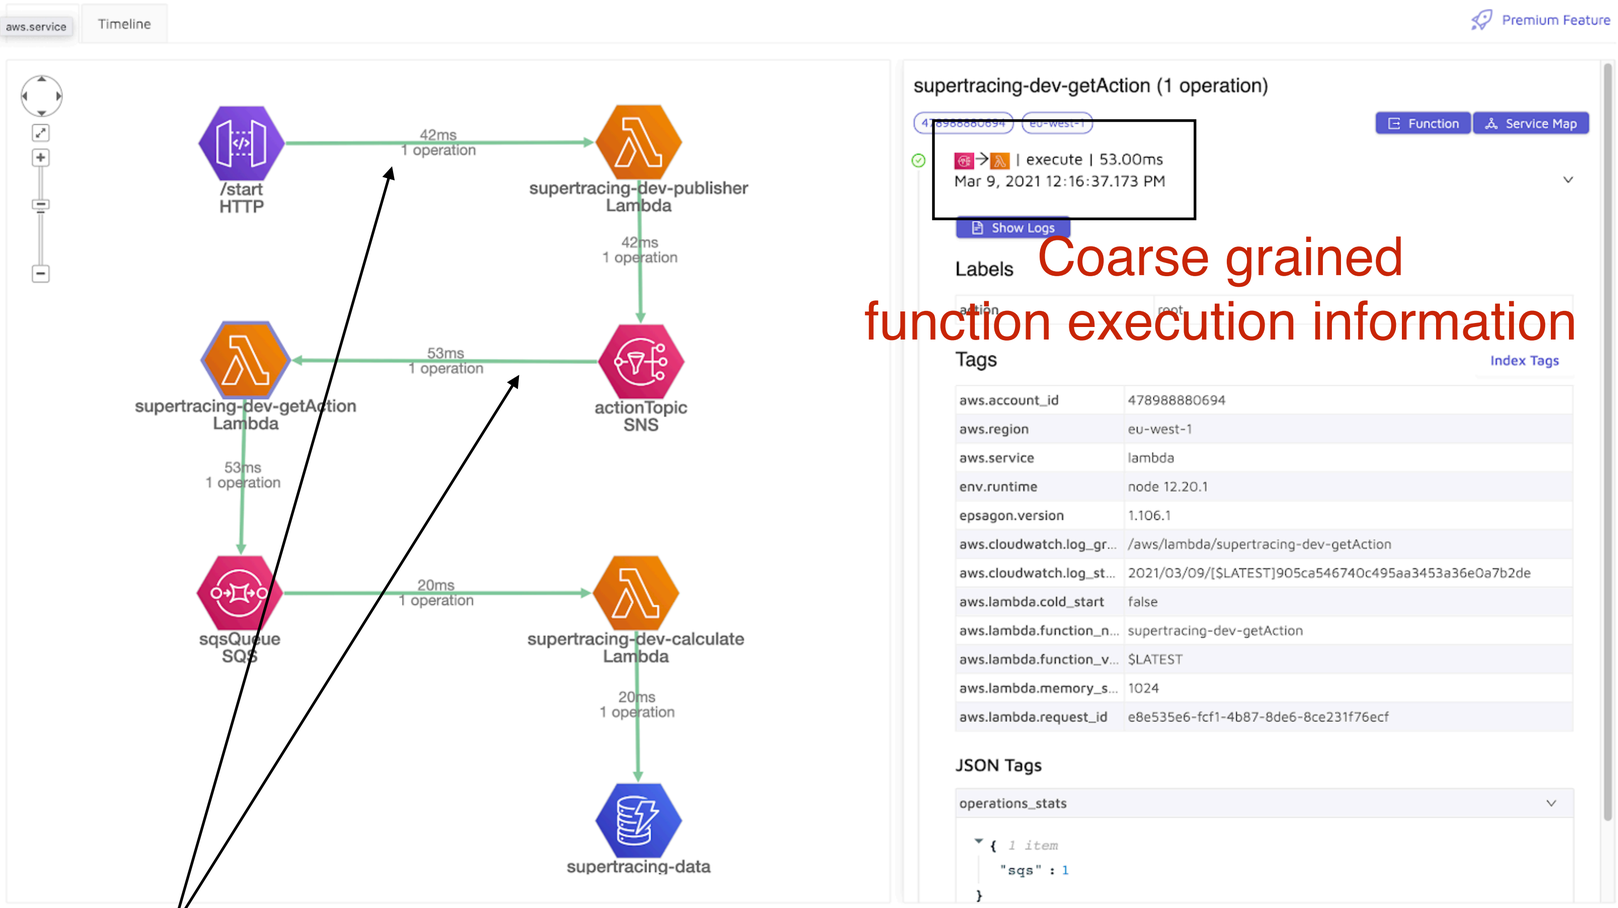

Coarse grained function execution information

Coarse grained function communications

Fine-grained Serverless Auditing

Fine-grained Serverless Auditing

Alastor: A complete provenance collection and auditing framework for serverless platforms

Alastor Architecture

Serverless provenance:
What are the specific agents, entities, and activities that must be **audited?**

Alastor Architecture

Serverless provenance:
What are the specific agents, entities, and activities that must be **audited?**

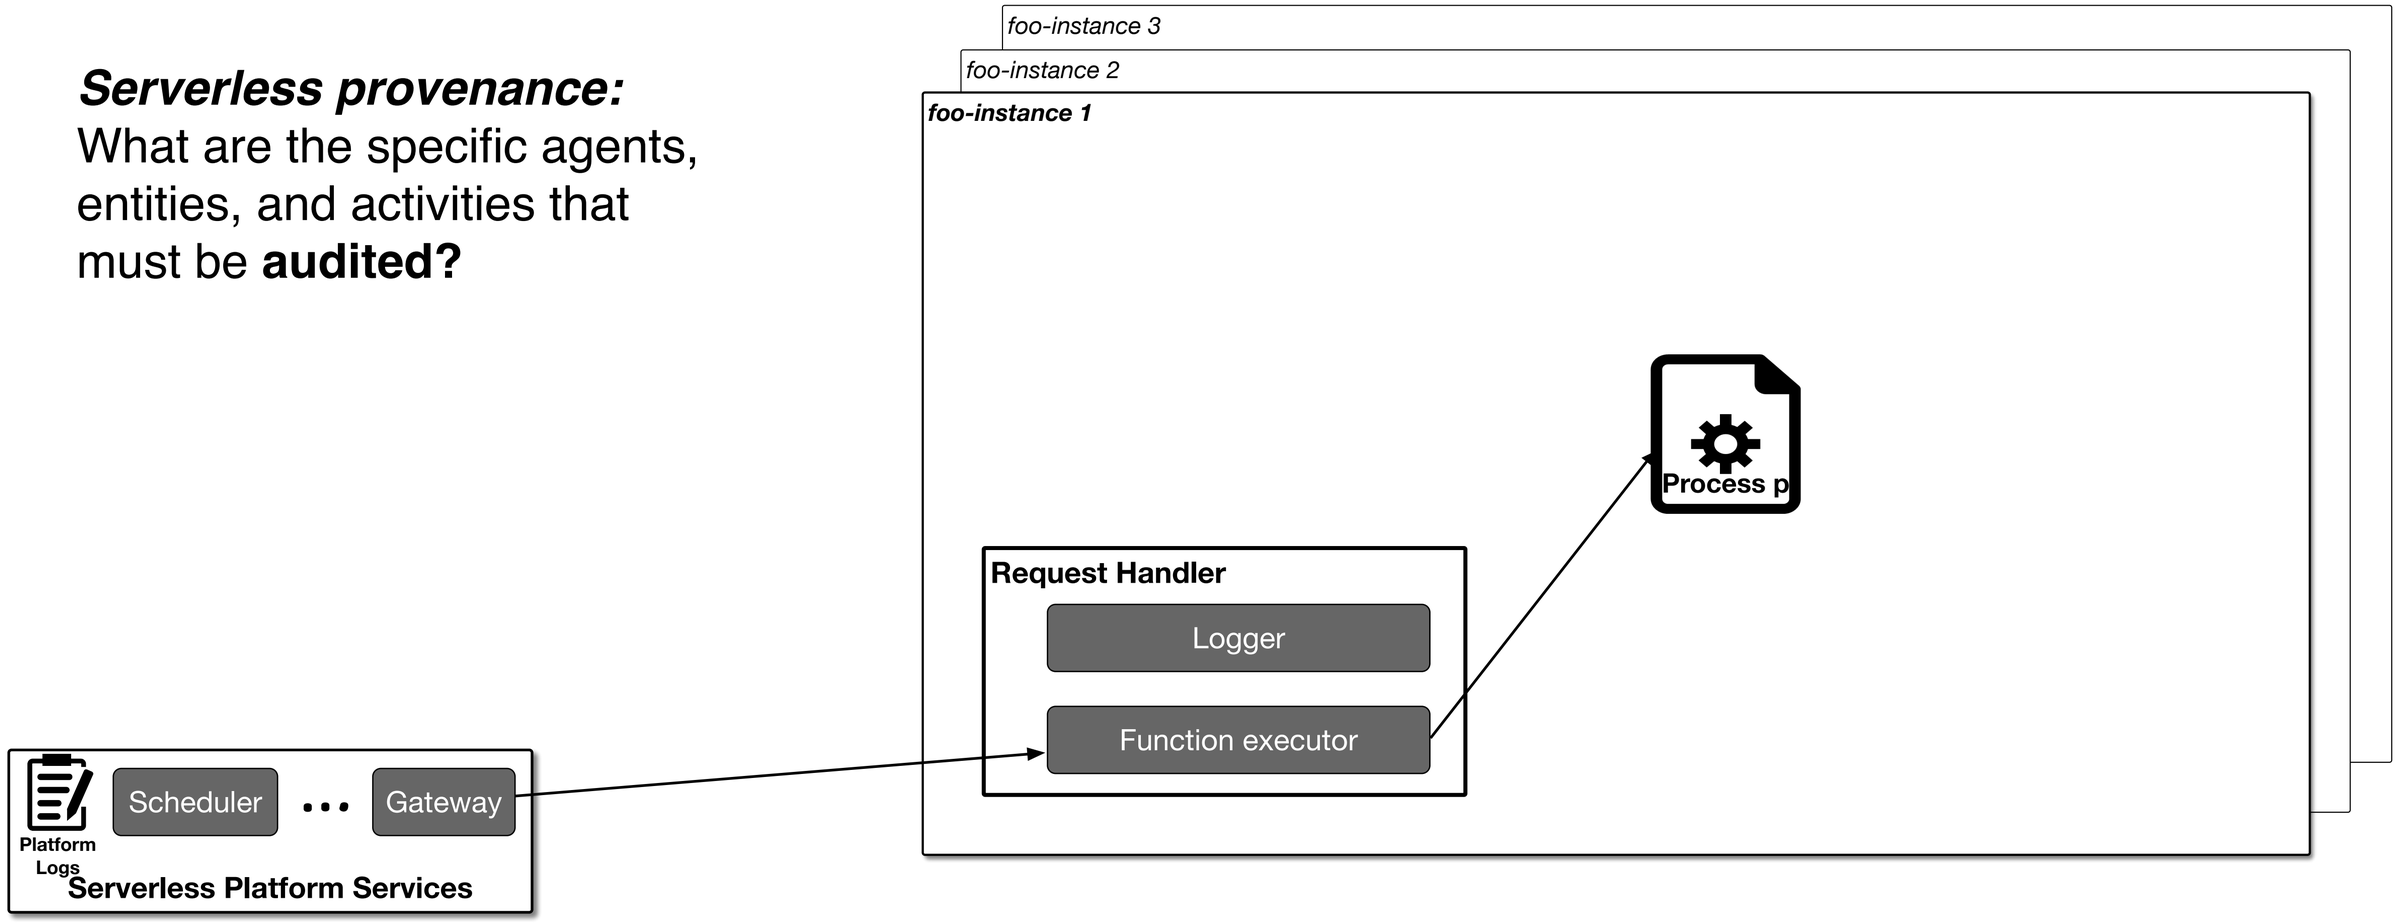

Alastor Architecture

Serverless provenance:
What are the specific agents, entities, and activities that must be **audited?**

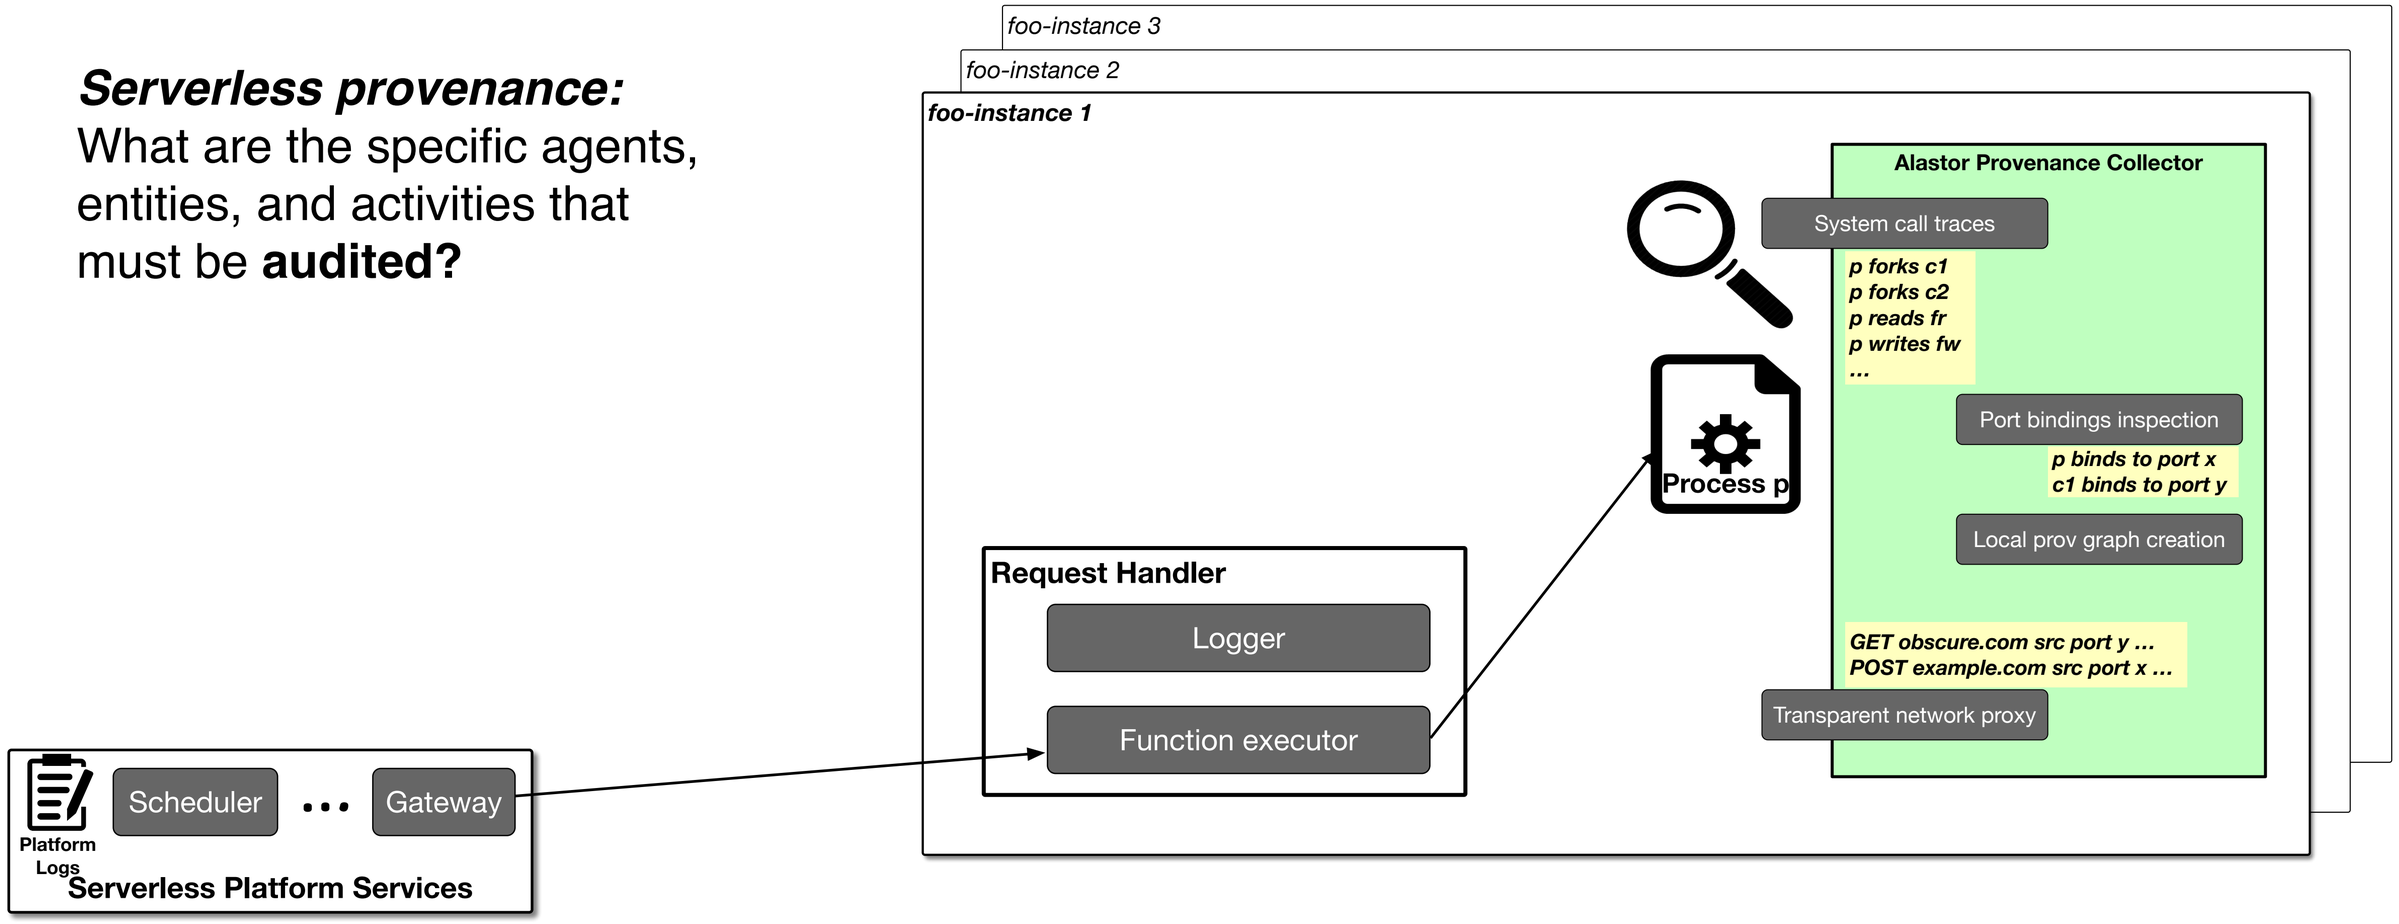

Alastor Architecture

Serverless provenance:
What are the specific agents, entities, and activities that must be **audited?**

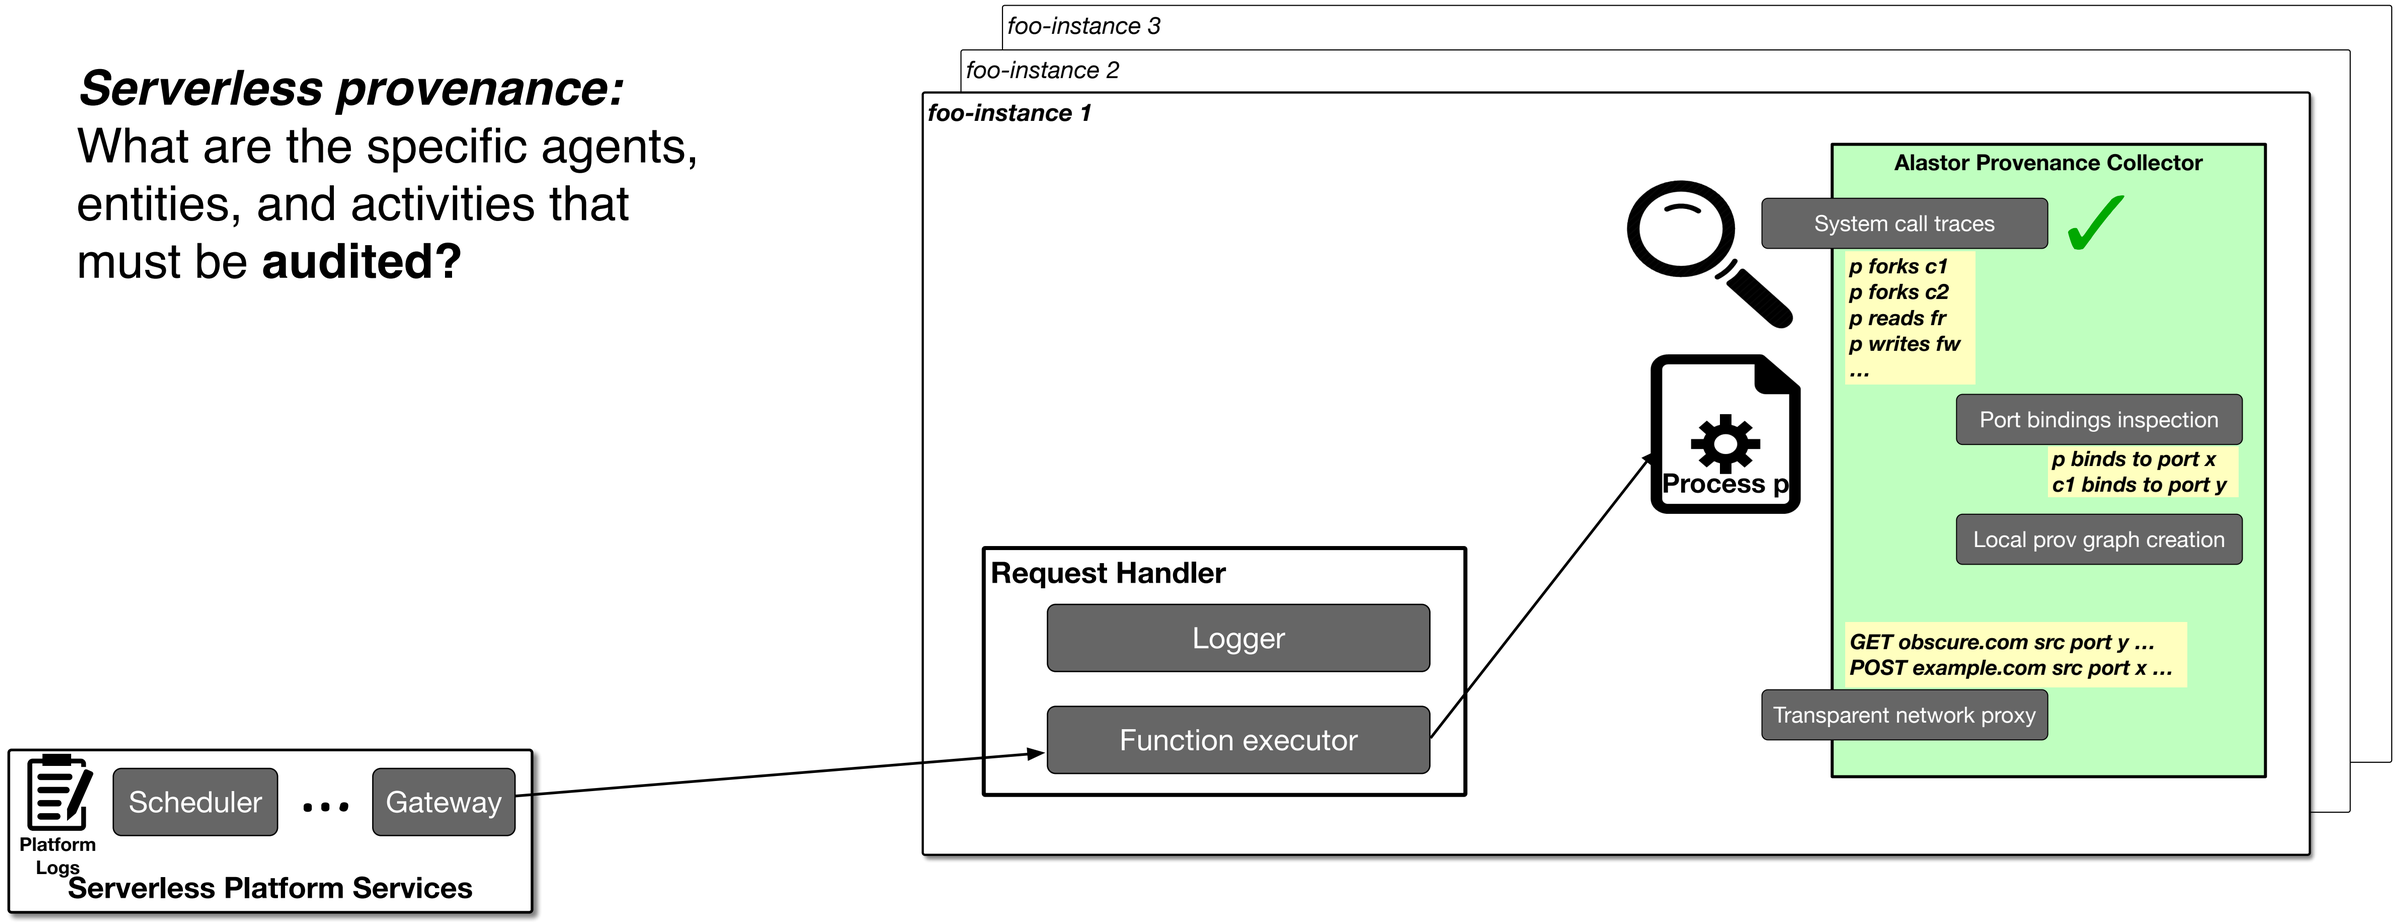

Alastor Architecture

Serverless provenance:
What are the specific agents, entities, and activities that must be **audited?**

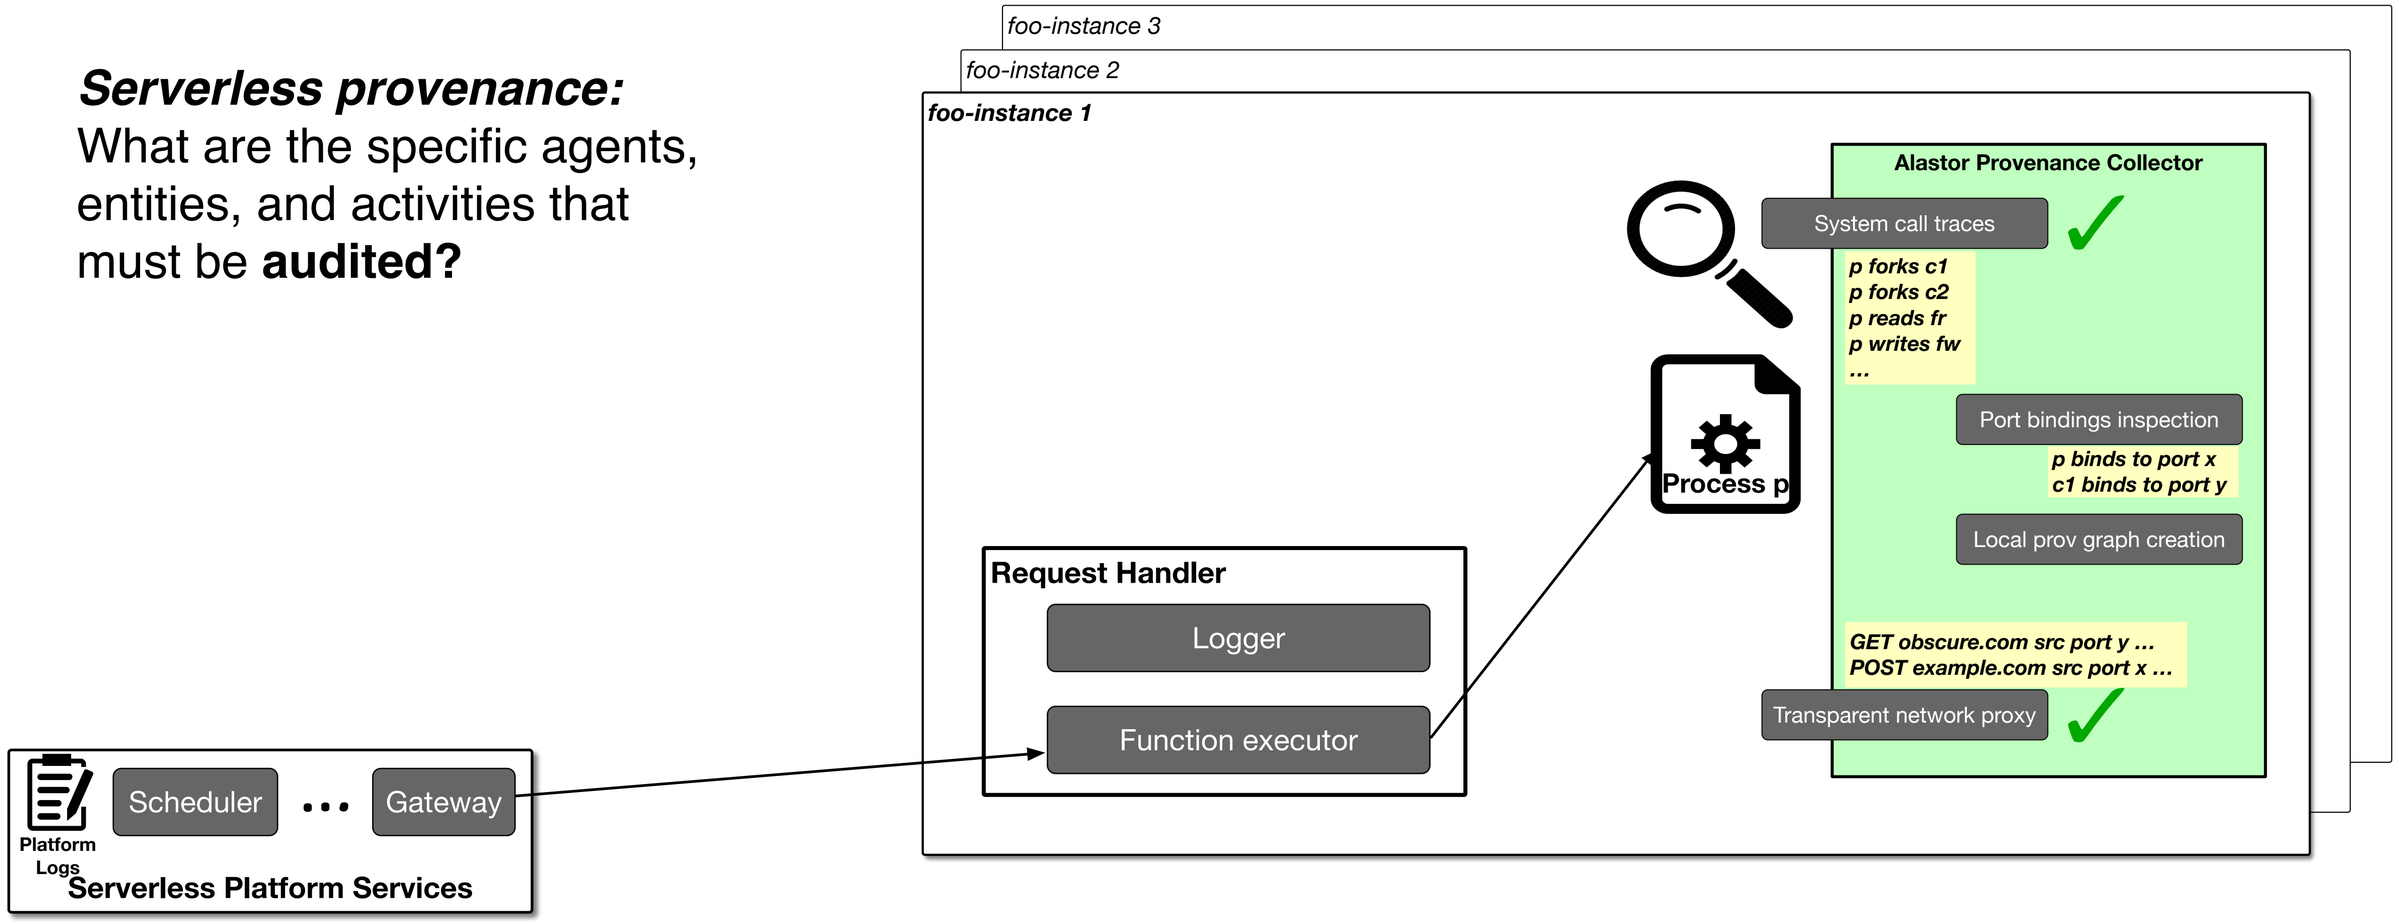

Alastor Architecture

Serverless provenance:
What are the specific agents, entities, and activities that must be **audited?**

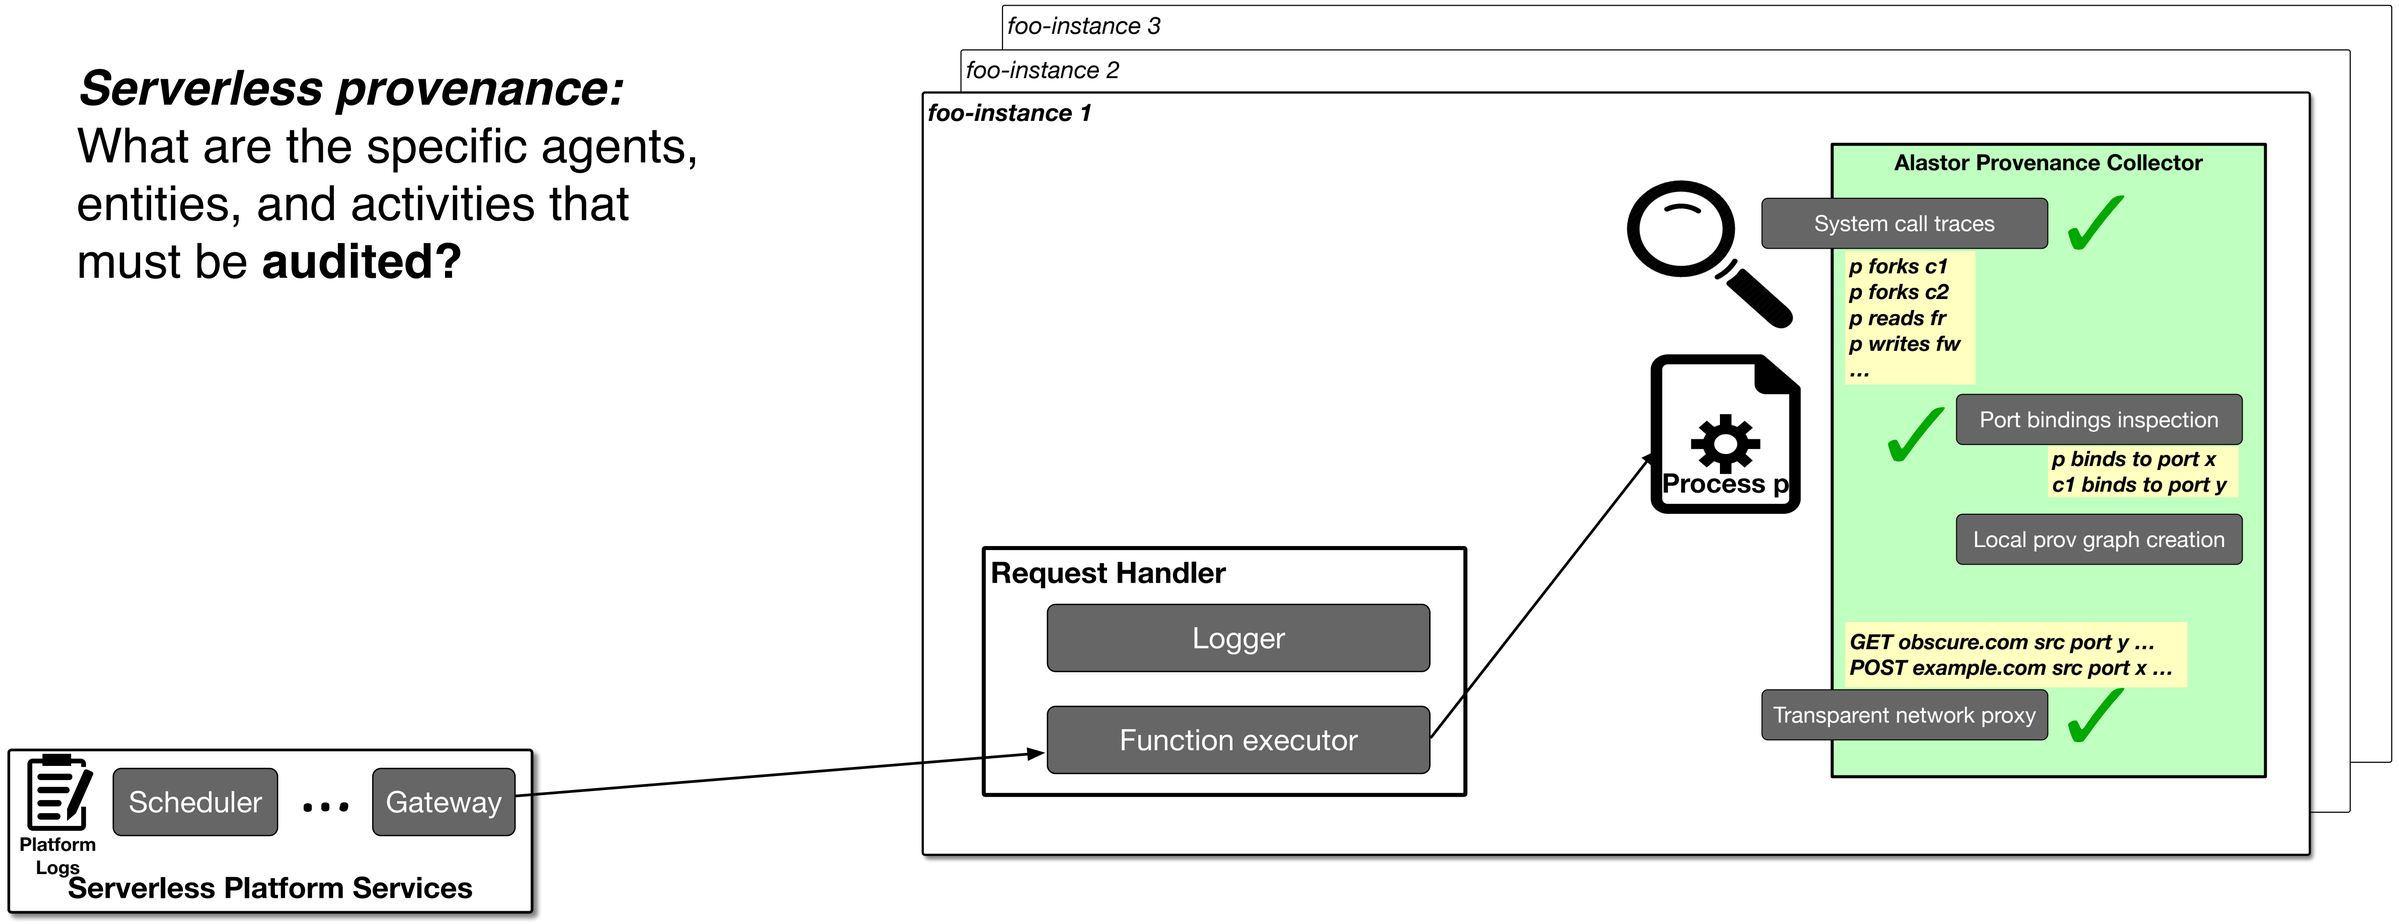

Alastor Architecture

Serverless provenance:
What are the specific agents, entities, and activities that must be **audited?**

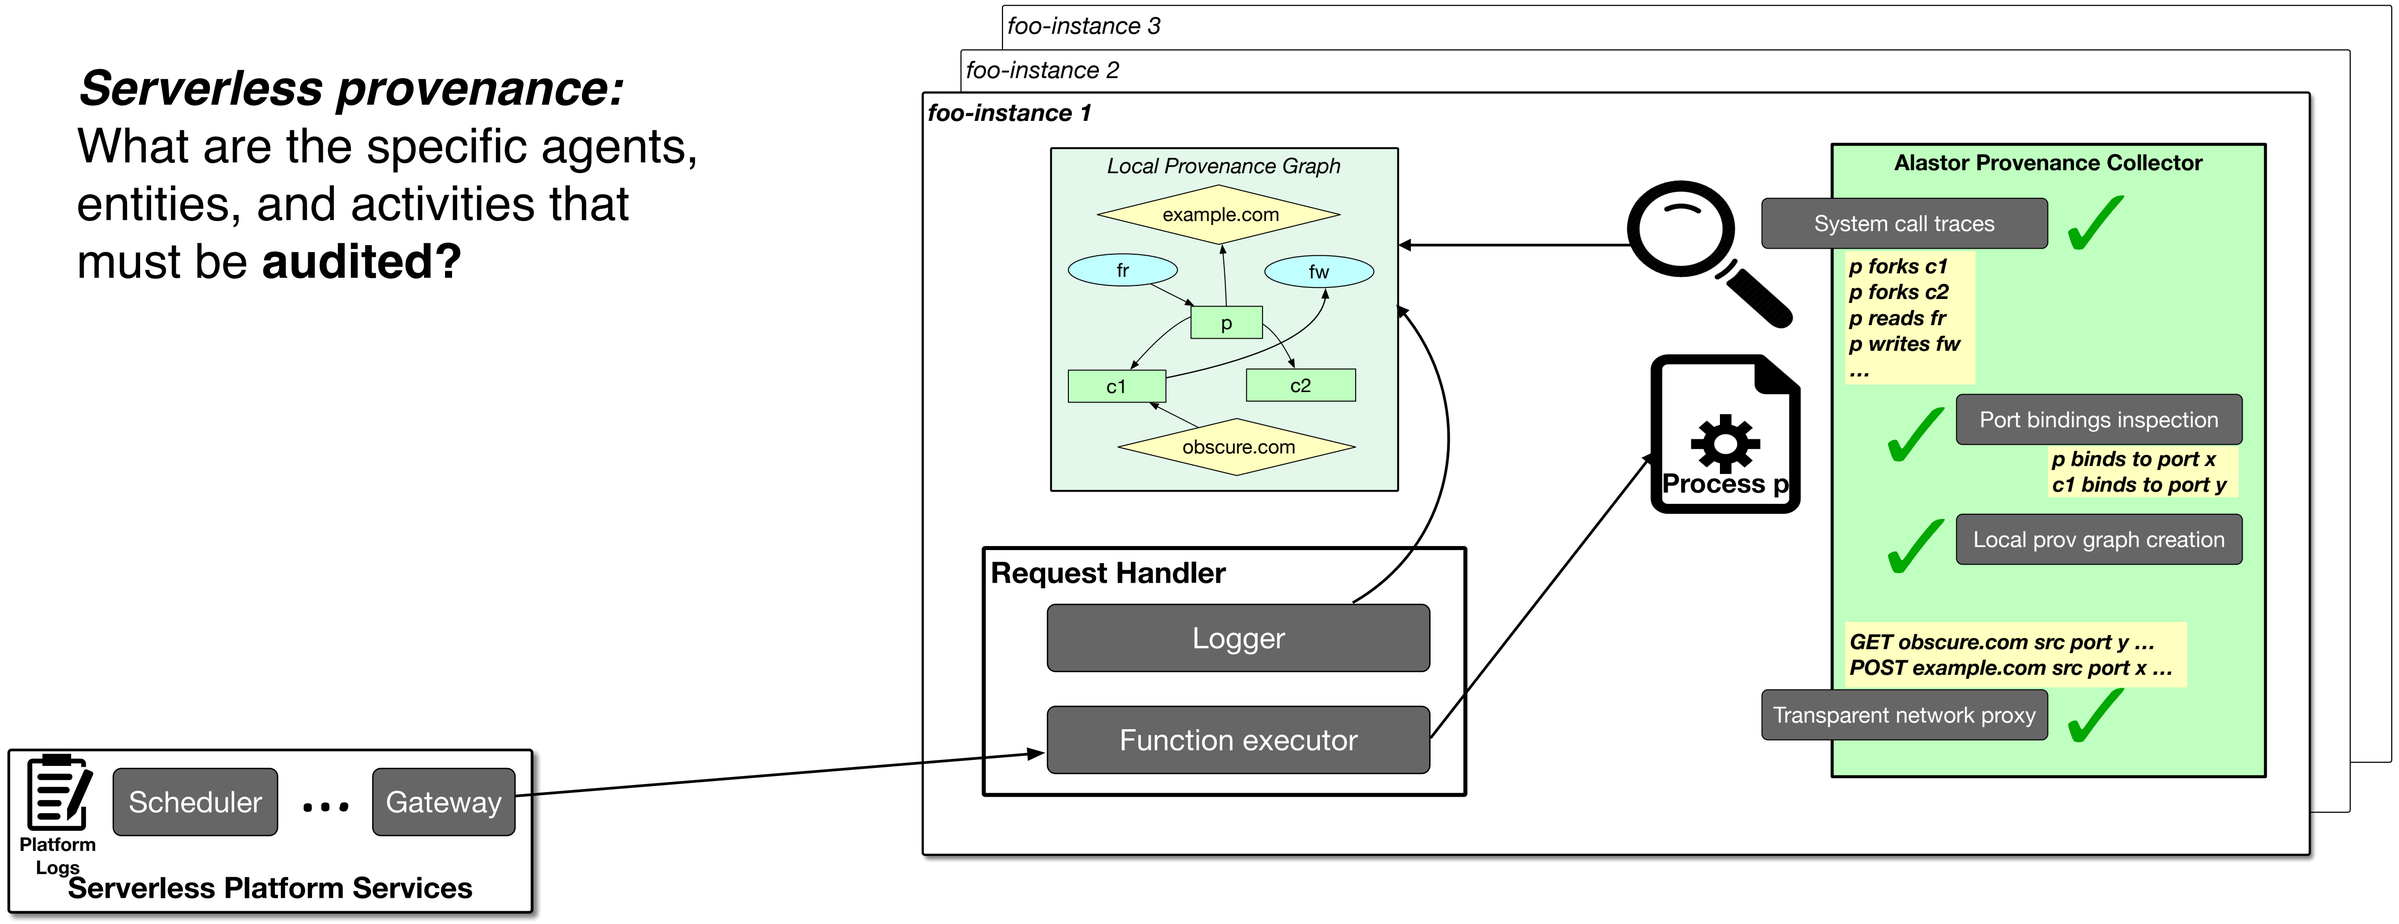

Alastor Architecture

Universal auditing for serverless:

How can forensic evidence at various levels of the serverless stack be integrated to facilitate effective threat investigation?

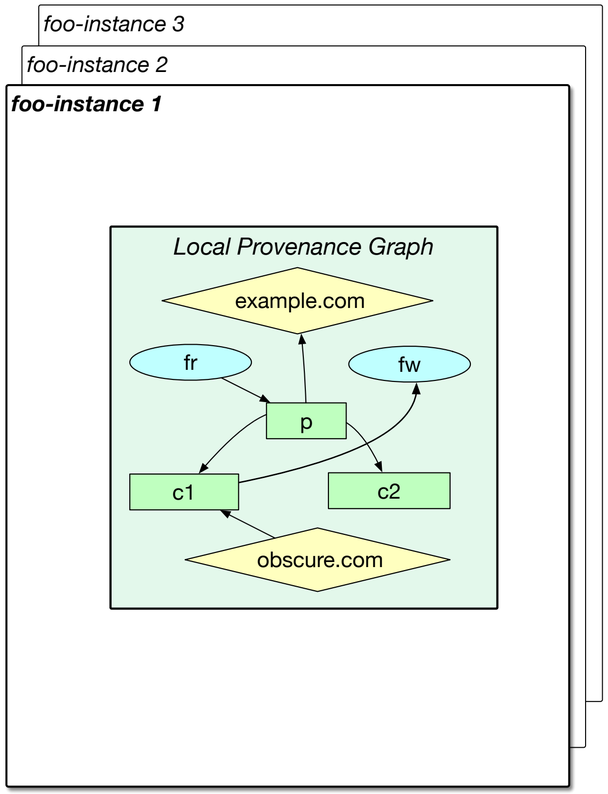

Alastor Architecture

Universal auditing for serverless:

How can forensic evidence at various levels of the serverless stack be integrated to facilitate effective threat investigation?

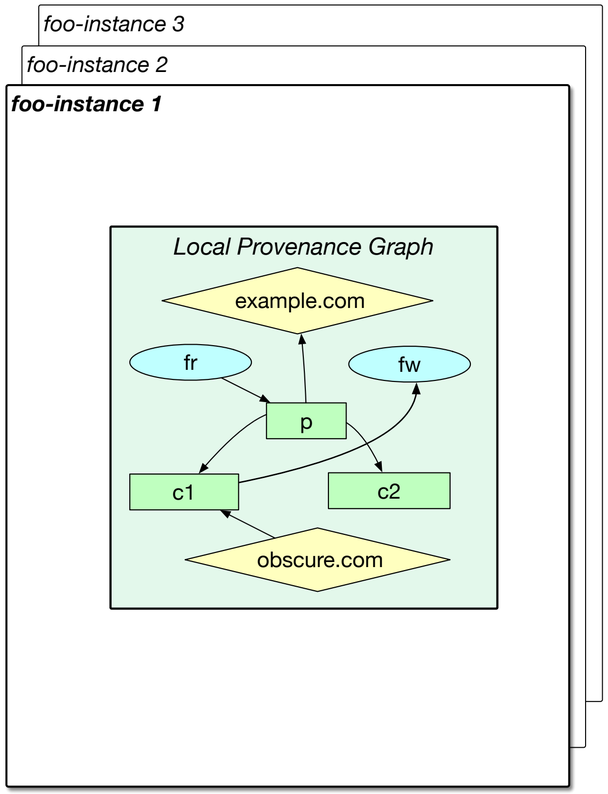

<IP, instanceID>
<instanceID, {reqID}>

Alastor Architecture

Universal auditing for serverless:

How can forensic evidence at various levels of the serverless stack be integrated to facilitate effective threat investigation?

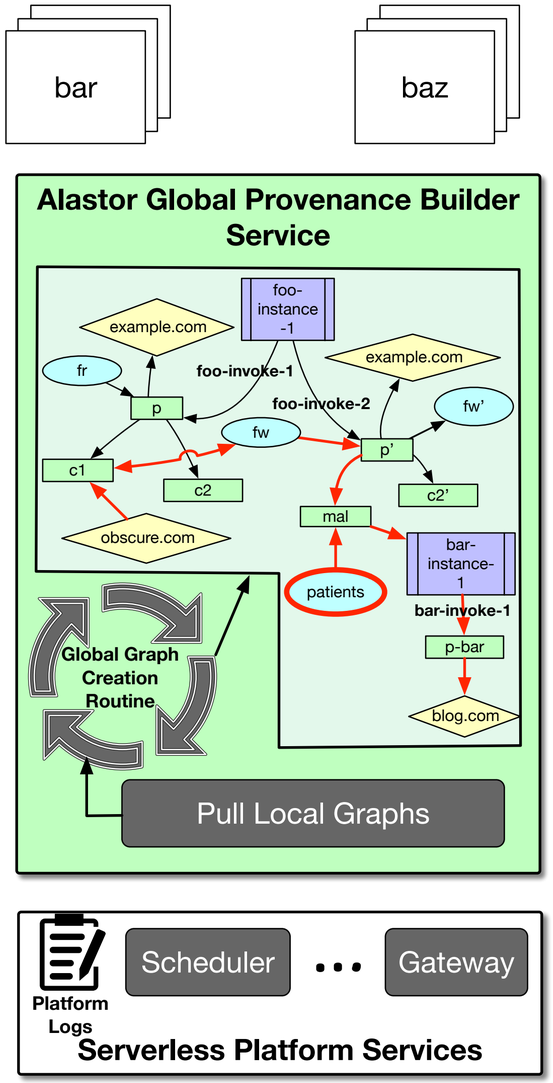

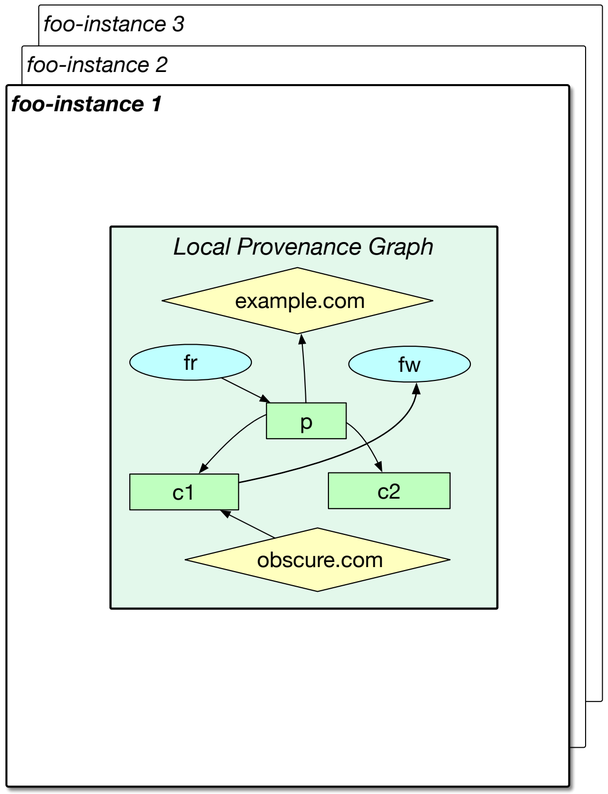

Retail Serverless Application

Hacking Serverless Runtimes: Profiling AWS Lambda, Azure Functions, and More. Krug et. al., BlackHat USA 2017.
Gone in 60 Milliseconds: Intrusion and Exfiltration in Server-less Architectures. Rich Jones, 33C3 2016.

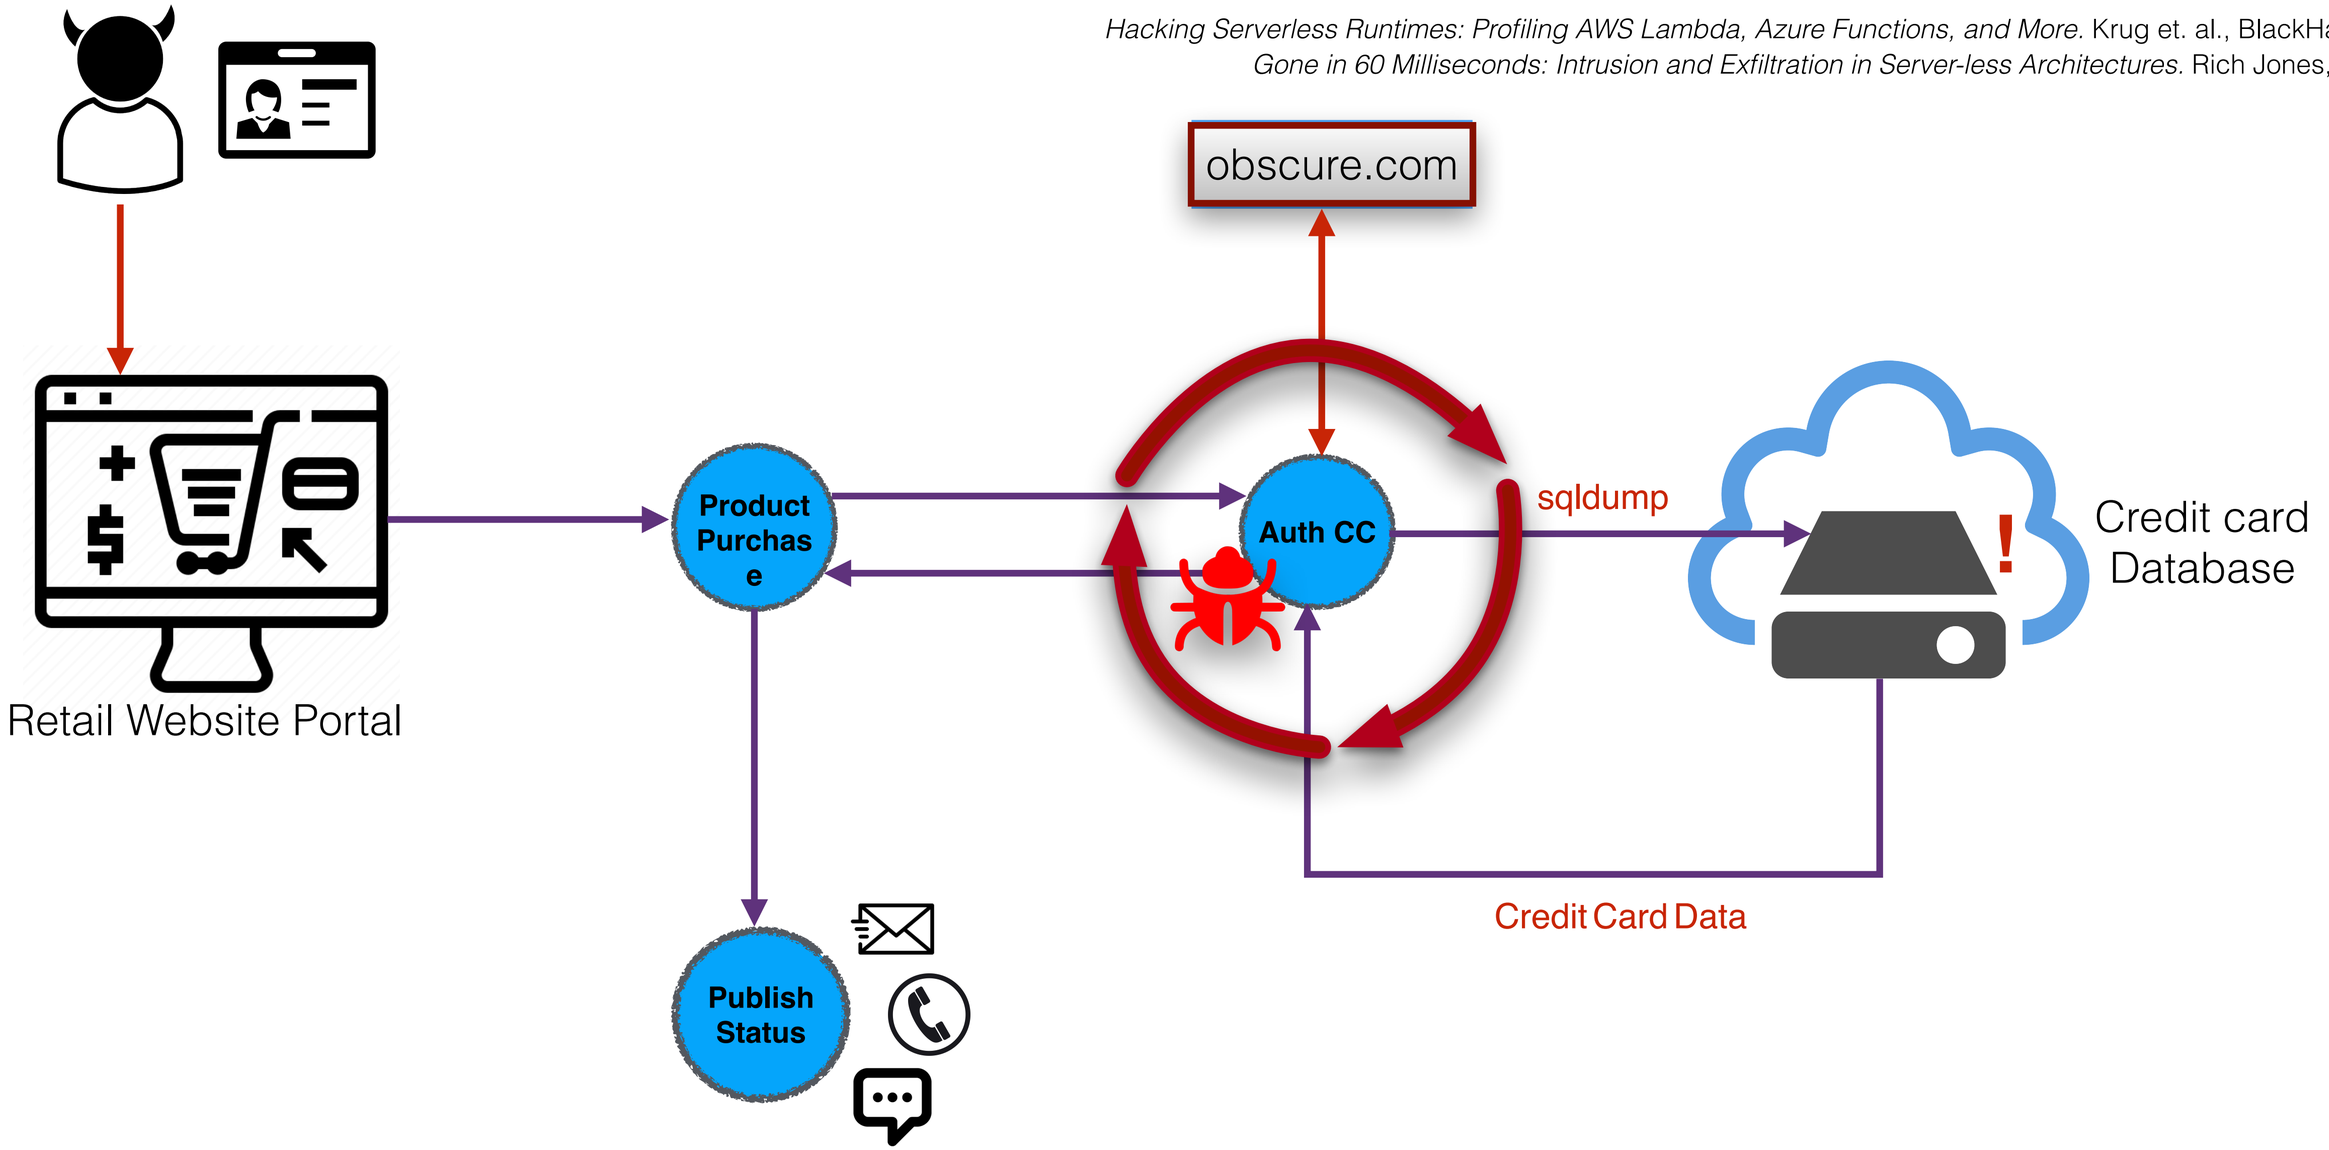

Attack Path Reconstruction Using Epsagon

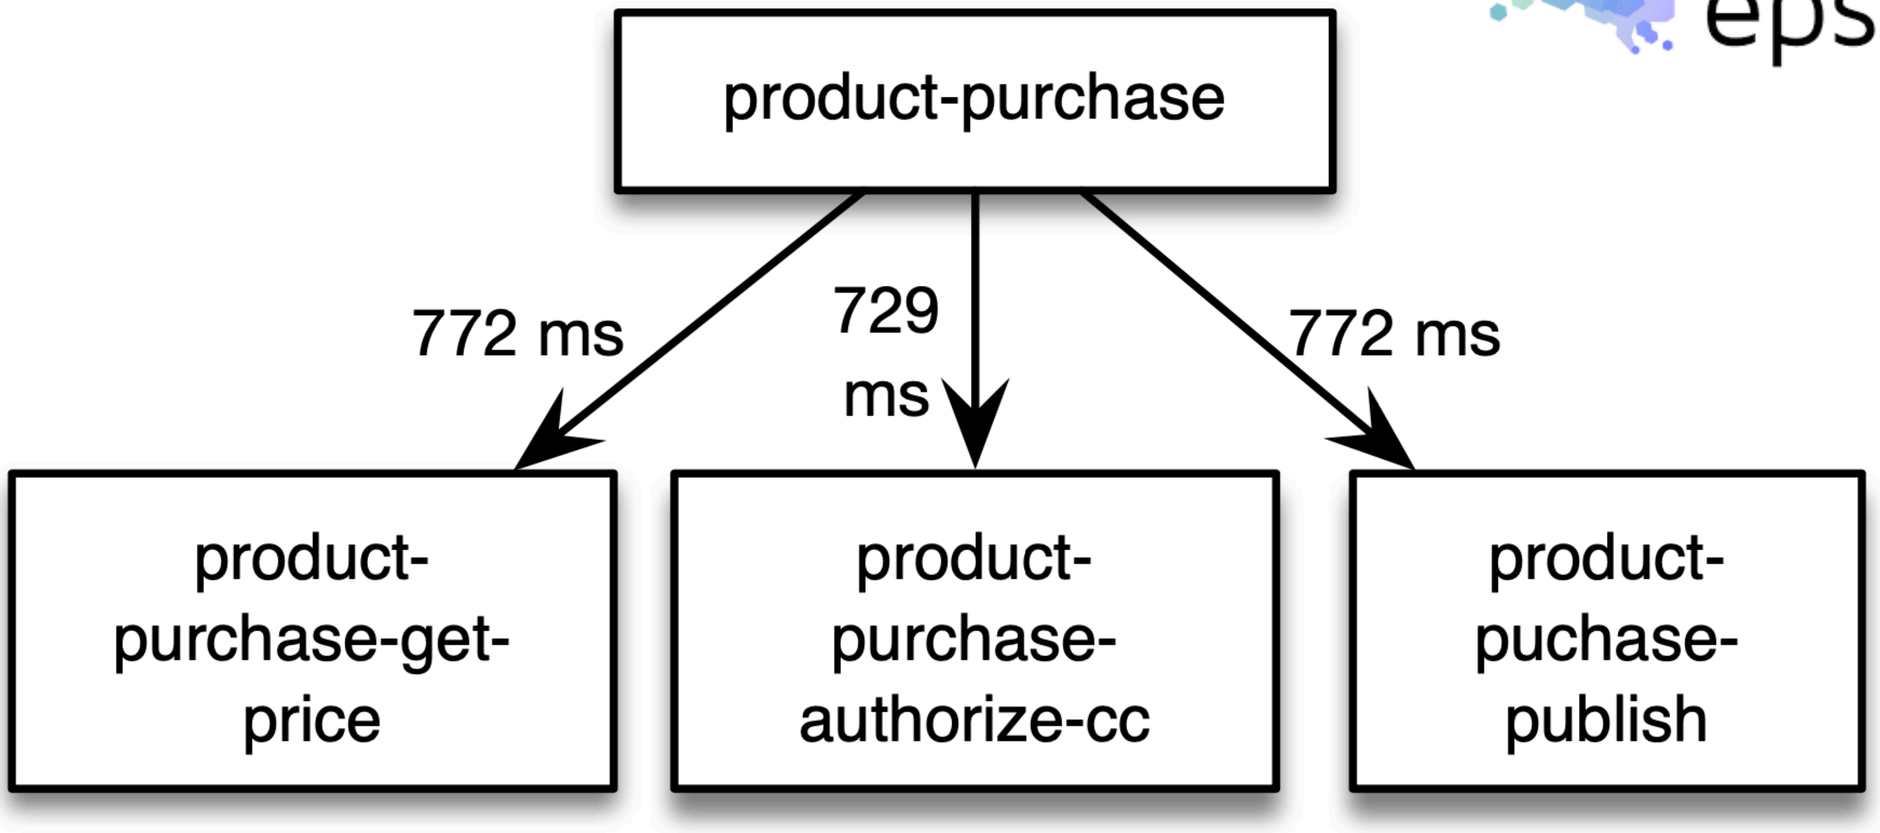

Attack Path Reconstruction Using Epsagon

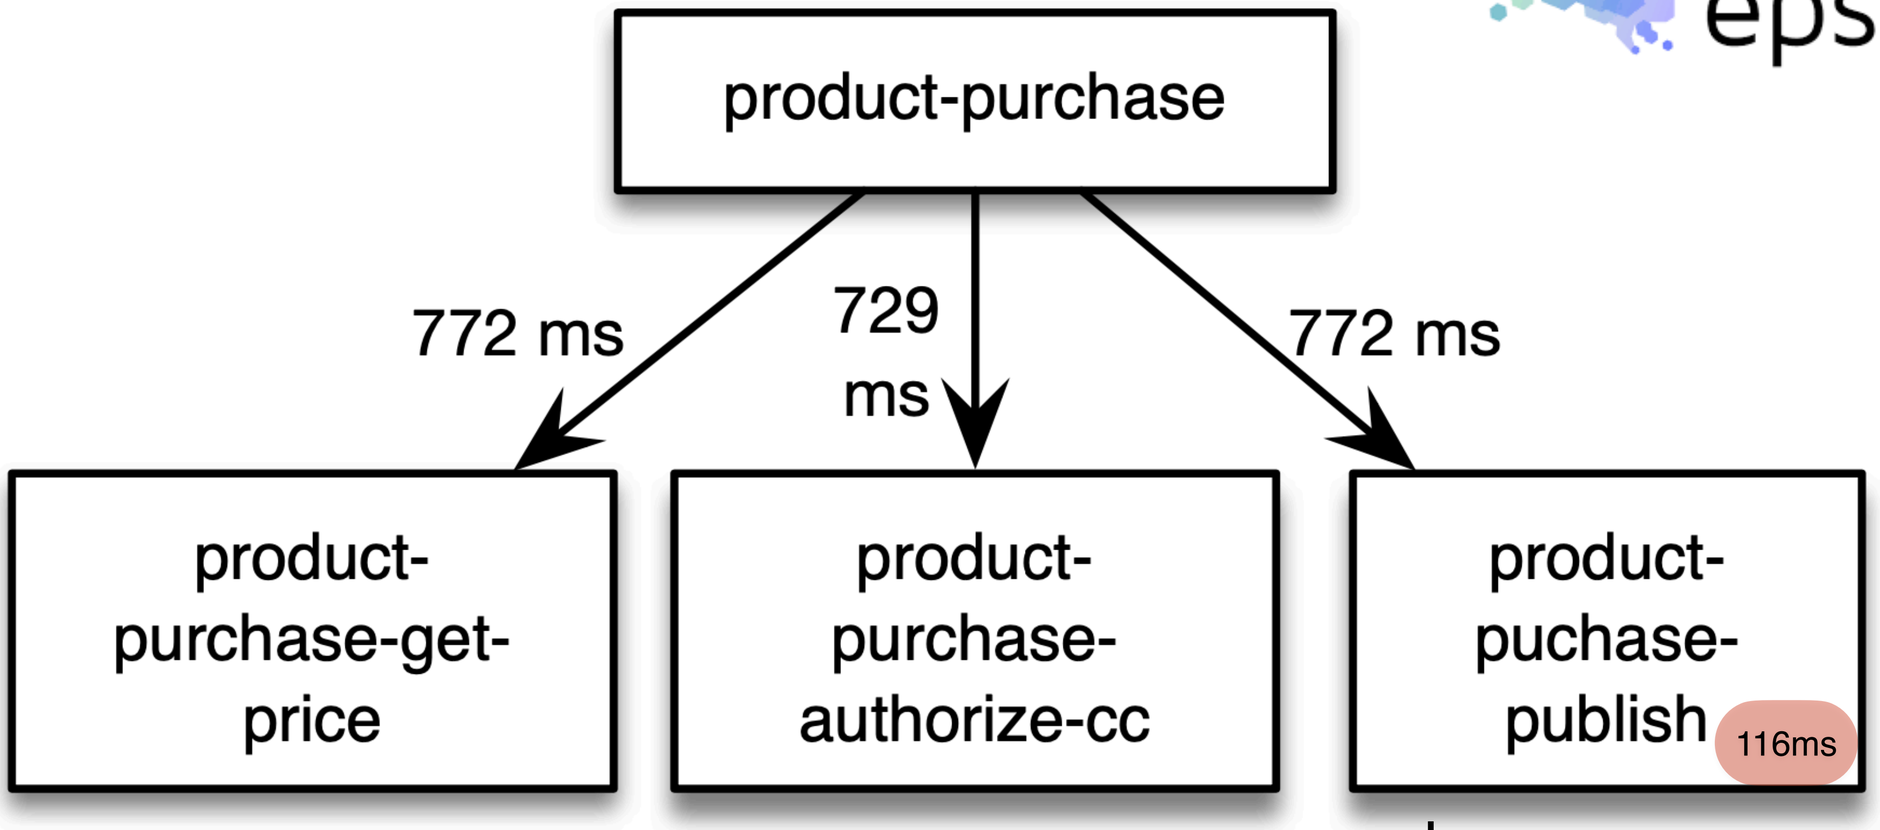

above average exec time

Attack Path Reconstruction Using Alastor

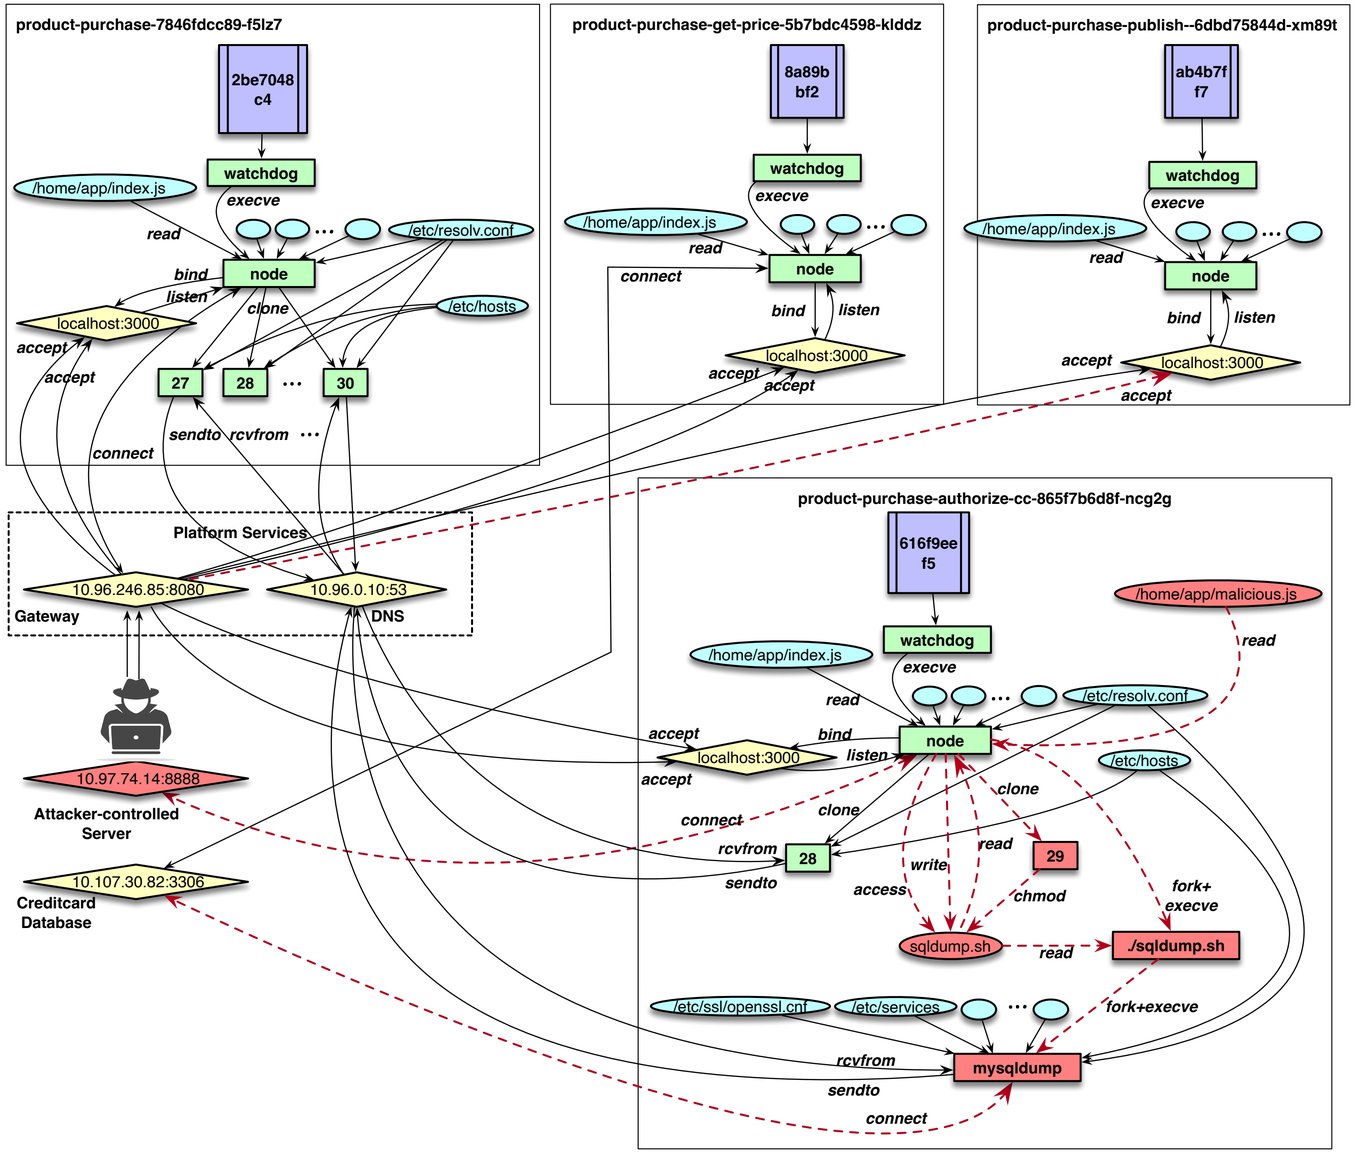

Attack Path Reconstruction Using Alastor

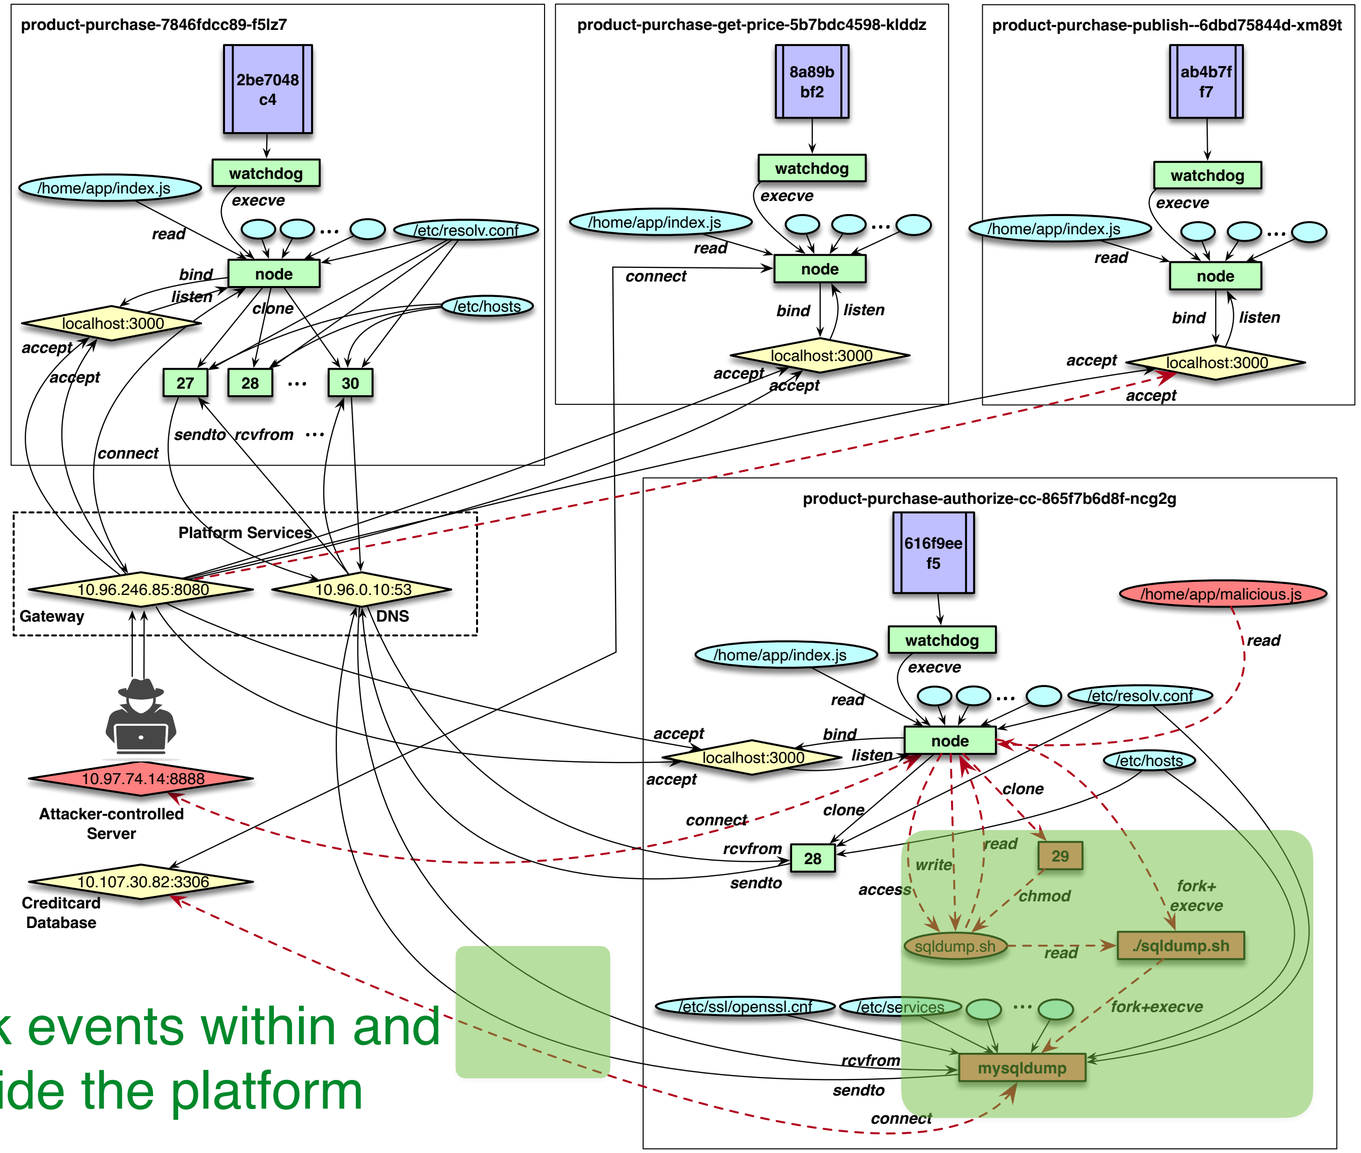

Network events within and outside the platform

Events inside function instance

Attack Path Reconstruction Using Alastor

Information on
each container spawned
each request served

Inter function
communications

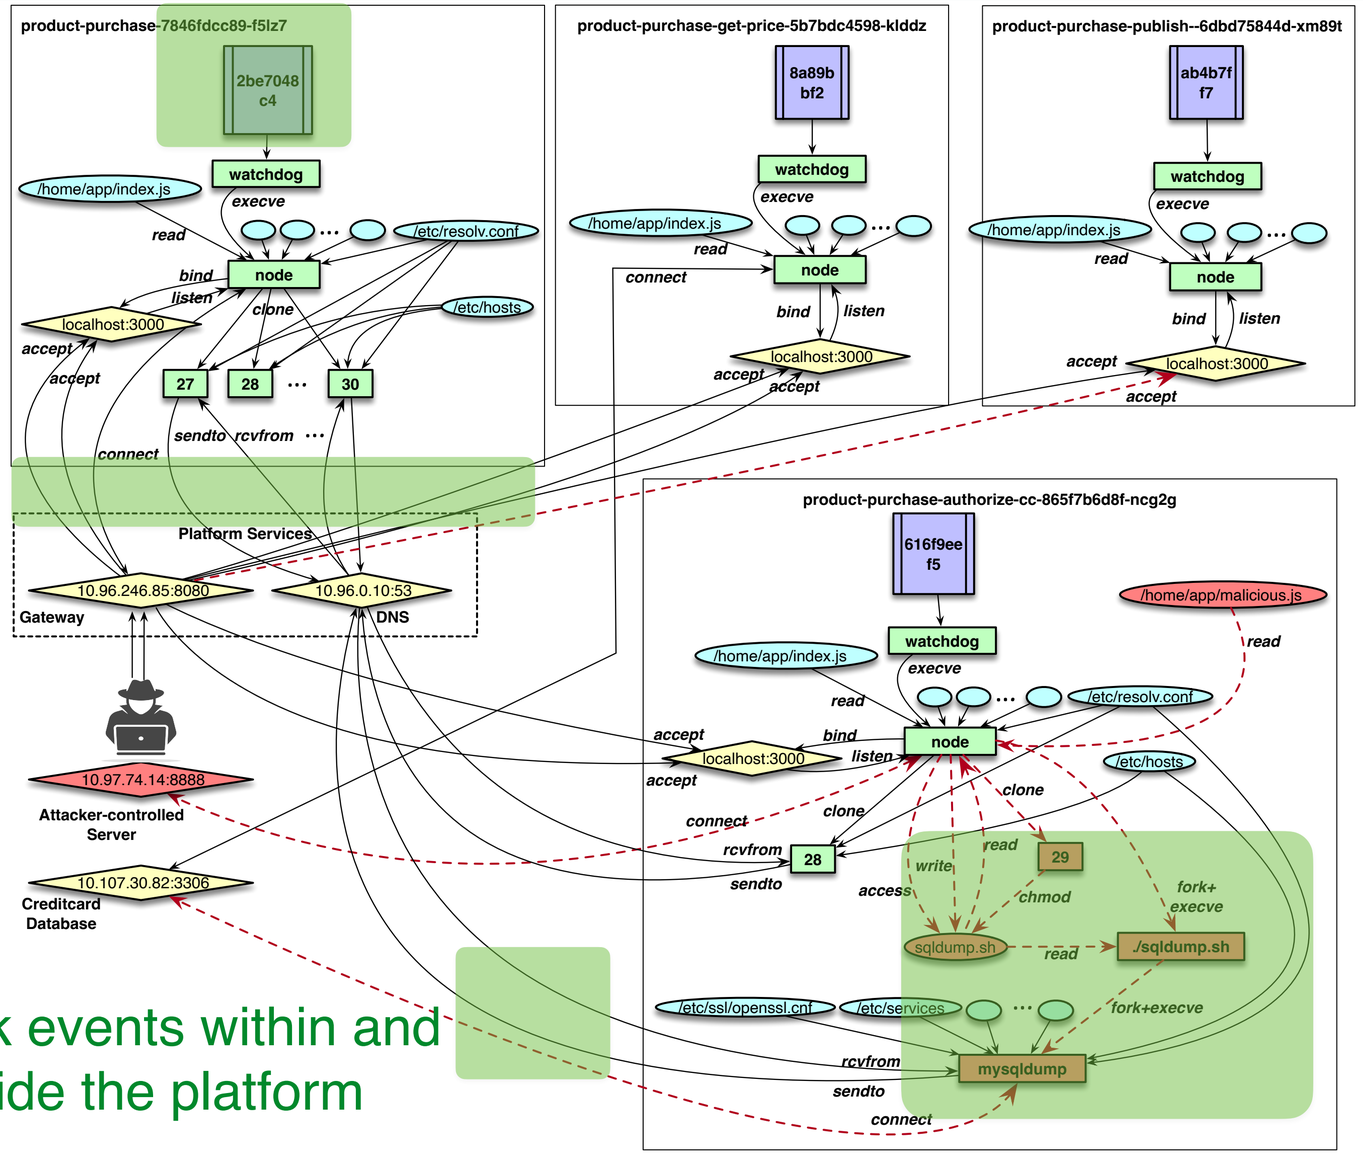

Network events within and
outside the platform

Events inside
function instance

Attack Path Reconstruction Using Alastor

Request 1

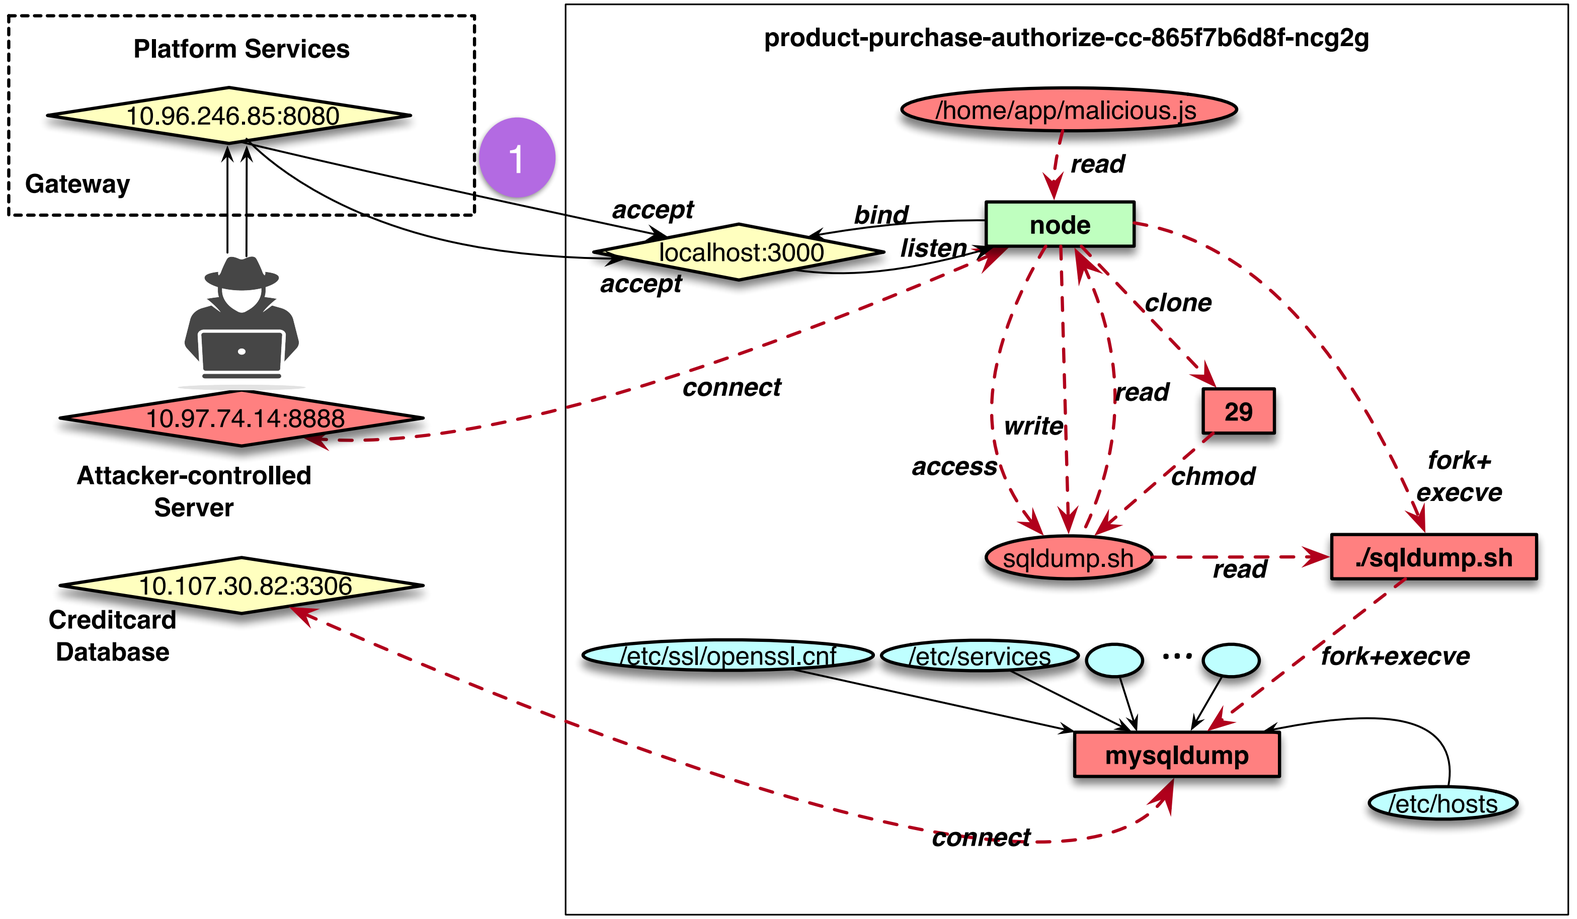

Attack Path Reconstruction Using Alastor

Request 1

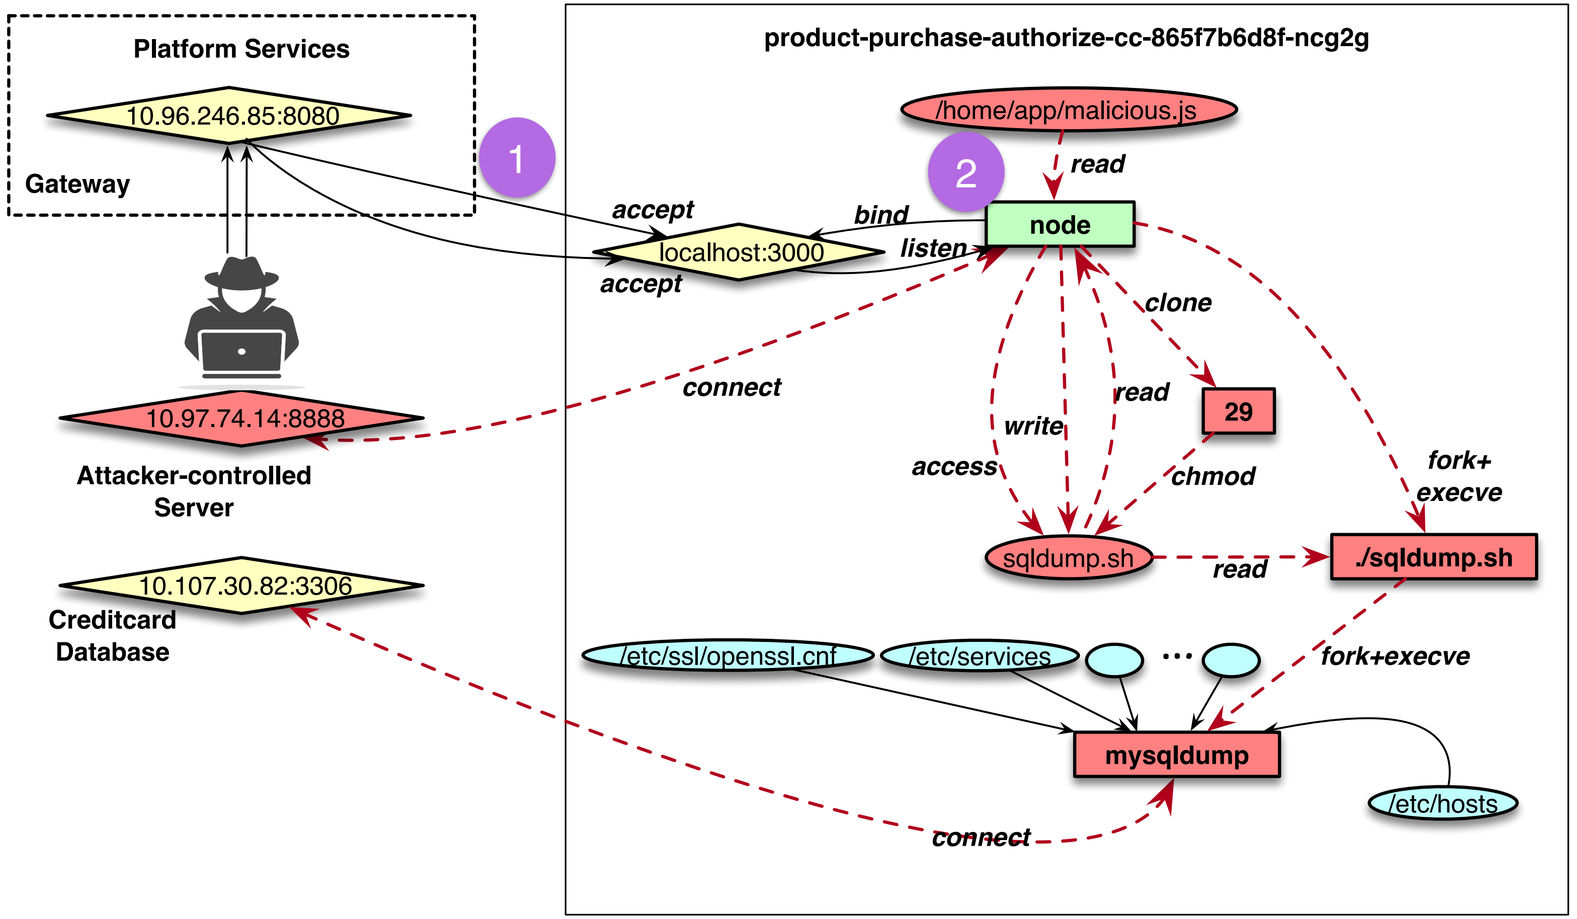

Attack Path Reconstruction Using Alastor

Request 1

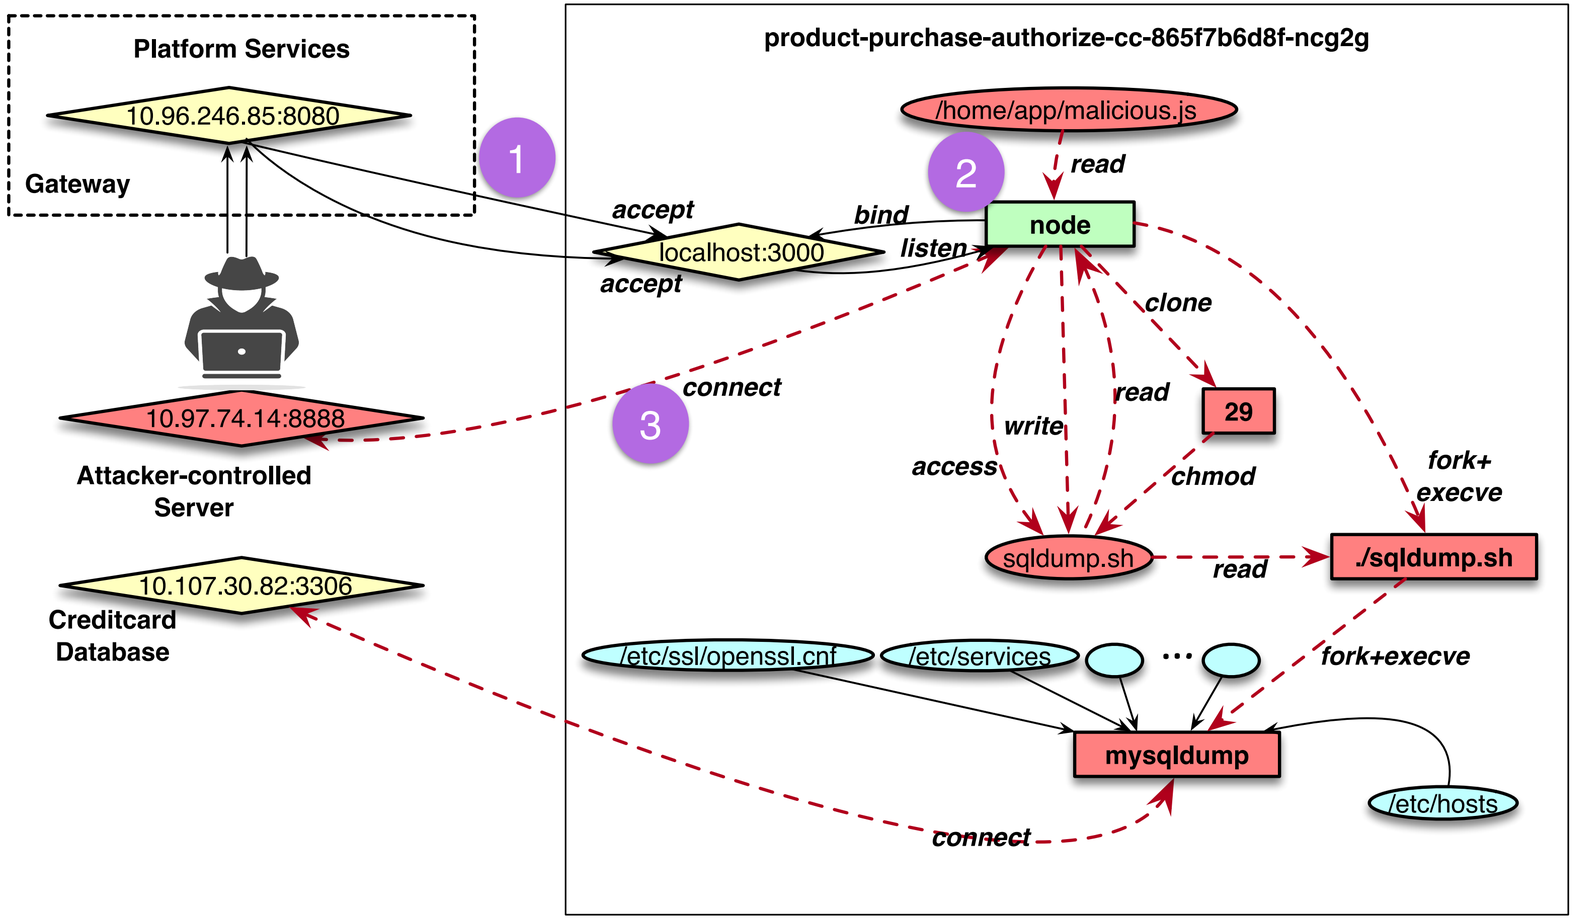

Attack Path Reconstruction Using Alastor

Request 1

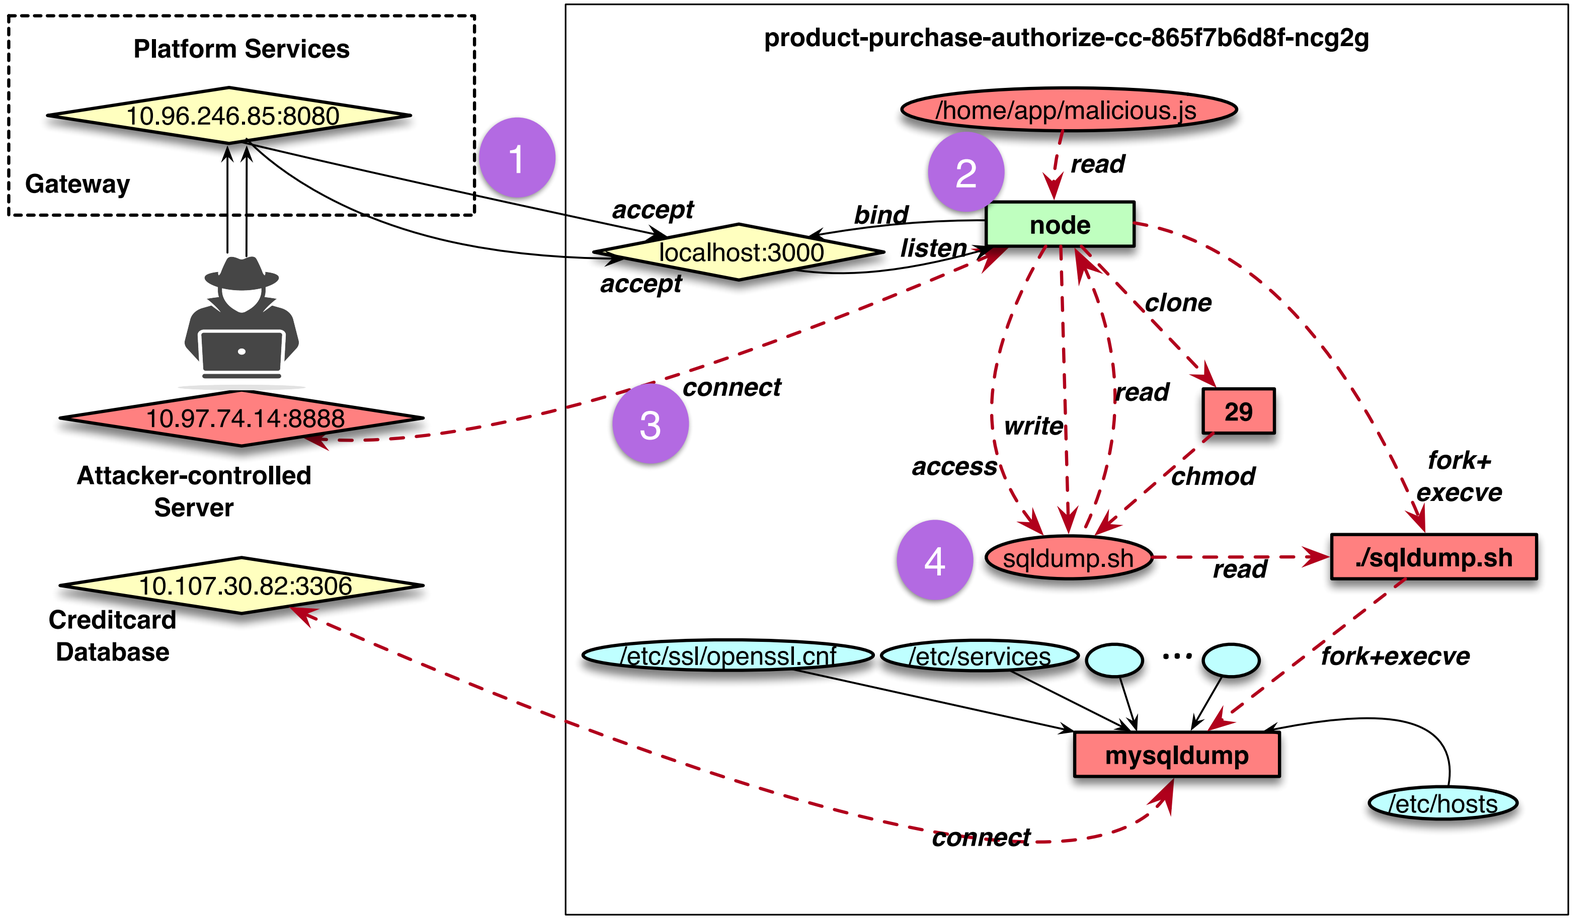

Attack Path Reconstruction Using Alastor

Request 1
Request 2

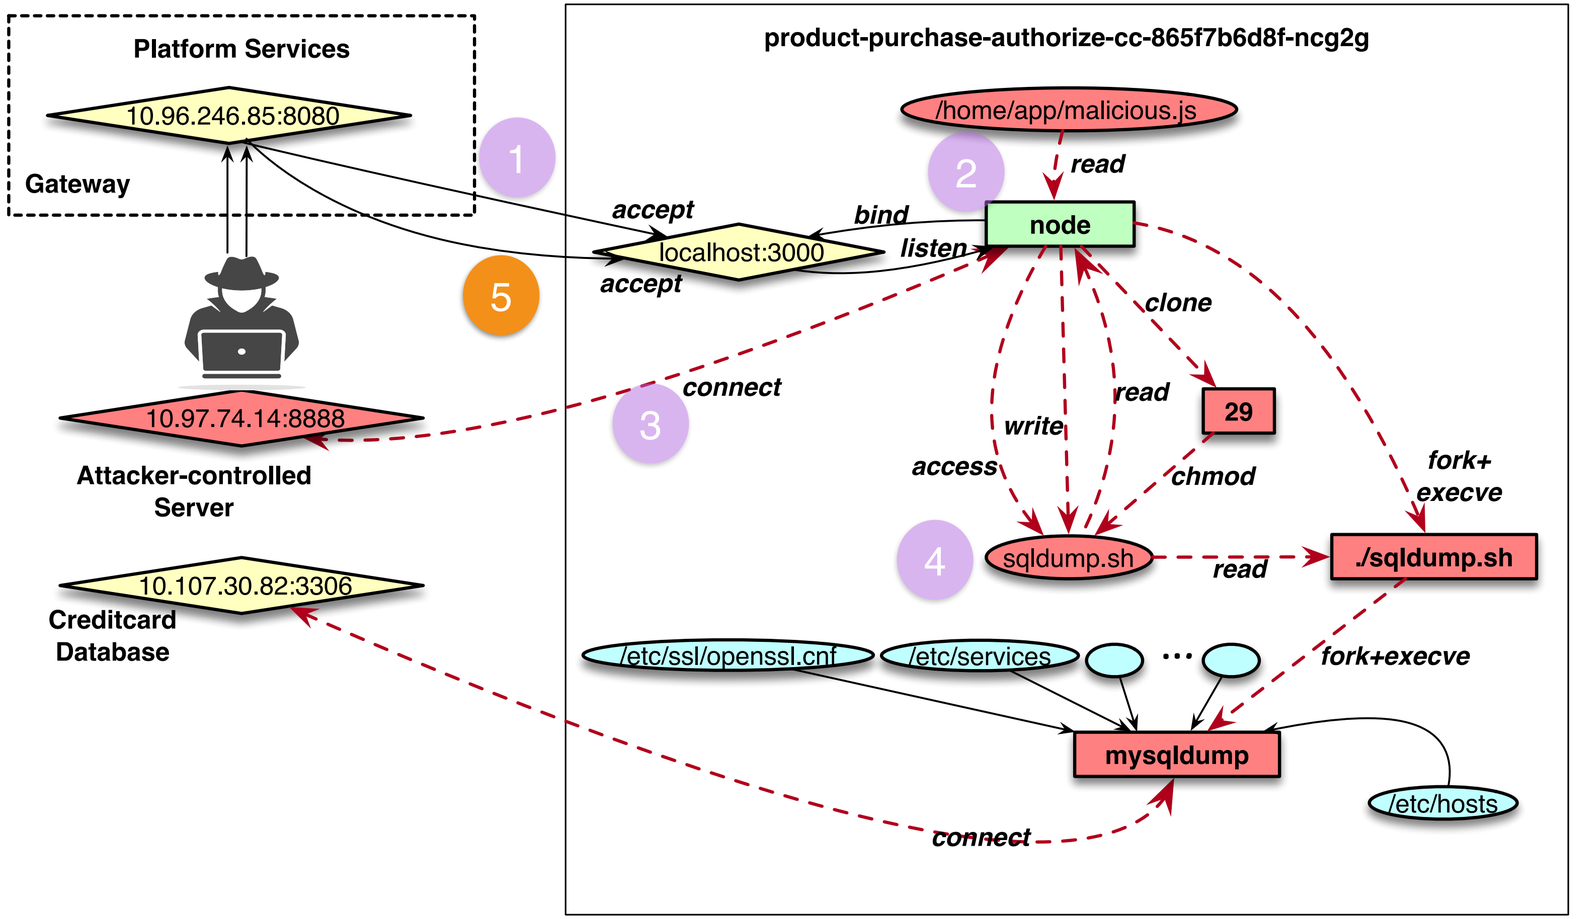

Attack Path Reconstruction Using Alastor

Request 1
Request 2

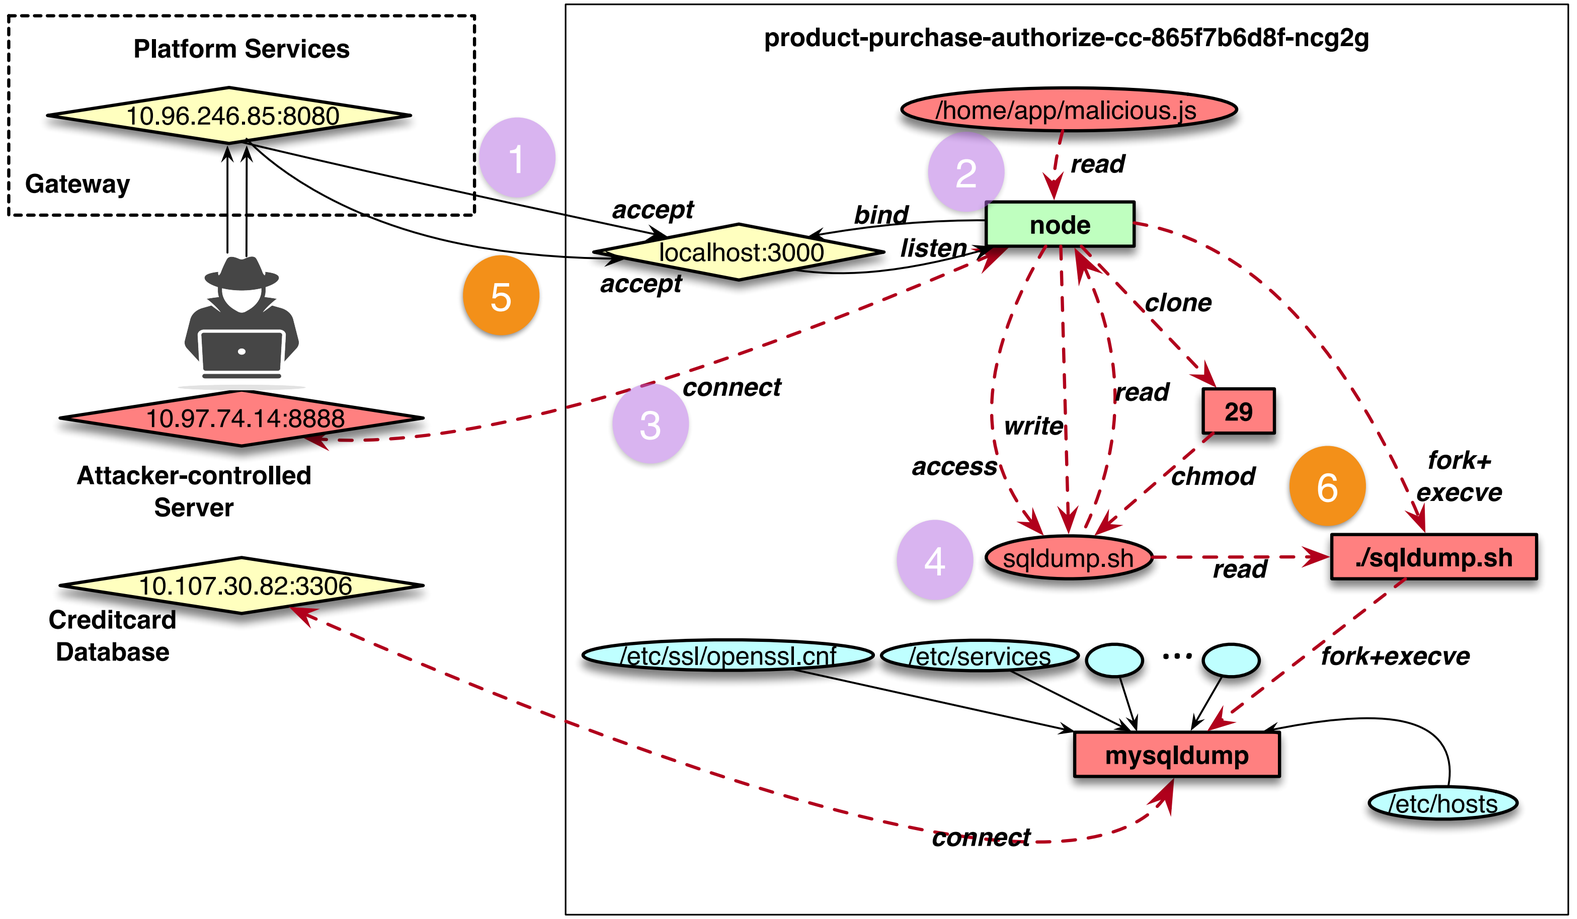

Attack Path Reconstruction Using Alastor

Request 1
Request 2

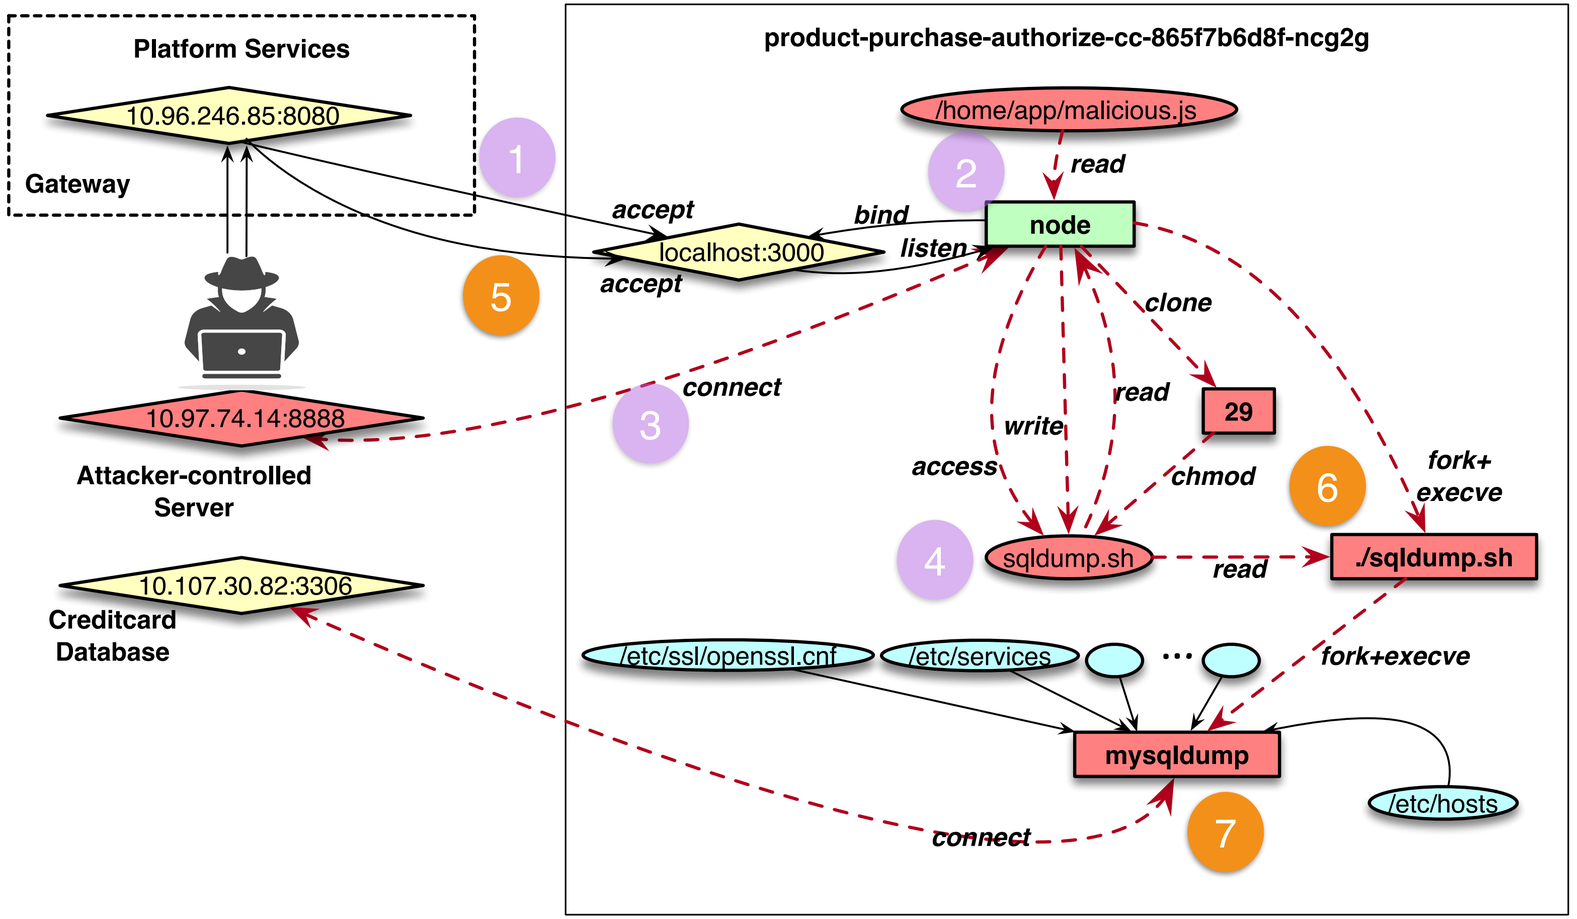

Attack Path Reconstruction Using Alastor

Container Reuse Attack

Request 1
Request 2

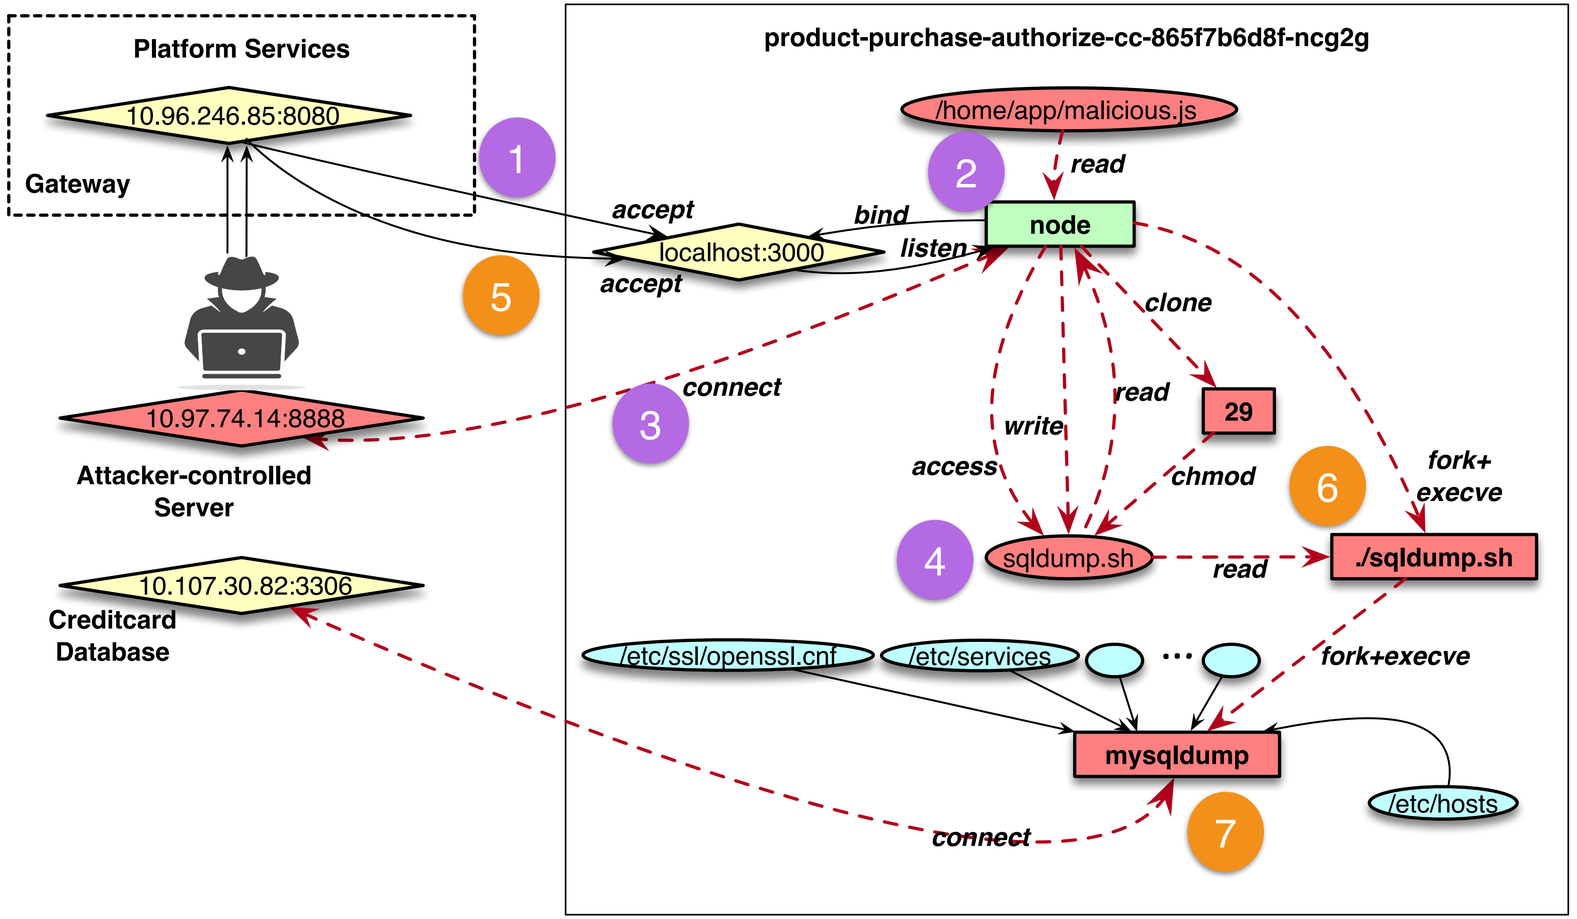

Conclusion

- Alastor can be used to diagnose the serverless-specific attack techniques like Container reuse and Exfiltration through function workflows.
- Alastor imposes manageable performance overhead of 13.74% in exchange for improved forensic capabilities.
- Alastor logs can be used with other IDS tools (e.g., Deeplog).

Thanks & Questions

pdatta2@illinois.edu

Intrusion Detection with DeepLog and Alastor

We trained Du et al's DeepLog on Alastor traces.

DeepLog operates on unstructured, free-text log entries.

	P	N
P	140	10
N	5	112

Metric	Values
Accuracy	0.944
Precision	0.966
Recall	0.933
F1	0.949

Complexity of Alastor Provenance Graphs

	Global Graph		Local Graph			
	Nodes	**Edges**	**Purchase**	**Get-price**	**Auth-CC**	**Publish**
Benign	69	80	N 18 E 28	N 13 E 13	N 21 E 14	N 13 E 13
Attack1	76	117	N 19 E 36	N 13 E 13	N 26 E 34	N 13 E 13
Attack2	49	65	N 18 E 29	N 13 E 13	-	N 13 E 13