Poster: Toward Zero-Trust Path-Aware Access Control

Joshua H. Seaton
Sena Hounsinou
jseaton@uccs.edu
University of Colorado Colorado Springs
Colorado Springs, CO, USA

Timothy Wood
George Washington University
Washington, DC, USA

Shouhuai Xu
Philip N. Brown
Gedare Bloom
University of Colorado Colorado Springs
Colorado Springs, CO, USA

ABSTRACT

In this poster, we introduce path-aware risk scores for access control (PARSAC), a novel context-sensitive technique to enrich access requests with risk scoring of the path taken by those requests between the authenticated user and the resources they access. These path-aware risk scores enable another layer of security for traditional access control systems that addresses the need for fine-grained monitoring and enforcement within a zero-trust architecture. We define rules for general functions that can be used to determine risk and instantiate a specific approach to calculate path risk scores. We evaluate our approach with realistic network graphs; PARSAC finds more paths with lower risk when compared with traditional routing algorithms that select the shortest path.

CCS CONCEPTS

• **Security and privacy → Access control**; **Distributed systems security**.

KEYWORDS

access control, zero-trust, path-aware networking

ACM Reference Format:
Joshua H. Seaton, Sena Hounsinou, Timothy Wood, Shouhuai Xu, Philip N. Brown, and Gedare Bloom. 2022. Poster: Toward Zero-Trust Path-Aware Access Control. In *SACMAT '22: ACM Symposium on Access Control Models and Technologies, June 08–10, 2022, New York City, NY*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3532105.3535036

1 INTRODUCTION

A fundamental new approach that aims to combat the threats posed by adversaries that penetrate boundary security is to adopt a *zero-trust* security model that moves the focus of defense from the network boundary to the access control system, i.e., to focus on principals and resources [10]. Organizations with a zero-trust architecture avoid trusting users or assets based on their location, i.e., whether or not they are within the boundary. The zero-trust paradigm represents a new way to think about how access control should work; Bertino [1] suggests that adoption of zero-trust architecture will necessitate adoption of attributes to handle the expected scale-out of fine-grained controls and associated policies. We agree and furthermore see the need for attributes to support the kinds of dynamic risk-based decisions that zero-trust architectures motivate.

In this poster, we introduce *path-aware risk scores for access control* (PARSAC) that builds on state-of-the-art advances being made in path-aware security for Internet routing [4, 6–8]. With path-aware security, applications can express enforceable security policies for path authorization, packet source authentication, and endpoint validation of path routing with strong cryptographic security guarantees and modest, practical performance overhead [6]. The open challenge, and our driving research question, is that it is not yet known: *what are the best mechanisms for end-to-end access control in path-aware network systems aiming to achieve zero-trust*? The contributions we make in this poster include:

- augmentation of access control to consider risks induced by the path that requests take from source to destination;
- mathematical formulation of PARSAC;
- evaluation of PARSAC risk scores for realistic network graphs.

2 PARSAC SYSTEM MODEL

Traditionally, distributed access control approaches [2, 5, 11] have relied on centralized policy engines where a single policy decision point (PDP) holds the authoritative access control policies while secondary decision points and policy enforcement points (PEP) may be distributed closer to the resources that are protected. In this work we assume a Dolev-Yao threat model with the added capability that an adversary has valid credentials for issuing requests to the PDP/PEP. Access control in the zero trust model does not allow for implicit trust for any user, device, data, or application until its request passes the PDP/PEP, which should be located in close proximity to each system resource to reduce the trusted domain to that resource and to allow the implementation of least privilege access and (granular) policies specific to each resource.

PARSAC provides a framework for assessing the risk associated with routing paths in two steps. First, PARSAC uses information about known trust relationships among nodes to compute a *node risk score* for each node along the routing path using a *node risk* function *NR* (see Definition 2.1). This score estimates the risk that the incoming traffic may have been modified in some adverse way as it passed through the node in question. Second, given the sequence of computed node risk scores along the routing path, PARSAC aggregates these node risk scores into a *path risk score* to capture the overall risk of tampering associated with the specific network routing path using a *path risk* function *PR* (Definition 2.3). This path risk score represents the aggregate risk associated with the given network path; a high path risk score may be grounds for denial of access to sensitive resources.

SACMAT '22, June 08–10, 2022, New York City, NY

ACM ISBN 978-1-4503-9357-7/22/06.
https://doi.org/10.1145/3532105.3535036

2.1 Network Model

A network model $G = (V, E_r, E_o, T)$ comprises a set V of nodes and a set $E_r \subseteq V \times V$ of physical network links connecting pairs of nodes in V. If $(i, j) \in E_r$, then nodes i and j can send traffic directly to each other. The graph (V, E_r) represents the network's physical topology, which we call the *routing graph.* We assume that this routing graph is undirected, connected, and simple. Let $E_o \subseteq V \times V \times T$ denote the set of *overlay links* corresponding to the directed trust relationships between nodes, where $(i, j, t) \in E_o$ means that node i assigns node j a risk scoreof $t \in (0, \bar{t}\,] \subseteq (0, 1)$ given *untrusted risk* score $\bar{t} \in (0, 1)$. The risk score t represents node i's assessment of the risk associated with network traffic passing through j; higher values of t correspond with a higher risk. The untrusted risk $\bar{t}$ corresponds to the risk associated with a totally unknown node. Thus, (V, E_o) represents a directed, weighted graph of trust relationships among the node. This graph is directed and simple, but need not be connected, fully known, or static.

2.2 Node-level Risk Score Assignment

To assess the risk of a routing path, PARSAC first computes a risk score for each node on the packet's routing path based on its relationship with the destination node u. Let u be a network packet's destination node and v be a node on a routing path that node u assesses. Node u aims to assign a risk score to node v; in the event that u does not know v directly (i.e., $(u, v, t) \notin E_o$), this risk assignment is done on the basis of the relationships encoded in the overlay network (V, E_o) as follows and as depicted in Figure 1a. Let τ be a directed path in the overlay network (V, E_o) from node u to node v, where we treat τ explicitly as its associated sequence of risk scores in the overlay. Let $\mathcal{T}$ be the set of all sequences of risk scores of bounded length. A *node risk function NR* is a function $NR : \mathcal{T} \to (0, \bar{t}\,]$ that maps directed paths to risk scores, so that $NR(\tau)$ represents the risk score assigned by node u to node v on the basis of overlay path τ. We use the symbol t to refer to an individual risk score, and τ to refer to a sequence of risk scores (an overlay path). Given overlay paths τ and τ', we write (τ, τ') to denote the concatenation of τ and τ', and slightly abuse notation by writing (τ, t) to denote the risk score t concatenated to the end of overlay path τ. We define a "natural" node risk function as:

Definition 2.1. A function $NR : \mathcal{T} \to (0, \bar{t}\,]$ is an *admissible node risk function* if it satisfies all of the following properties:

- **Brevity:** If path τ has only one edge with score t, the edge's score is returned: $NR(\tau) = t$.
- **Extension:** Appending an edge to the end of a path cannot decrease risk: for every path τ and any risk score t, it holds that $NR((\tau, t)) \geq NR(\tau)$.
- **Decay:** Close relationships confer no more risk than distant relationships: if $t_i \geq t_j$ and τ_1, τ_2, and τ_3 are any (possibly empty) paths, then for paths $\tau = (\tau_1, t_i, \tau_2, t_j, \tau_3)$ and $\tau' = (\tau_1, t_j, \tau_2, t_i, \tau_3)$, it holds that $NR(\tau) \geq NR(\tau')$.

Brevity ensures that the risk scores in the overlay network are consistent with actual trust relationships among nodes, Extension ensures that no path can be falsely sanitized by appending spurious low-risk path elements, and Decay captures the natural property that I should not trust the friend of a stranger more than I trust someone who is a stranger to my friend.

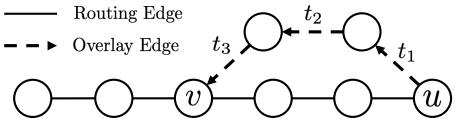

(a) Node u receives a request along a routing path (solid lines) and computes a risk score for each node along that path based on the overlay graph (dashed lines). The risk score that node u assigns to node v is $NR(t_1, t_2, t_3)$.

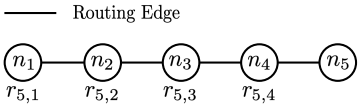

(b) The destination node n_5 first uses NR to compute a risk score $r_{5,i}$ for each node on the routing path before computing aggregate risk as $PR(r_{5,1}, r_{5,2}, r_{5,3}, r_{5,4})$.

Figure 1: Calculating node risk and path risk.

Now, let $\mathcal{T}_{u,v}$ denote the set of shortest overlay paths from u to v, and let NR be an admissible node risk function. The *node risk score* $R_{u,v}$ assigned by node u to node v is given by

$$R_{u,v} = \begin{cases} \min_{\tau \in \mathcal{T}_{u,v}} NR(\tau) & \text{if } \mathcal{T}_{u,v} \neq \emptyset \\ \bar{t} & \text{otherwise.} \end{cases} \tag{1}$$

That is, the risk score assigned by u to v is the *lowest* of that of all shortest paths. If no overlay path exists from u to v (i.e., no information is available about node v), the untrusted risk $\bar{t}$ is assigned.

Proposition 2.2. *The following functions are all admissible node risk functions satisfying Definition 2.1. In the following, $\tau = (t_1, t_2, \ldots, t_k)$ denotes an arbitrary directed overlay path of length k.*

- *Maximum:* $NR_{\max}(\tau) = \max_{t \in \tau} t$
- *Maximum with length:* $NR_{\mathrm{ml}}(\tau) = \min\{\bar{t}, k \max_{t \in \tau} t\}$
- *Sum with saturation:* $NR_{\mathrm{ss}}(\tau) = \min\{\bar{t}, \sum_{t \in \tau} t\}$
- *Max with order penalty using decay parameter $\alpha < 1$:* $NR_{\mathrm{mop}}(\tau; \alpha) = \max_{i \in \{1, \ldots, k\}} \left(\bar{t} - \alpha^{k-i}(\bar{t} - t_i)\right)$.

2.3 Path-level Risk Score Assignment

Given a physical network path $p = (n_1, n_2, \ldots, n_k)$, for each intermediate node $i \in \{1, \ldots, k-1\}$, the destination node k computes its node risk score $r_{k,i}$ according to Eq.(1) using an admissible node risk function. Let $\rho_p := (r_{k,1,}, r_{k,2}, \ldots, r_{k,k-1})$ denote the sequence of risk scores along path p computed by the destination node k. Let $\mathcal{T}$ be the set of all sequences of risk scores of bounded length. Analogous to the node risk function in Section 2.2, a *path risk function* is a function $PR : \mathcal{T} \to (0, 1)$ that aggregates sequences of node risk scores to a single path risk score, which can then be used to condition access control decisions by the destination node k. Figure 1b depicts a generic example of path risk computation.

Throughout, ρ denotes a sequence of node risk scores, and $r_{k,i}$ denotes the individual risk score of node i as assessed by destination k (we often omit the subscript k when clear from context). Given risk sequences ρ and ρ', we write (ρ, ρ') to denote the concatenation

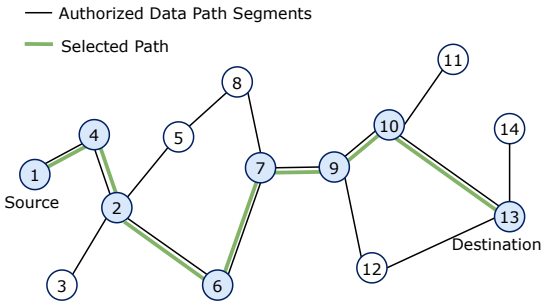

(a) Paths authorized in the Control Plane.

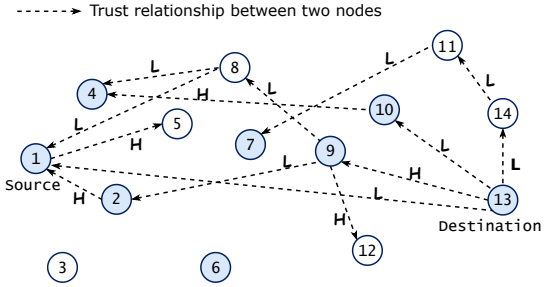

(b) Overlay of trust relationships.

Figure 2: Network Model. Although the source node is low-risk (L), the packet can only be forwarded along authorized segments depicted in Figure 2a (thick green line). The path selected for the packet is $((1,4),(4,2),(2,6),(6,7),(7,9),(9,10),(10,13))$**.**

of ρ and ρ', and slightly abuse notation by writing (ρ, r) to denote the risk score r concatenated to the end of risk sequence ρ.

As with node risk, we provide a simple axiomatic definition of a "natural" path risk aggregation function followed by a simple, admissible path risk function in Proposition 2.4.

Definition 2.3. A function $PR : \mathcal{T} \to (0,1)$ is an *admissible path risk function* if it satisfies all of the following properties. In the following, let ρ_1 and ρ_2 be any sequences of risk scores, at most one of which is empty.

- **Brevity:** If node risk sequence $\rho = (r)$ for some $r \in (0,1)$ (i.e., ρ has only one entry), that entry is returned: $PR(\rho) = r$.
- **Nesting:** The risk of a path is no less than the risk of any subpath: if ρ is any non-empty risk sequence, then $PR(\rho) < PR(\rho_1, \rho, \rho_2)$.
- **Monotonicity:** Increasing (decreasing) the risk of a single node increases (decreases) path risk: if $r < r'$, then for paths $\rho = (\rho_1, r, \rho_2)$ and $\rho' = (\rho_1, r', \rho_2)$, it holds that $PR(\rho) < PR(\rho')$.

Proposition 2.4. *Let ρ denote an arbitrary sequence of node risk scores. If we define the* path probability *path risk function as $PR_{\text{pp}}(\rho) = 1 - \prod_{r_i \in \rho}(1 - r_i)$, it holds that PR_{pp} satisfies Definition 2.3.*

Note that several simple functions such as average risk score or maximum risk score are not admissible path risk functions.

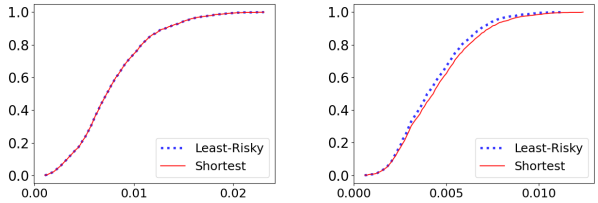

(a) Minimal degree **(b) Medium degree**

Figure 3: CDF of path risks with varying degrees.

3 CONCLUSION

We introduced PARSAC as a novel method for access control in the context of zero-trust. Future work can investigate composing PARSAC with access control models, characterize the impact PARSAC has on forensic analyses [3], enrich PARSAC's risk scoring methodology with threat intelligence and security analytics [9], and implement and evaluate PARSAC for practical application.

ACKNOWLEDGMENTS

Project sponsored by the National Security Agency under Grant Number H98230-21-1-0155. The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein.

REFERENCES

[1] Elisa Bertino. 2021. Zero Trust Architecture: Does It Help? *IEEE Security & Privacy* 19, 05 (Sept. 2021), 95–96. Publisher: IEEE Computer Society.

[2] Gedare Bloom and Rahul Simha. 2014. Hardware-enhanced distributed access enforcement for role-based access control. In *Proceedings of the 19th ACM symposium on Access control models and technologies (SACMAT '14)*. ACM, New York, NY, USA, 5–16.

[3] Nahid Juma, Xiaowei Huang, and Mahesh Tripunitara. 2020. Forensic Analysis in Access Control: Foundations and a Case-Study from Practice. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. ACM, New York, NY, USA, 1533–1550.

[4] Tobias Klenze, David Basin, and Christoph Sprenger. 2021. Formal Verification of Secure Forwarding Protocols. In *34th IEEE Computer Security Foundations Symposium (CSF 2021)*.

[5] Marko Komlenovic, Mahesh Tripunitara, and Toufik Zitouni. 2011. An empirical assessment of approaches to distributed enforcement in role-based access control (RBAC). In *Proceedings of the first ACM conference on Data and application security and privacy (CODASPY '11)*. ACM, New York, NY, USA, 121–132.

[6] Markus Legner, Tobias Klenze, Marc Wyss, Christoph Sprenger, and Adrian Perrig. 2020. EPIC: Every Packet Is Checked in the Data Plane of a Path-Aware Internet. In *29th USENIX Security Symposium*. USENIX Association, 541–558.

[7] Guyue Liu, Hugo Sadok, Anne Kohlbrenner, Bryan Parno, Vyas Sekar, and Justine Sherry. 2021. Don't Yank My Chain: Auditable NF Service Chaining. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*. USENIX Association.

[8] Jad Naous, Michael Walfish, Antonio Nicolosi, David Mazieres, Michael Miller, and Arun Seehra. 2011. Verifying and enforcing network paths with ICING. In *Proceedings of the Seventh Conference on Emerging Networking Experiments and Technologies*. 1–12.

[9] Isaac Polinsky, Pubali Datta, Adam Bates, and William Enck. 2021. SCIFFS: Enabling Secure Third-Party Security Analytics using Serverless Computing. In *Proceedings of the 26th ACM Symposium on Access Control Models and Technologies (SACMAT '21)*. ACM, New York, NY, USA, 175–186.

[10] Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly. 2020. *Zero Trust Architecture*. Technical Report. National Institute of Standards and Technology.

[11] Mahesh V. Tripunitara and Bogdan Carbunar. 2009. Efficient access enforcement in distributed role-based access control (RBAC) deployments. In *Proceedings of the 14th ACM symposium on Access control models and technologies (SACMAT '09)*. ACM, New York, NY, USA, 155–164.