Elisa Bertino
Purdue University

Zero Trust Architecture: Does It Help?

Today there is a lot of emphasis on zero trust architecture (ZTA). ZTA has been introduced as a fine-grained defense approach paradigm that shifts defenses from static, network-based perimeters to users, assets, and resources.[1] It assumes that no entities outside and inside the protected system can be trusted and therefore requires articulated and high-coverage deployment of security controls, such as authentication and access control. In a way, ZTA is not new; the idea that securing a system requires pervasive, fine-grained, and continuous deployment of layered security controls is quite obvious. However, the current emphasis on ZTA is important as it pushes systematic approaches to cybersecurity.

However, deploying ZTA is complex from both the technical and organizational points of view as ZTA makes security management much more complex than already is. Because ZTA enforces fine-grained security controls, we can expect that huge numbers of policies have to be specified, implemented, deployed, and managed. These policies will likely be attribute based, that is, based on properties of subjects, protected resources, and contexts. For a system controlled by attribute-based policies, it is critical that these attributes be trustworthy. When one is not sure about some attributes' trustworthiness, then some risk-based criteria must be used for control-related decisions. Such criteria are often difficult to formalize and automate as they are application dependent. Perhaps some artificial intelligence/machine learning techniques could help here, which, however, would then introduce the problem of properly training and securing these models from attacks.

Also, policies must be correct, consistent, minimal, and complete. Correctness requires that the policies comply with their intended goals. Typically, those goals are derived based on high-level organizational policies, risks, potential attacks, and so forth. Eliciting such goals is far from trivial. Consistency is critical for policies that support explicit denials, such as policies for attribute-based access control; it requires that a set of policies does not have both a permission and denial for the same subject, action, and protected resource. When it is not possible to statically enforce consistency on a set of policies—for example, when the set of all possible attributes of subjects and protected resources and the set of possible contexts are not known in advance, mechanisms for conflict resolution need to be included in the access control mechanism. Minimality refers to making sure that the set of policies does not include redundant policies. Redundant policies increase the administrative work for managing policies; for example, if a subject is not any longer allowed to access a given network, all of the policies covering such access must be properly modified and/or removed, thus increasing the security risk if these revocations are not properly executed. Completeness requires that for any action to be executed by the subjects of the system there is a corresponding policy controlling such execution requests. If, for a given access request, there is no corresponding policy, the default decision usually taken by access control systems is to deny the access. Such an approach, however, may lead to situations in which subjects that have a legitimate reason for accessing the protected resource are unable to access the resource. This in turn requires interactions with administrators, which are expensive and introduce delays and, as a result, applications can become brittle.

Additional challenges include the impact of ZTA control on application performance, the difficulty of upgrading legacy infrastructures to ZTA and the integration

continued on p. 95

Digital Object Identifier 10.1109/MSEC.2021.3091195
Date of current version: 3 September 2021

Last Word

continued from p. 96

of ZTA with cloud infrastructures.[2] Addressing these challenges requires not only technical solutions but also incremental deployments of ZTA. So, it would seem that one may end up with “ZTA islands” within systems.

In summary, I believe that ZTA can help with security, but its management and deployment will not be easy. ■

References

1. “Zero trust architecture.” NIST, Gaithersburg, MD, Aug. 2020. [Online]. Available: https://csrc.nist.gov/publications/detail/sp/800-207/final
2. S. Shackelford, “Zero-trust security: Assume that everyone and everything on the internet is out to get you—and maybe already has,” The Conversation. https://theconversation.com/zero-trust-security-assume-that-everyone-and-everything-on-the-internet-is-out-to-get-you-and-maybe-already-has-160969 (accessed May 20, 2021).

Elisa Bertino is the Samuel D. Conte Professor of Computer Science at Purdue University, West Lafayette, Indiana, 47907, USA. Prior to joining Purdue, she was a professor and department head at the Department of Computer Science and Communication of the University of Milan, Italy. Her recent research focuses on cybersecurity and privacy of cellular networks and Internet of Things systems, and edge analytics and machine learning for cybersecurity. Bertino is a Fellow of IEEE, the Association for Computing Machinery, and the American Association for the Advancement of Science. Contact her at bertino@purdue.edu.

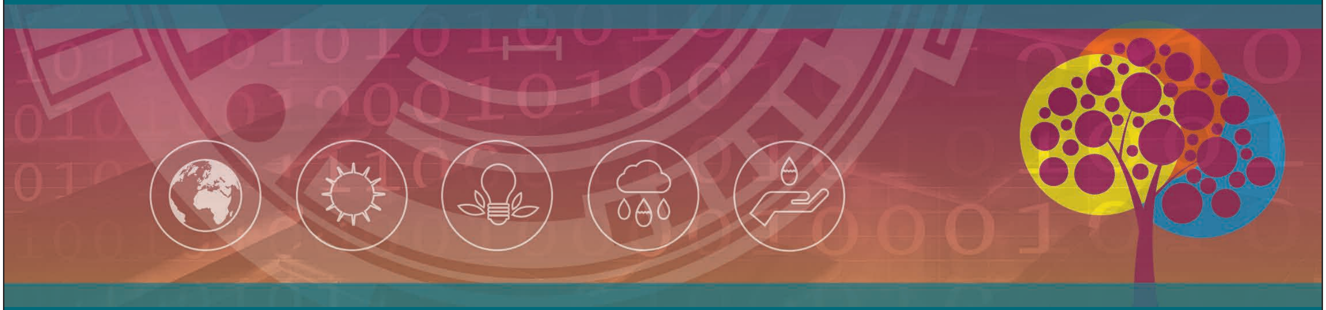

Digital Object Identifier 10.1109/MSEC.2021.3103624