PRACTITIONER TOOLKIT

Reference List

Maintained version of the book's sources, as of September 28, 2026

Every source cited in the book's chapter notes, listed once and grouped by type. Chapter numbers show where each source is cited. Standards and products in this field change quickly; corrections and newer versions are recorded in the change log at arvikpress.com/bounding-agency.

Standards and Specifications

A2A Project. "Agent2Agent (A2A) Protocol Specification." Version 1.0. https://github.com/a2aproject/A2A/blob/main/docs/specification.md. Accessed September 28, 2026. [Chapters 4, 5]

Campbell, Brian, John Bradley, Nat Sakimura, and Torsten Lodderstedt. "OAuth 2.0 Mutual-TLS Client Authentication and Certificate-Bound Access Tokens." RFC 8705. IETF, February 2020. https://www.rfc-editor.org/rfc/rfc8705. [Chapter 5]

Fett, Daniel, Brian Campbell, John Bradley, Torsten Lodderstedt, Michael Jones, and David Waite. "OAuth 2.0 Demonstrating Proof of Possession (DPoP)." RFC 9449. IETF, September 2023. https://www.rfc-editor.org/rfc/rfc9449. [Chapter 5]

IETF WIMSE Working Group. "AI Identity Management System." draft-ietf-wimse-aims-00. IETF, September 15, 2026. https://datatracker.ietf.org/doc/draft-ietf-wimse-aims/. Accessed September 28, 2026. Internet-Draft, work in progress. [Chapters 4, 5, 7, 9]

ISO/IEC 42001:2023. *Information technology — Artificial intelligence — Management system*. Geneva: ISO, 2023. Voluntary standard. [Chapter 11]

Jones, Michael B., Anthony Nadalin, Brian Campbell, John Bradley, and Chuck Mortimore. "OAuth 2.0 Token Exchange." RFC 8693. IETF, January 2020. https://www.rfc-editor.org/rfc/rfc8693. [Chapters 2, 7]

Kasselman, Pieter, Jeff Lombardo, Yaroslav Rosomakho, Brian Campbell, Nick Steele, and Aaron Parecki. "AI Agent Authentication and Authorization." draft-klrc-aiagent-auth-03. IETF, July 2026. https://datatracker.ietf.org/doc/draft-klrc-aiagent-auth/. Internet-Draft, work in progress. [Chapters 4, 5]

Lodderstedt, Torsten, John Bradley, Andrey Labunets, and Daniel Fett. "Best Current Practice for OAuth 2.0 Security." RFC 9700 (BCP 240). IETF, January 2025. https://www.rfc-editor.org/rfc/rfc9700. [Chapter 5]

Lodderstedt, Torsten, Justin Richer, and Brian Campbell. "OAuth 2.0 Rich Authorization Requests." RFC 9396. IETF, May 2023. https://www.rfc-editor.org/rfc/rfc9396. [Chapter 6]

Lodderstedt, Torsten, Stefanie Dronia, and Marius Scurtescu. "OAuth 2.0 Token Revocation." RFC 7009. IETF, August 2013. [Chapters 5, 6, 10]

Model Context Protocol. "Authorization." Specification revision 2026-07-28. https://modelcontextprotocol.io/specification/2026-07-28/basic/authorization. Accessed September 28, 2026. [Chapter 5]

Model Context Protocol. "Key Changes." 2025-11-25 specification changelog. https://modelcontextprotocol.io/specification/2025-11-25/changelog. Accessed September 28, 2026. [Chapter 15]

OpenID Foundation. "Authorization API 1.0 Final Specification Approved." January 12, 2026. https://openid.net/authorization-api-1-0-final-specification-approved/. [Chapter 6]

OpenID Foundation. "Three Shared Signals Final Specifications Approved." September 2, 2025. https://openid.net/three-shared-signals-final-specifications-approved/. [Chapters 5, 10]

Richer, Justin. "OAuth 2.0 Token Introspection." RFC 7662. IETF, October 2015. [Chapters 5, 6, 10]

Schwenkschuster, Arndt, Pieter Kasselman, Kelley Burgin, Michael J. Jenkins, Brian Campbell, and Aaron Parecki. "OAuth Identity and Authorization Chaining Across Domains." IETF OAuth Working Group, July 2026. https://datatracker.ietf.org/

doc/draft-ietf-oauth-identity-chaining/. Internet-Draft, approved for publication; in the RFC Editor queue as of September 2026. [Chapter 7]

SPIFFE Project. “SPIFFE Overview.” https://spiffe.io/docs/latest/spiffe-about/overview/. Accessed September 28, 2026. [Chapters 4, 5]

Tulshibagwale, Atul, George Fletcher, and Pieter Kasselman. “Transaction Tokens.” draft-ietf-oauth-transaction-tokens-11. IETF OAuth Working Group, July 30, 2026. https://datatracker.ietf.org/doc/draft-ietf-oauth-transaction-tokens/. Internet-Draft, work in progress; not an RFC. [Chapter 6]

Government, Regulatory and Public-Sector Sources

Booth, Harold, Bill Fisher, Ryan Galluzzo, and Joshua Roberts. “Accelerating the Adoption of Software and AI Agent Identity and Authorization.” National Cybersecurity Center of Excellence (NCCoE), NIST, February 2026. https://www.nccoe.nist.gov/publications/other/accelerating-adoption-software-and-ai-agent-identity-and-authorization-concept. PDF, https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf. Accessed September 28, 2026. Concept paper, not guidance. [Chapters 1, 11]

Chandramouli, Ramaswamy, and Zack Butcher. *A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Location Environments*. NIST SP 800-207A. September 2023. [Chapter 6]

Cybersecurity and Infrastructure Security Agency (CISA). *Zero Trust Maturity Model*. Version 2.0. April 2023. https://www.cisa.gov/zero-trust-maturity-model. PDF, https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf. [Chapters 3, 12, 13, 14]

Directive (EU) 2022/2555 (NIS2). *Official Journal of the European Union*, L 333, December 27, 2022. https://eur-lex.europa.eu/eli/dir/2022/2555/oj. [Chapter 11]

European Commission. “AI Omnibus Enters into Force.” July 2026. https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force. Accessed September 28, 2026. [Chapter 11]

Fisher, Bill, and Ryan Galluzzo. “Back to the Future: Why Agentic AI Needs a Strong Identity Foundation.” *Cybersecurity Insights* (blog), NIST, August 27, 2026. https://www.nist.gov/blogs/cybersecurity-insights/back-future-why-agentic-ai-needs-strong-identity-foundation. Accessed September 27, 2026. [Chapters 1, 3]

Hu, Vincent C., et al. *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*. NIST Special Publication 800-162. January 2014 (updated August 2019). https://doi.org/10.6028/NIST.SP.800-162. [Chapter 6]

Infocomm Media Development Authority (IMDA), Singapore. *Model AI Governance Framework for Agentic AI*. January 22, 2026 (updated May 20, 2026). https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf. Non-binding government guidance. [Chapters 4, 6, 8, 10, 11, 12, 15, 16]

National Cybersecurity Center of Excellence (NCCoE). “Software and AI Agent Identity and Authorization.” Project page. NIST. https://www.nccoe.nist.gov/projects/software-and-ai-agent-identity-and-authorization. Accessed September 28, 2026. [Chapters 1, 3, 4, 9]

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. January 2023. https://doi.org/10.6028/NIST.AI.100-1. Voluntary framework. [Chapter 11]

National Institute of Standards and Technology. *The NIST Cybersecurity Framework (CSF) 2.0*. NIST CSWP 29. February 2024. https://doi.org/10.6028/NIST.CSWP.29. Voluntary framework. [Chapters 11, 12]

NIST National Vulnerability Database. “CVE-2025-32711.” https://nvd.nist.gov/vuln/detail/CVE-2025-32711. Accessed September 28, 2026. [Chapters 1, 10]

Regulation (EU) 2024/1689 (Artificial Intelligence Act). *Official Journal of the European Union*, L series, July 12, 2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj. [Chapter 11]

Regulation (EU) 2026/1744. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202601744. Accessed September 28, 2026. [Chapter 11]

Rose, Scott, Oliver Borchert, Stu Mitchell, and Sean Connelly. *Zero Trust Architecture*. NIST Special Publication 800-207. August 2020. https://doi.org/10.6028/NIST.SP.800-207. [Chapters 6, 13]

Industry and Community Guidance

Cloud Security Alliance. “Agentic AI Identity and Access Management: A New Approach.” August 18, 2025. https://cloudsecurityalliance.org/artifacts/agentic-ai-identity-and-access-management-a-new-approach. Accessed September 27, 2026. [Chapters 3, 5, 7, 9, 13]

Cloud Security Alliance. “Introducing the CSA AI Controls Matrix.” July 10, 2025. https://cloudsecurityalliance.org/blog/2025/07/10/introducing-the-csa-ai-controls-matrix-a-comprehensive-framework-for-trustworthy-ai. Voluntary control framework. [Chapter 11]

Huang, Ken. “Agentic AI Threat Modeling Framework: MAESTRO.” Cloud Security Alliance, February 6, 2025. https://cloudsecurityalliance.org/blog/2025/02/06/agentic-ai-threat-modeling-framework-maestro. Accessed September 28, 2026. [Chapter 10]

MITRE. *ATLAS (Adversarial Threat Landscape for Artificial-Intelligence Systems)*. https://atlas.mitre.org/. Accessed September 28, 2026. [Chapter 10]

OWASP GenAI Security Project. “LLM06:2025 Excessive Agency.” *OWASP Top 10 for LLM Applications 2025*. https://genai.owasp.org/llmrisk/llm062025-excessive-agency/. [Chapter 10]

OWASP GenAI Security Project. “OWASP Top 10 for Agentic Applications.” December 9, 2025. https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/. Announcement, https://genai.owasp.org/2025/12/09/owasp-top-10-for-agentic-applications-the-benchmark-for-agentic-security-in-the-age-of-autonomous-ai/. Accessed September 28, 2026. [Chapters 2, 6, 10]

OWASP. “Non-Human Identities Top 10 – 2025.” https://owasp.org/www-project-non-human-identities-top-10/. Accessed September 28, 2026. [Chapters 4, 8, 10]

South, Tobin, et al. *Identity Management for Agentic AI: The New Frontier of Authorization, Authentication, and Security for an AI Agent World*. OpenID Foundation whitepaper, October 2025. https://openid.net/wp-content/uploads/2025/10/Identity-Management-for-Agentic-AI.pdf. Also arXiv:2510.25819. [Chapters 6, 7, 8, 9, 10, 13]

Research

Birgisson, Arnar, Joe Gibbs Politz, Úlfar Erlingsson, Ankur Taly, Michael Vrable, and Mark Lentczner. “Macaroons: Cookies with Contextual Caveats for Decentralized Authorization in the Cloud.” *Network and Distributed System Security Symposium (NDSS)*, 2014. [Chapters 6, 7]

Dalugoda, Asiri. “HDP: A Lightweight Cryptographic Protocol for Human Delegation Provenance in Agentic AI Systems.” arXiv:2604.04522, March 2026. https://arxiv.org/abs/2604.04522. Research preprint, not peer-reviewed. [Chapter 7]

Greshake, Kai, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. “Not What You’ve Signed Up For: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection.” In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security (AISec ’23)*. New York: ACM, 2023. https://doi.org/10.1145/3605764.3623985. [Chapter 10]

Hardy, Norman. “The Confused Deputy (or why capabilities might have been invented).” *ACM SIGOPS Operating Systems Review* 22, no. 4 (October 1988): 36–38. [Chapter 2]

Huang, Ken, et al. “Agent Name Service (ANS): A Universal Directory for Secure AI Agent Discovery and Interoperability.” arXiv:2505.10609, May 2025. [Chapter 5]

Huang, Ken, Vineeth Sai Narajala, John Yeoh, Jason Ross, Ramesh Raskar, Youssef Harkati, Jerry Huang, Idan Habler, and Chris Hughes. “A Novel Zero-Trust Identity Framework for Agentic AI: Decentralized Authentication and Fine-Grained Access Control.” arXiv:2505.19301, May 2025. https://arxiv.org/abs/2505.19301. Research preprint, not peer-reviewed. [Chapters 3, 5, 7, 9]

Kasirzadeh, Atoosa, and Iason Gabriel. “Characterizing AI Agents for Alignment and Governance.” arXiv:2504.21848, 2025. https://arxiv.org/abs/2504.21848. [Chapters 1, 4]

Lampson, Butler, Martín Abadi, Michael Burrows, and Edward Wobber. “Authentication in Distributed Systems: Theory and Practice.” *ACM Transactions on Computer Systems* 10, no. 4 (November 1992): 265–310. [Chapters 2, 7]

Milanta, Marco, and Luca Beurer-Kellner. “GitHub MCP Exploited: Accessing Private Repositories via MCP.” Invariant Labs, May 26, 2025. https://invariantlabs.ai/blog/mcp-github-vulnerability. [Chapters 9, 10]

Pang, Ruoming, et al. “Zanzibar: Google’s Consistent, Global Authorization System.” In *Proceedings of the 2019 USENIX Annual Technical Conference (USENIX ATC ’19)*, July 2019. [Chapter 6]

Reddy, Pavan, and Aditya Sanjay Gujral. "EchoLeak: The First Real-World Zero-Click Prompt Injection Exploit in a Production LLM System." arXiv:2509.10540, 2025. https://arxiv.org/abs/2509.10540. Accessed September 28, 2026. [Chapters 2, 10]

South, Tobin, Samuele Marro, Thomas Hardjono, Robert Mahari, Cedric Deslandes Whitney, Alan Chan, and Alex Pentland. "Position: AI Agents Need Authenticated Delegation." *Proceedings of the 42nd International Conference on Machine Learning*, PMLR 267 (2025): 82211–82231. https://proceedings.mlr.press/v267/south25a.html. [Chapters 2, 7, 11]

Incidents and Vulnerability Disclosures

AI Incident Database. "Incident 1152: LLM-Driven Replit Agent Reportedly Executed Unauthorized Destructive Commands During Code Freeze, Leading to Loss of Production Data." Incident date July 18, 2025. https://incidentdatabase.ai/cite/1152/. Accessed September 28, 2026. [Chapters 2, 6, 10]

Google Threat Intelligence Group and Mandiant. "Widespread Data Theft Targets Salesforce Instances via Salesloft Drift." August 26, 2025 (updated August 28, 2025). https://cloud.google.com/blog/topics/threat-intelligence/data-theft-salesforce-instances-via-salesloft-drift. Accessed September 28, 2026. [Chapters 5, 8, 10, 16]

Microsoft Security Response Center. "CVE-2025-32711: M365 Copilot Information Disclosure Vulnerability." June 11, 2025. https://msrc.microsoft.com/update-guide/vulnerability/CVE-2025-32711. Accessed September 28, 2026. [Chapters 1, 2, 6, 10]

Vendor, Analyst and Other Sources

A2A Project. "A New Chapter for A2A: Joining the Agentic AI Foundation." August 27, 2026. https://a2a-protocol.org/latest/blog/2026/08/27/a-new-chapter-for-a2a-joining-the-agentic-ai-foundation/. Accessed September 28, 2026. [Chapter 9]

Baker McKenzie. "Singapore: IMDA Updates Model AI Governance Framework for Agentic AI." June 2026. https://www.bakermckenzie.com/en/insight/publications/2026/06/singapore-imda-updates-model-ai-governance-framework-for-agentic-ai. [Chapter 11]

ForgeRock. "ForgeRock Unveils a New Offering to Create a Smarter Identity Management and Governance Solution." Press release, December 10, 2019. https://www.globenewswire.com/news-release/2019/12/10/1958748/0/en/ForgeRock-Unveils-a-New-Offering-to-Create-a-Smarter-Identity-Management-and-Governance-Solution.html. [Chapter 3]

Gartner. "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027." Press release, June 25, 2025. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027. [Chapters 1, 3, 15]

Google Cloud. "Agent Identity overview." Google Cloud IAM documentation. https://docs.cloud.google.com/iam/docs/agent-identity-overview. Accessed September 28, 2026. [Chapter 4]

Microsoft. "Administrative relationships in Microsoft Entra Agent ID (owners, sponsors, and managers)." Microsoft Learn. https://learn.microsoft.com/en-us/entra/agent-id/agent-owners-sponsors-managers. Accessed September 28, 2026. [Chapter 8]

Microsoft. "Agent identity sponsor tasks in Lifecycle Workflows." Microsoft Learn. https://learn.microsoft.com/en-us/entra/id-governance/agent-sponsor-tasks. Accessed September 28, 2026. [Chapter 8]

Microsoft. "What's new in Microsoft Entra Agent ID." Microsoft Learn. https://learn.microsoft.com/en-us/entra/agent-id/whats-new-agent-id. Accessed September 28, 2026. [Chapters 1, 4]

Model Context Protocol. "The 2026-07-28 Specification." MCP blog, July 28, 2026. https://blog.modelcontextprotocol.io/posts/2026-07-28/. Accessed September 28, 2026. [Chapters 9, 15]

Replit. "Introducing a Safer Way to Vibe Code with Replit Databases." https://blog.replit.com/introducing-a-safer-way-to-vibe-code-with-replit-databases. Accessed September 28, 2026. [Chapter 10]

Replit. "Introducing Plan Mode: A Safer Way to Vibe Code." https://blog.replit.com/introducing-plan-mode-a-safer-way-to-vibe-code. Accessed September 28, 2026. [Chapter 10]

SailPoint. "SailPoint Research Highlights Rapid AI Agent Adoption, Driving Urgent Need for Evolved Security." Press release, May 28, 2025. https://www.sailpoint.com/press-releases/sailpoint-ai-agent-adoption-report. [Chapters 10, 12]

Tinkoff, Dan, Lieven Van der Veken, Michael Chui, and Tara Balakrishnan. "The State of AI in 2026: On the Road to ROI." McKinsey & Company, August 25, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai. Report PDF, https://www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai/the-state-of-ai-in-2026-on-the-road-to-roi.pdf. Accessed September 28, 2026. [Chapters 1, 15]

© 2026 Rahul Bhardwaj. Published by Arvik Press. You may copy and adapt the tools for use within your own organization. They may not be resold or republished.