Towards a Model for Zero Trust Data

Jason M. Pittman
Booz Allen Hamilton, USA
pittman_jason@bah.com

Shaho Alaee
Booz Allen Hamilton, USA
alaee_shaho @bah.com

Courtney Crosby
Booz Allen Hamilton, USA
crosby_courtney @bah.com

Tom Honey
Booz Allen Hamilton, USA
honey_tom @bah.com

Geoffrey M. Schaefer
Booz Allen Hamilton, USA
schaeffer_geoffrey@bah.com

***Abstract*— The world has realized traditional cybersecurity models are flawed because users and systems behind the perimeter are implicitly trusted. The response has been to treat access requests and behaviors post-access as untrusted. Thus, the aim of such zero trust architecture is to establish a borderless access-control framework. Accordingly, existing research is centered around network perimeters and communications layers. That is, data access channels or endpoints and not data itself. Consequently, we conducted a systematic review of relevant literature and developed a model illustrating a potential application of zero trust tenets and principles to data objects instead of data access pathways based on the findings. Concurrently, given the rising popularity of employing artificial intelligence to zero trust frameworks, our zero trust data concept targets artificial intelligence training and real-world evaluation data segments.**

Keywords—zero trust, cybersecurity, data, artificial intelligence

I. Introduction

The modern enterprise is bereft of certainty in terms of operational security. Friends are foes and adversaries appear as friends. Further, the differentiation between use and misuse is touted as quantitative but remains qualitative at best. As a result, a shift in the cybersecurity paradigm has begun which seeks to position trust itself as a vulnerability. More specifically, the traditional cyber- security model of implicitly trusting users and systems once they are within the enterprise perimeter is eschewed in favor of not trusting at all [5].

This movement - zero trust architecture- was born from the observation [22] that traditional cybersecurity continually fails because of fundamental misappropriation of trust. Remarkably, 72% of organizations planned to implement zero trust capabilities in 2020 [6]. A year later, the number rose to 76% [18]. The remarkability of these percentages becomes clear when evaluated in the context of zero trust architecture becoming a defined cybersecurity concept barely a decade earlier [22]. We take this as evidence that zero trust architecture has momentum as a cybersecurity paradigm as well as a meaningful rate of adoption to substantiate the hype.

However, a general problem is zero trust architecture applies to data access points (i.e., endpoints) but data objects are not discussed throughout the literature [6]. There is a prima facie difference between access to data and data being accessed. As much as the difference is recognized in the zero trust architecture literature [5, 6, 37], we see no effort to work at the level of data. Instead, the operational focus continues to be on the network perimeter [12, 19, 34, 39] and adjacent access endpoints such as Internet of Things or Cloud [30, 8].

Coupled with this general problem, the same research suggests, "[l]ooking out further, generative adversarial networks will continuously verify the efficacy of zero trust protection by generating synthetic attacks and threats" [6, p. 113]. Additional literature [8, 19, 11, 12, 38] has likewise pointed towards artificial intelligence as having a critical role in the future of zero trust architecture. The same literature has yet to consider the role of data objects as a necessary input to the very artificial intelligence systems purported to play such a role in trust-based cybersecurity constructs. Such a gap in the research is indicative of a specific problem. Thus, the purpose of this work is to demonstrate a model for zero trust at the level of data objects within artificial intelligence training and evaluation operational segments.

II. Related Work

The essence and substance of zero trust data rests upon four existing knowledge domains. Foremost, we operationalize *trust* in two contexts: as trust relates to knowledge as a general case and as trust relates to technology as a specific case. Then, we introduce the concept of *approximate epistemology*. Zero trust architecture relies upon conditional reasoning and subjective logic. Approximate epistemology is a necessary bridge then between *trust* and the technological instantiation of zero trust. The third knowledge domain encapsulates zero trust architecture. We found two relevant categories in the zero trust literature: one category contains the fundamental elements of zero trust and the other details the growing role of AI in zero trust. Lastly, a consequence of involving AI in zero trust necessitates discussing adversarial AI.

A. Trust

Knowledge relies on trust. In fact, we can articulate four discrete components of epistemic trust: belief, communication, reliance, and confidence [25]. From there, we can subgroup the components into epistemic variables (belief and communication) and trust variables (reliance and confidence) [25]. For the purposes of this work, we are most interested in trust variables but recognize epistemic variables cannot be fully decoupled from our concepts. To that end, trust in this context, exemplifies the social aspect of knowledge insofar as we do not directly experience trust but hold trust as valid because of the collective position of validity.

In this way, trust is a non-Boolean proxy for human behavior [36]. Furthermore, technological mediation is the embodiment of that proxy. Meaning, humans trust humans but pass trust judgement vis-vis the technologies created and used by people. Furthermore, perceived trust to be integral to society [35]. That is, trust as a knowledge construct, exists in many disciplines and permeates our cognitive existence [27]. Additionally, there is an argument to be made that, by using

technology, we implicitly place trust in such technology [23]. Nonetheless, trust we do. Certainly, part of such trust is due to the mediation technology provides. As well, trust in technology and trust from technology are integral functions. At the same time, we must be cautious in establishing concepts leading to technological trust, especially when trust is first positioned as a vulnerability. Such caution is warranted insofar as research [16, 17] has suggested that technological trust first-and-foremost stems from our relation to the technology.

B. Approximate Trust

Furthermore, research [20, 21, 2] suggests approximate trust is an extension of technological trust because modern technology operates within environments harboring high degrees of uncertainty. In the face of such un- certainty, trust is an emergent judgement of action based on assumed truth [29, 9]. If we treat data with zero trust, then data have zero truth. Therefore, it follows data must be treated as always false. If we accept such a conclusion as following from the premises, when combined with the notion of trust existing as a continuous or non-Boolean value, we can pose meaningful questions relative to the overarching topic. For instance, an operational question is how do we derive an approximate truthful (i.e., trusted) conclusion from known false premises?

Research in adjacent areas [33, 2, 26] demonstrates implementations of approximate or *fuzzy* logic. Notably, the implementations represented in the literature are con- fined to authentication and networking-based evaluations of trust. This extends naturally into zero trust architecture wherein traditional binomial conditional reasoning is employed. For instance, just as a common perimeter technology such as a firewall relies on binomial conditional reasoning [5, 20], so too does zero trust architecture. The difference being the former assumes trust whereas the latter assumes untrust.

Moreover, insofar as these architectures attempt to *diagnose* trust, probabilistic reasoning is not included as a rational foundation. This compounds potential issues since especially where human operators are involved since, "there is the possibility in the inability to take into account the analyst's levels of confidence in the probability arguments and the inability to handle the situation when the analyst fails to produce probabilities for some of the input arguments" [21, p. 462].

C. Zero Trust Architecture

Abstractly, NIST [31], codifies zero trust architecture implementation as (a) enhanced governance; (b) micro-segmentation; (c) and software-defined network perimeters. Succinctly, the emphasis of zero trust architecture is on borderless access-controls [1]. As an architecture, the goal of is, "fine-grained identify-based access control" [1, p. 9] to prevent lateral movement across the enterprise. Functionally, zero trust architecture functions in alignment with this goal by forcing the explicit verification, authentication, authorization, and continuous monitoring of access to data [6, 31, 1, 5]. Notably, existing re- search does not demonstrate a means of applying zero trust architecture *to data*.

Nonetheless, the literature does maintain a consistent articulation of the technological components essential for any zero trust architecture. Overarchingly, zero trust implementation requires a centralized controller which verifies access requests [40, 5]. Subordinately, the same re-search details how a policy enforcement point acts as a proxy service for these requests and communicates internally with a zero trust engine component. In turn, the engine cross-references access policies in its allocated pol- icy storage and communicates a *trust, no trust* semantic back upstream to the policy enforcement point [40, 5]. In this way, the default assumption of untrusted is equivalent to a default deny all in a firewall.

In addition to the implementation goal and overall architecture, there is consensus across the literature concerning the tenets and principles zero trust architecture embodies. For instance, common tenets [1] are (a) segmentation of access; (b) authentication for all access; (c) end-to-end encryption; (d) least privilege always; and (e) continuous monitoring of all endpoints. Stated differently but with the same intent [5], it can be said that zero trust (a) must apply to all data and services; (b) all access must be secured; (c) trust is never a default state; (d) it must be characteristics, behaviors, and environmental attributes which earn trust and not identity credentials; (e) and access is always temporary.

Despite the rising popularity and rate of adoption, zero trust architecture is not foolproof. Adversaries can by- pass zero trust architecture controls [1] if they are able to sufficiently alter the underlying policy or present them-selves as conforming to the policy as a form of trojan horse appearance. The potential issues are compounded by the necessity to construct zero trust models as self- regulating and immutable. To this end, the literature is pointing towards incorporation of AI into the PEP and engine layers.

Indeed, an AI agent is an appropriate technology to handle the complexity of mediating untrusted access [38]. The addition of AI is not without its own perils though. Chiefly, AI is tightly coupled to data and the standard AIOPs workflow includes two segments vulnerable to manipulation [7, 14]. The level of irony associated with employing AI to zero trust while not securing trust within the AI itself cannot be overstated.

D. Adversarial AI

The relation of adversarial AI to zero trust architecture may seem tenuous at first. However, an outstanding challenge in adversarial AI is to mitigate threats originating from areas of uncertainty [3, 4]. Meanwhile, zero trust architecture research [6, 8, 19, 11, 12, 38] is calling for deeper incorporation of AI. Critically, data are not capable of proving trust as might be the case for a user accessing a network segment. Thus, we understand some portion of the uncertainty to be related to data used during AI model training and, separately but coupled to the same idea, data ingested by the AI during operative evaluation.

In brief, adversarial training encompasses a set of, "intentionally worst-case perturbations to examples from the dataset" [10, p. 1]. A variety of examples exist in the research, most notably methods to perturb or *poison* AI with the intent of corrupting classification modalities [10, 24]. While such perturbative techniques can be effective, the research community is actively developing countermeasures [3]. Unfortunately, the countermeasures appear effective only within spaces governed by certainty.

Based on this, we argue for reconceptualizing adversarial training attacks within the framework of zero trust. In other

words, training data are not *poisoned* as much as such data are *untrusted*. In doing so, the evolving re- search in the adversarial AI space remains adjacent while also allowing for zero trust data to co-evolve as a distinct area of investigation.

III. Method

The goal of this research was to demonstrate a model for zero trust data. To affect this goal, we followed a systematic literature review methodology [15, 28]. Specifically, we (1) searched for literature; (2) selected results from the search based on inclusion criteria; (3) extracted relevant features; (4) and synthesized those features into findings.

Further, to guide and organize the work, we developed three research questions:

Q1: What zero trust architecture tenets or principles are applicable to zero trust data?

Q2: What tenets or principles do not exist in zero trust architecture that are necessary for zero trust data?

Q3: How do the responses to questions 1-3 converge into a zero trust data model?

A. Literature Search and Selection

We operationalized a series of literature searches as input to the systematic review methodology. The searches were conducted against prominent research databases such as ACM, IEEE, EBSCOHost, Springer, dblp, Microsoft Re- search, and Arxiv. We also leveraged Google Scholar to search these and other academic indices. Specific search strings included, but were not limited to, *zero trust*, *zero trust architecture*, *trust and technology*, *zero trust and (tenets or principles)*, *adversarial training approximate trust and technology*.

Overall, the search uncovered more than 874,000 articles. As with all literature searches, a smaller subset was created through a manual review of title and abstract. Then, the operational corpus was selected based on a reading of specific article sections such as introduction and results or findings.

B. Literature Inclusion Criteria

As with all literature searches, not all results are relevant or of sufficient quality to be valuable to the systematic review. Therefore, a definitive protocol to guide include (or exclude for that matter) returned literature becomes operational vital. In view of this, we included zero trust architecture and adversarial AI relevant research from the past ten years. The fields are significantly new and thus date criteria did not negatively impact the literature searches. On the other hand, trust and approximate truth relevant research spans a wide timeline. Therefore, it did not make sense to apply date criteria.

Furthermore, we did not actively include or exclude based on stated research methodology. In all cases we included previous reviews as well because of the leverage provided by such literature. Against the background of existing literature review, we focused on theoretical re- search given the exploratory nature of this work. While we had a primary interest in research demonstrating zero trust architecture models or prototypes of models, we included any relevant research attempting to situate such theory in application (e.g., Internet of Things, Cloud, and so forth).

Likewise, the discriminatory criteria applied to the ad- jacent literature categories foremost operationalized the theoretical utility of the topic. For instance, there is theoretical utility in developing a basis for trust insofar as trust is utilized in zero trust without a need to consider the applicable situation of the former.

IV. Findings

The following sections present our findings based on the outcomes of our systematic review. The findings are grouped and presented according to our guiding research questions for organizational purposes, not to imply state or prioritization in any way.

A. Question 1 - Tenets and Principles

Tenets and principles form the core of any exploratory effort with a goal of developing a model. To that end, we found nine articles [38, 32, 31, 22, 8, 6, 5, 1, 13] which explicitly asserted a set of tenets or principles for zero trust architecture. The set of tenets or principles ranged from three to seven assertions with three being common. Overall, we found the following applicable to our work:

1) trust is not a default state

Whereas the literature applies trust to users or systems accessing services and endpoints, the same applies to data objects.

2) access must be segmented

The literature conceptualizes access in terms of network communications. We refactor this principle to apply in terms of programmatic access between internal software components.

3) activity must be continuously monitored

Unlike with the previous two principles, activity is a general concept with unaltered applicability to our research goal. We do contextualize activity within the specific bounds of AI model training and evaluation (Fig. 1, Appendix A).

The remaining tenets and principles in the literature were related to authentication, least privilege, or to non-data objects. Consequently, those are not viable within the context of our research purpose. However, this is not to say there are no other elements potentially relevant.

B. Question 2 - Missing Elements

Indeed, given the identified problem stated in the intro- duction, we understood from the beginning of the study there might be missing elements from the set of tenets and principles. Accordingly, a concurrent phase of the systematic review was to infer absent features or characteristics in relation to the specific context for the potential zero trust data model.

1) data as an object

Without any doubt, the first missing element from existing zero trust architecture principles are data. Instead of relegating data to a motivation for implementing conventional zero trust architecture, this work positions data as the focal point of the architecture. Moreover, it may be important to distinguish between individual data elements

(e.g., a specific feature) and complete datasets (e.g., a data model).

2) internal components

The existing literature does not specifically outline architectural concepts or implementations relative to internal software components. This is rational when discussing zero trust at the network layer. However, particularly in the context of AI, internal components are critical elements to incorporate as such are simultaneously the end- point and the access mechanism in relation to data.

C. Question 3 - Converged Model

Perhaps we can demonstrate the converged model of tenets and principles with missing elements by leveraging the logic of approximate trust. In this manner, we use a contrapositive quantifier: assume values are false and compute the following:

$$\forall x(\neg Q \rightarrow \neg P)$$

Fortunately, the adversarial AI literature demonstrates the traditional paradigm of: If an object is red, is an octagon, and has the text "STOP", then the object is a stop sign.

We can take the contraposition of the traditional paradigm as: If the object is not a stop sign, then it is not red, not an octagon, and does not have the text STOP.

On one hand, this is a classification problem: trust versus untrusted or true versus false. On the other hand, this is a regression problem: spectrum or degree of truth (i.e., certainty). Thus, our assertion is the concept of how the zero trust data model operates can be described as a function with two inputs as U or the untrusted set of data and T as the trusted set of data. The delta between the intersection of the lower approximation of these sets is deltaed against the intersection of the lower approximation of the same:

$$f(U, T) \rightarrow (\underline{U} \cap \underline{T}) \Delta (\underline{U} \cap \underline{T})$$

The intended outcome is output approximately more trusted or approximately less poisoned. A simple example embodying this expression is the case of a poisoned stop sign artifact. If the zero trust data system assumes all road signs are untrusted (i.e., poisoned), it can use the approximation of (1) an object appearing as a sign and (2) approximating a stop sign given a stop signs unique features contained in the trusted data, the system ought to approximate stop sign from the intersection road sign features and observed features.

Doing so elevates categorization beyond trust and to truth. Put simply, the zero trust data model asserts it is true enough an object appears like a stop sign to be a stop sign. This is in contrast to trusting an object is or is not a stop sign.

With the stop sign example in mind, we contend a zero trust data solution must include two components, one for each segment of the overall AIOPs pipeline (Fig. 2, Appendix A). On the training segment, we see the potential for employing a Bayesian Network with an embedded subjective logic gate. The Bayesian Network would quantify approximate truth through two mechanisms (see Appendix B, Fig. 3). On the evaluation segment, the zero trust data model takes a queue from existing research [4] in leveraging a GAN [4].

V. Conclusion

Zero trust is a state, not a discrete technology, pol- icy zone, or protocol layer. Despite the brief time since its inceptions [22], zero trust has evolved a stable set of core tenets and principles governing its architecture. While specific assertions vary across the literature [38, 32, 31, 22, 8, 6, 5, 1, 13], there is consistency in implementations across a variety of technological plat- forms [12, 30, 36]. However, research [1] suggests zero trust architecture may be imparting a false sense of security because the dominant architecture focuses on end- points.

The purpose of this work was to rebalance the traditional paradigm wherein systems and data within a de- fined boundary are implicitly trusted. The rebalancing is achieved by inverting trusted to untrusted at the level of data rather than at the level of access to such data. In this sense, there is applicability of zero trust data in protecting against adversarial AI manipulations at the data layer.

While conducting the systematic review, we realized zero trust data may be useful outside the scope of cybersecurity. For instance, in general, machine learning depends heavily on unbiased, normalized data. It may be possible to apply a zero trust data solution to guard against such data-based issues. The result would be machine learning models with higher accuracy and lower false positives during evaluation periods. Furthermore, additional future work should include development of an applied prototype leveraging the zero trust data function demonstrated in our findings. In doing so, the inevitable convergence between the related work may be realized. Finally, an important factor not considered thus far is the human-computer interaction consequences of implementing zero trust both in the traditional context as well as the zero trust data framework.

References

[1] Alevizos, L., Ta, V. T., and Hashem Eiza, M. Augmenting zero trust architecture to endpoints using blockchain: A state-of-the-art review. Security and Privacy 5, 1 (2022).

[2] Arora, G., Mathur, I., and Gandhi, S. Quantifying trust evaluation based on approximate rea- soning. In 2015 2nd International Conference on Computing for Sustainable Global Development (IN-DIACom) (2015), IEEE, pp. 1448–1451.

[3] Bai, T., Luo, J., Zhao, J., Wen, B., and Wang, Q. Recent advances in adversarial train- ing for adversarial robustness. arXiv preprint arXiv:2102.01356 (2021).

[4] Bai, T., Zhao, J., Zhu, J., Han, S., Chen, J., Li, B., and Kot, A. Towards efficiently evaluat- ing the robustness of deep neural networks in iot systems: A gan-based method. IEEE Internet of Things Journal (2021).

[5] Buck, C., Olenberger, C., Schweizer, A., Volter, F., and Eymann, T. Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust. Com- puters & Security 110 (2021), 102436.

[6] Campbell, M. Beyond zero trust: trust is a vulnerability. Computer 53, 10 (2020), 110–113.

[7] Dang, Y., Lin, Q., and Huang, P. Aiops: real- world challenges and research innovations. In 2019 IEEE/ACM 41st International Conference on Soft- ware Engineering: Companion Proceedings (ICSE- Companion) (2019), IEEE, pp. 4–5.

[8] Dimitrakos, T., Dilshener, T., Kravtsov, A., La Marra, A., Martinelli, F., Rizos, A., Rosett, A., and Saracino, A. Trust aware continuous authorization for zero trust in consumer in- ternet of things. In 2020 IEEE 19th International Conference on Trust, Security and Privacy in

Computing and Communications (TrustCom) (2020), IEEE, pp. 1801–1812.

[9] Douven, I., and Kelp, C. Truth approximation, social epistemology, and opinion dynamics. Erkenntnis 75, 2 (2011), 271–283.

[10] Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014).

[11] Hale, B., Van Bossuyt, D. L., Papakonstantinou, N., and O'Halloran, B. A zero-trust methodology for security of complex systems with machine learning components. In International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (2021), vol. 85376, American Society of Mechanical Engineers, p. V002T02A067.

[12] Hireche, O., Benza¨ıd, C., and Taleb, T. Deep data plane programming and ai for zero-trust self- driven networking in beyond 5g. Computer Net- works (2021), 108668.

[13] Horne, D., and Nair, S. Introducing zero trust by design: Principles and practice beyond the zero trust hype.

[14] Hornik, K., Stinchcombe, M., and White, H. Multilayer feed forward networks are universal approximators. Neural networks 2, 5 (1989), 359–366.

[15] Hunt, H., Pollock, A., Campbell, P., Est- court, L., and Brunton, G. An introduction to overviews of reviews: planning a relevant research question and objective for an overview. Systematic reviews 7, 1 (2018), 1–9.

[16] Ihde, D. Technology and the lifeworld: From garden to earth.

[17] Ihde, D. Technics and praxis: A philosophy of technology, vol. 24. Springer Science & Business Media, 2012.

[18] Invanti. 2021 zero trustprogress report. Tech. rep., 2021.

[19] Jin, Q., and Wang, L. Zero-trust based distributed collaborative dynamic access control scheme with deep multi-agent reinforcement learn- ing. EAI Endorsed Transactions on Security and Safety 8, 27 (2020).

[20] Josang, A. Conditional reasoning with subjective logic. Journal of Multiple-Valued Logic and Soft Computing 15, 1 (2008), 5–38.

[21] Jøsang, A. Generalising bayes' theorem in subjective logic. In MFI (2016), pp. 462–469.

[22] Kindervag, J., Balaouras, S., et al. No more chewy centers: Introducing the zero trust model of information security. Forrester Research 3 (2010).

[23] Kiran, A. H., and Verbeek, P.-P. Trusting our selves to technology. Knowledge, Technology & Pol- icy 23, 3 (2010), 409–427.

[24] Kireev, K., Andriushchenko, M., and Flammarion, N. On the effectiveness of adversar- ial training against common corruptions. arXiv preprint arXiv:2103.02325 (2021).

[25] McCraw, B. W. The nature of epistemic trust. Social epistemology 29, 4 (2015), 413–430.

[26] Miao, T., Shen, J., Lai, C.-F., Ji, S., and Wang, H. Fuzzy-based trustworthiness evaluation scheme for privilege management in vehicular ad hoc networks. IEEE Transactions on Fuzzy Systems 29, 1 (2020), 137–147.

[27] Origgi, G. Is trust an epistemological notion? Episteme 1, 1 (2004), 61–72.

[28] Pollock, A., Campbell, P., Brunton, G., Hunt, H., and Estcourt, L. Selecting and implementing overview methods: implications from five exemplar overviews. Systematic reviews 6, 1 (2017), 1–18.

[29] Ramsey, J. L. Towards an expanded epistemology for approximations. In PSA: Proceedings of the bien nial meeting of the philosophy of science association (1992), vol. 1992, Philosophy of Science Association, pp. 154–164.

[30] Rodigari, S., O'Shea, D., McCarthy, P., Mc- Carry, M., and McSweeney, S. Performance analysis of zero-trust multi-cloud. In 2021 IEEE 14th International Conference on Cloud Computing (CLOUD) (2021), IEEE, pp. 730–732.

[31] Rose, S., Borchert, O., Mitchell, S., and Connelly, S. Zero trust architecture. Tech. rep., National Institute of Standards and Technology, 2020.

[32] Sanders, G., Morrow, T., Richmond, N., and Woody, C. Integrating zero trust and devsecops. Tech. rep., 2021.

[33] Schmidt, S., Steele, R., Dillon, T. S., and Chang, E. Fuzzy trust evaluation and credibility development in multi-agent systems. Applied Soft Computing 7, 2 (2007), 492–505.

[34] Sengupta, B., and Lakshminarayanan, A. Distritrust: Distributed and low-latency access validation in zero-trust architecture. Journal of Informa- tion Security and Applications 63 (2021), 103023.

[35] Simmel, G. The philosophy of money. Routledge, 2004.

[36] Sood, A. K., Huang, Y., Simon, R., White, E., and Cleary, K. Zero trust intrusion containment for telemedicine. Tech. rep., 2002.

[37] Teerakanok, S., Uehara, T., and Inomata, A. Migrating to zero trust architecture: reviews and challenges. Security and Communication Networks 2021 (2021).

[38] Walker-Roberts, S., and Hammoudeh, M. Artificial intelligence agents as mediators of trustless security systems and distributed computing appli- cations. In Guide to Vulnerability Analysis for Computer Networks and Systems. Springer, 2018, pp. 131–155.

[39] Wang, L., Ma, H., Li, Z., Pei, J., Hu, T., and Zhang, J. A data plane security model based on zero-trust architecture.

[40] Yan, X., and Wang, H. Survey on zero-trust network security. In International Conference on Artificial Intelligence and Security (2020), Springer, pp. 50–60.

Appendix A

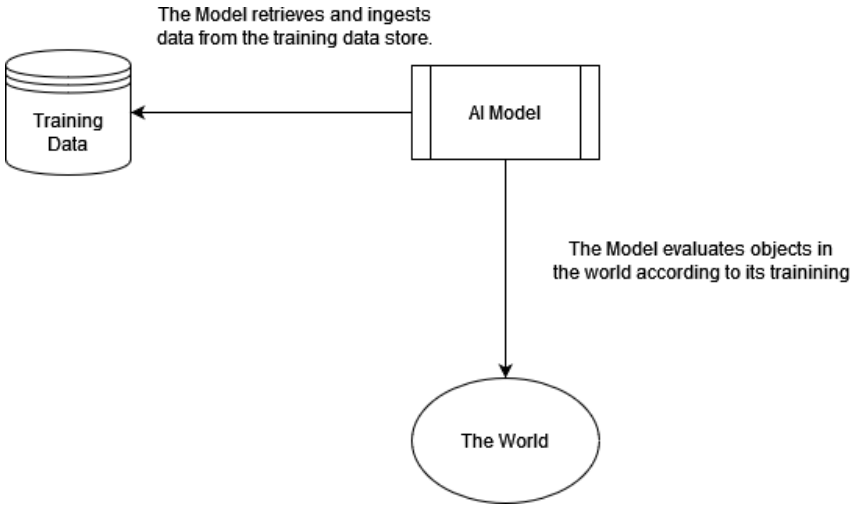


Fig. 1. The Standard AIOPs Pipelines

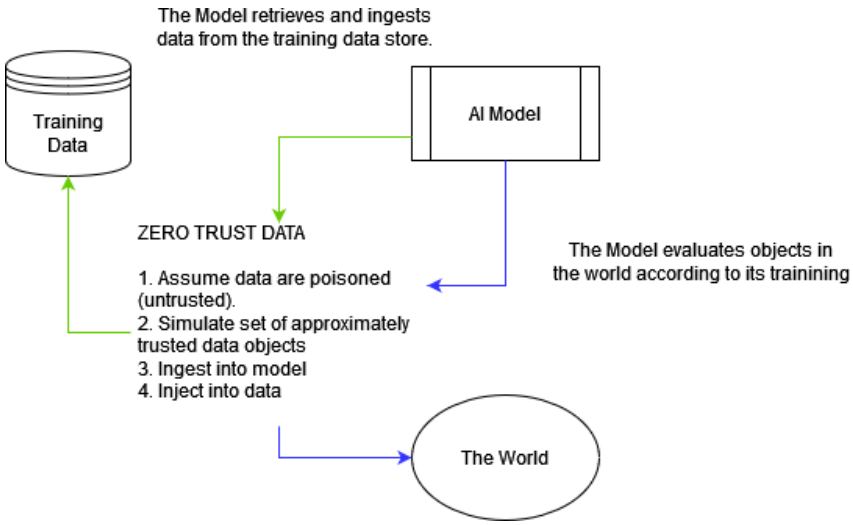


Fig. 2. The Standard AIOPs Pipelines with ZTD

Appendix B

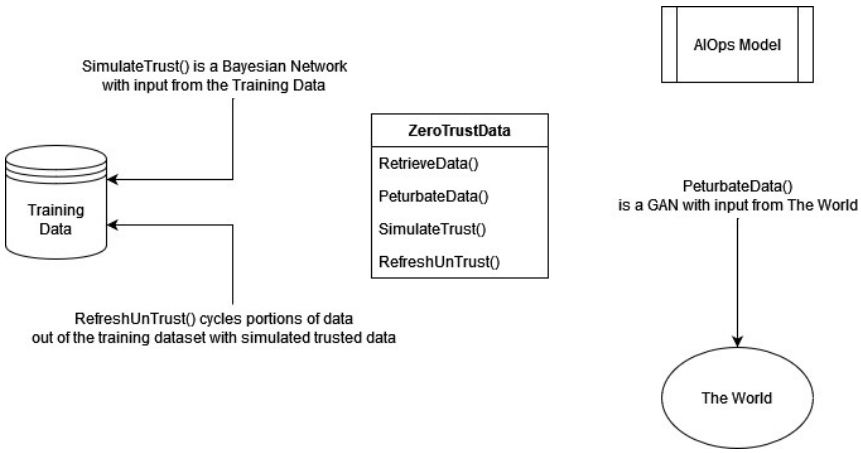


Fig. 3. The Methods for Zero Trust Data