

Zero Trust Architecture for Cyber Resilient Industrial IoTFirdaus Hakimi^{*}, Rivaldi Zola²,¹ Teknik Informatika, Fakultas Teknik, Telkom University, Indonesia² Teknik Informatika, Fakultas Teknik, Telkom University, Indonesia

firdaushakimi90@gmail.com, rivaldi9409@gmail.com

Received	07 01 2026	Revised	26 01 2026
Accepted	30 01 2026	Published	02 02 2026

Abstract Industrial Internet of Things deployments connect sensors, controllers, gateways, engineering workstations, analytics services, and remote support channels to production processes. This connectivity improves visibility and automation but weakens assumptions that a trusted internal network is sufficient protection. Zero trust architecture offers a resource centered model based on explicit authentication, least privilege, continuous evaluation, and constrained access. Applying it to operational technology is not a direct copy of enterprise information technology practice because availability, deterministic communication, safety, long equipment life, and limited device capability impose different constraints. This technical review develops a zero trust architecture for manufacturing small and medium enterprises. The design combines asset identity, network segmentation, policy enforcement, secure gateways, monitored remote access, data integrity, and incident recovery while preserving local control during security service failure. A threat and control matrix links representative attack paths to preventive, detective, and recovery measures. The proposed migration sequence begins with asset and flow discovery, removes uncontrolled remote access, establishes zones and conduits, introduces device and workload identity, and then adds adaptive policy. Rather than restating NIST zero trust principles in general terms, this review's contribution is translating them into a resource-constrained migration sequence and a policy-availability trade-off appropriate to manufacturing SMEs that cannot resource an enterprise-scale security operations function. The article emphasizes that continuous verification does not mean continuous disruption: control traffic should use precomputed policy, local enforcement, and fail safe behavior appropriate to process risk. The main conclusion is that zero trust becomes credible in industrial environments when it is implemented as a risk based engineering program with explicit safety dependencies, measurable policy outcomes, and tested recovery rather than as a single security product.

Keywords *industrial internet of things; zero trust architecture; operational technology security; microsegmentation; device identity; manufacturing SMEs*

1 Introduction

Industrial Internet of Things systems connect physical processes with sensing, analytics, and remote services. Typical components include programmable logic controllers, sensors, actuators, industrial gateways, human machine interfaces, historians, maintenance laptops, cloud applications, and vendor support connections. The combination supports condition monitoring, flexible production, energy management, and data driven improvement. It also creates paths across systems that were previously isolated. Research on industrial IoT consistently identifies heterogeneous devices, long lifecycles, insecure protocols, limited resources, and safety dependencies as major security challenges [1]-[4].

Perimeter security assumes that traffic inside a protected network is more trustworthy than traffic outside it. That assumption becomes weak when users work remotely, vendors connect for maintenance, mobile equipment crosses zones, cloud services exchange data with the plant, and compromised credentials originate from an approved device. Zero trust shifts the decision from

network location to the identity and state of the subject, device, workload, resource, and request. NIST describes zero trust as removing implicit trust based only on location or ownership and making authentication and authorization discrete functions before a resource session [5].

Operational technology cannot adopt enterprise controls without adaptation. A delayed or blocked command may affect production quality or safety. Legacy controllers may not support modern authentication, encryption, or endpoint agents. Systems may operate for decades and can only be patched during rare shutdowns. Incident response must preserve safe physical states. NIST guidance for operational technology therefore emphasizes security measures that account for performance, reliability, and safety requirements [6]. A practical zero trust program needs local enforcement, deterministic behavior, protected fallbacks, and compensating controls for devices that cannot participate directly.

This article presents a technical architecture and staged roadmap for manufacturing small and medium enterprises. It does not claim that zero trust eliminates risk or that all controls are appropriate for every process. The contribution is a set of trust boundaries, policy decisions, evidence sources, enforcement patterns, and recovery requirements that can be tailored through risk assessment. The framework aligns resource centered zero trust principles with industrial zones and conduits, device lifecycle management, remote access, data flows, and operational governance.

Existing guidance on zero trust [5] and on operational technology security [6], [7] is written at the level of general principle and is largely resource-agnostic, and it does not resolve how a manufacturing SME with a small, cross-functional team should sequence adoption without a dedicated security operations capability. This review addresses that gap: it does not introduce a new cryptographic or policy-evaluation mechanism, but instead specifies which decisions can be centralized versus which must be precomputed and enforced locally to preserve process availability, and it sequences these decisions into a staged migration with measurable outcomes that a small team can realistically execute and sustain.

2 Review Methodology

This review synthesizes zero trust and operational-technology security guidance from NIST, IEC, ISO/IEC, and ENISA, together with peer-reviewed literature on industrial IoT security, cyber-physical systems, and industrial automation, retrieved from IEEE Xplore, ScienceDirect, ACM Digital Library, and the publishing standards bodies' own repositories, using combinations of the terms *zero trust*, *industrial IoT*, *operational technology security*, *IEC 62443*, *microsegmentation*, and *device identity*. Authoritative standards and framework documents were included regardless of publication date because they define current regulatory and engineering expectations; peer-reviewed technical literature was prioritized from the last ten years to reflect current industrial IoT threat and device characteristics. Sources were included when they specified a reproducible architecture, control, or documented failure mode relevant to industrial deployment; sources limited to product marketing or lacking methodological detail were excluded. Because the evidence base spans standards, surveys, and applied case reports with different scopes and evaluation criteria, synthesis was narrative rather than statistical. Where general enterprise zero trust guidance conflicted with the availability and determinism requirements of operational technology, the more conservative, availability-preserving position was adopted and expressed as a design requirement in Sections 3 and 4, consistent with the approach used in comparable technical-review articles in this journal.

3 Industrial Threat Model

A useful threat model begins with consequences. Loss of availability can stop production or interfere with control. Loss of integrity can change setpoints, logic, recipes, measurements, or maintenance records. Loss of confidentiality can expose intellectual property, credentials, network design, or customer information. Safety consequences may arise when cyber actions alter physical behavior or hide abnormal conditions, as emphasized in critical infrastructure and cyber physical system research [11], [12]. The organization should identify critical functions and safe states before listing controls, because a control that protects confidentiality but creates an unsafe outage is not acceptable.

Representative entry paths include phishing against office or engineering users, stolen remote support credentials, exposed virtual private network services, compromised vendor software, infected removable media, unauthorized wireless devices, insecure cloud application keys, and vulnerable internet facing gateways. After entry, an attacker may discover assets, reuse credentials, move laterally through flat networks, alter controller logic, suppress alarms, or exfiltrate data. The path matters because zero trust seeks to reduce both initial access and the ability to convert one compromise into broad control.

Industrial IoT creates additional machine identities and data dependencies. A gateway may aggregate hundreds of sensors and authenticate to a broker on their behalf. A cloud analytics service may receive data but should not have authority to write to control devices. A maintenance application may need temporary access to one machine, not the entire control network. Each relationship should be expressed as a resource transaction with an allowed subject, protocol, direction, time, purpose, and verification requirement. Implicit reachability is a poor substitute for explicit policy.

Table 1 Representative industrial attack paths and controls

Attack path	Preventive controls	Detection and recovery
Stolen vendor credential	Strong identity managed device approval named resource and time limit	Session recording behavior review rapid revocation
Compromised engineering workstation	Application allowlisting least privilege cell segmentation and protected credentials	Endpoint alert controller change log and trusted image restoration
Insecure IoT device	Gateway isolation restricted flows unique identity and procurement requirements	New peer detection device quarantine and replacement
Manipulated telemetry	Mutual authentication integrity protection and topic authorization	Cross sensor checks protected logs and replay from trusted records
Malicious software update	Signed artifacts dual approval and controlled maintenance window	Version attestation rollback and golden configuration restore

Threat modeling must include failures and mistakes, not only malicious action. A certificate can expire, a policy service can become unreachable, a firmware update can change protocol behavior, and an engineer can apply the wrong rule. Misconfiguration is especially important in small teams where one person may manage networking, applications, and security. The architecture should therefore minimize unique rules, support tested templates, validate changes before enforcement, and retain a rapid rollback. Security logs must be available without becoming a new production dependency.

4 Zero Trust Design Principles

The first principle is to define protected resources. In a plant, a resource can be a controller programming interface, a historian query, a recipe repository, an engineering workstation, a message broker topic, or a remote maintenance function. Protecting a named resource is more precise than treating an entire subnet as trusted. Each resource receives an owner, criticality, allowed transaction set, and evidence needed for access. Policies can then distinguish read only telemetry from logic modification even when both traverse the same gateway.

The second principle is strong identity for people, devices, and workloads. Human administrators should use unique accounts and phishing resistant multifactor authentication where practical. Managed computers require certificates or hardware protected keys and a record of security posture. Services and containers need workload identities that are separate from user credentials. Constrained field devices may use manufacturer installed identities, locally provisioned certificates, secure gateways, or allowlisted physical ports. Shared accounts should be removed or tightly controlled because they prevent attribution and revocation.

The third principle is least privilege with short, context aware authorization. Access is limited by resource, action, protocol, time, and operating condition. A vendor may receive an approved maintenance window through a monitored jump host, with file transfer and controller programming treated as separate privileges. A quality analytics service may read selected broker topics but cannot publish commands. Policy should be stable enough for operations while still revoking access when

identity, device health, or context changes. Continuous evaluation means reassessing evidence at meaningful events, not adding unnecessary latency to every control packet.

The fourth principle is assume breach and contain movement. Segmentation establishes zones for enterprise systems, manufacturing operations, supervisory control, cells or areas, safety systems, and external services. Conduits permit only required flows. Microsegmentation can further restrict workloads within a zone, but it should not create rule complexity beyond the team capacity. Local firewalls, industrial gateways, broker access control, and application authorization can reinforce one another. The design goal is that compromise of a sensor, workstation, or vendor account does not automatically create a route to every controller.

The fifth principle is evidence driven assurance. Identity provider logs, device inventory, network flow records, controller change records, endpoint alerts, gateway decisions, and physical maintenance events should support a shared incident timeline. Logging must be time synchronized, protected from modification, and retained according to risk. The organization should measure denied unauthorized paths, unused privileges, unmanaged assets, policy exceptions, and recovery performance. These outcomes align control selection with the risk and governance structures in the NIST Cybersecurity Framework, NIST SP 800-53, and ISO/IEC 27001 [10], [16], [17], and are more informative than counting installed security products.

5 Reference Architecture

The proposed architecture contains identity, policy, enforcement, visibility, and recovery services. The identity service manages people, devices, workloads, certificates, and lifecycle state. A policy administration function stores approved access rules. A policy decision function evaluates identity, resource, posture, context, and risk. Policy enforcement points exist at remote access gateways, industrial demilitarized zones, cell firewalls, application gateways, message brokers, and selected endpoints. Visibility services collect flows, events, configuration, and change evidence. Recovery services preserve backups, golden configurations, and tested restoration procedures.

Policy decisions should be as local as the operational requirement demands. A remote engineering session can tolerate an identity and posture decision before connection. A time critical controller exchange should not depend on a distant service for every packet. Its policy can be precomputed, distributed to a local enforcement point, and refreshed under change control. If the central policy service is unavailable, existing approved industrial flows continue for a defined period, while new privileged sessions are denied or require a documented emergency procedure. This approach preserves availability without granting indefinite trust.

Table 2 Architecture services and failure behavior

Service	Function	Required failure behavior
Identity	Lifecycle for human device and workload credentials	Existing critical flows use bounded cached policy
Policy decision	Evaluate subject resource posture and context	New privileged sessions are denied or use emergency procedure
Policy enforcement	Permit block limit and record transactions	Local rules remain active without central connectivity
Visibility	Collect flows events configuration and time context	Production continues and logs buffer locally
Recovery	Restore trusted logic configuration identity and services	Priority order and independent access are pretested

Legacy devices are placed behind compensating enforcement. An industrial gateway can authenticate upstream services, restrict commands, translate secure protocols, validate message ranges, and log transactions. A data diode or unidirectional gateway can be used where information must leave a highly critical zone without a return path. Network access control can bind a device identity to a port or switch context. These controls do not make an insecure device secure, but they reduce exposure and provide a manageable boundary until replacement.

Remote access receives a separate trust path. Users authenticate with individual accounts and strong factors, connect from managed devices, receive approval for a named asset and period, and enter through a recorded jump service. Direct inbound access to controllers is prohibited. File transfer

passes malware and integrity checks, and privileged actions generate change records. Emergency access uses distinct credentials, rapid expiration, and post event review. Vendor identities are disabled when contracts or personnel change, rather than remaining permanently active.

Data architecture also requires policy. MQTT or similar brokers should authenticate clients and authorize specific topics and actions in accordance with the protocol's security mechanisms [15]. Telemetry producers normally cannot subscribe to unrelated data, and analytics consumers cannot publish control commands. Data leaving the plant is filtered to the minimum required fields and protected in transit. Models, recipes, and configuration files receive integrity checks and version control. These measures address the risk that valid network connectivity is used to manipulate information rather than exploit a protocol vulnerability.

6 Controls and Technology Selection

Asset discovery is the first technical control because identity and policy cannot protect unknown devices. Passive network monitoring is preferred where active scanning could disturb sensitive equipment. The inventory should include owner, function, location, firmware, protocols, communication peers, criticality, backup status, and support lifecycle. Observed flows are compared with engineering documentation. Differences become investigation items, not automatically authorized rules. The inventory links logical identity to physical asset records and should be updated through procurement, commissioning, maintenance, and retirement.

Segmentation should follow process and consequence. The IEC 62443 family describes zones and conduits for industrial automation and control systems [7]. A practical design can separate enterprise access, operations management, supervisory systems, production cells, and safety related functions. Access control lists permit only documented protocols and peers, while an industrial demilitarized zone mediates enterprise to plant exchange. Internal name resolution, time services, patch repositories, and monitoring require explicit design because hidden dependencies often appear during cutover.

Endpoint security varies by capability. Workstations and servers can use secure configuration, application allowlisting, endpoint detection, protected credentials, and controlled patching. Controllers may rely on signed firmware, disabled unused services, physical protection, secure engineering procedures, and network enforcement. IoT device guidance emphasizes cybersecurity capabilities and manufacturer lifecycle activities [8], [9], while empirical surveys show why exposed device services and weak lifecycle controls remain material risks [18]. Procurement should require unique identity, secure update, vulnerability disclosure, supported cryptography, logging, and a defined support period. Devices that cannot meet minimum requirements need compensating controls and a replacement plan.

Detection should combine network and host evidence with process context. Industrial network monitoring can identify new peers, protocol changes, unexpected write commands, unusual scanning, and communication outside normal periods. Endpoint tools can detect credential theft or malicious files. Process monitoring can identify implausible sensor relationships or commands that conflict with operating state. No single signal proves an attack. Correlation and well defined escalation prevent a noisy anomaly system from overwhelming a small operations team.

Backup and recovery are security controls, not only continuity tools. Offline or protected copies should include controller logic, human machine interface projects, historian configuration, gateway policies, certificates, server images, and dependency documentation. Restoration must be tested on compatible equipment or a representative environment. The team should know which systems recover first, how identity services are restored, and how trusted configurations are verified. A backup that cannot be located, decrypted, or loaded during an incident does not reduce operational risk.

7 Migration for Manufacturing SMEs

A small enterprise should start with a narrow protected surface rather than attempt simultaneous transformation. The first phase identifies critical production services, assets, data flows, remote connections, and current accounts. Immediate risks such as shared remote credentials, internet

exposed devices, unsupported gateways, and absent backups receive priority. Industrial automation implementation experience and smart manufacturing guidance support this staged, context aware approach [13], [14]. The organization also defines a cross functional owner because production, maintenance, information technology, engineering, and management share responsibility for the resulting controls.

The second phase creates basic zones and a controlled remote access path. Existing flows are monitored before blocking begins. Proposed rules run in alert only mode, and engineers review unexpected dependencies. Cutover occurs during a planned window with rollback commands prepared. Exceptions include an owner, purpose, scope, and expiry. This discipline avoids turning temporary troubleshooting access into permanent trust. The goal is a small set of understandable conduits that match actual process communication.

The third phase introduces identity and least privilege for administrators, services, and capable devices. Shared engineering accounts are replaced, multifactor authentication protects remote and privileged access, and certificates identify managed gateways or workstations. Broker topics and application roles are restricted. Legacy devices remain behind gateways with documented compensating controls. The team tests certificate renewal and revocation because an identity system that works only at enrollment will eventually create an outage.

The fourth phase adds posture, analytics, and adaptive decisions where evidence supports them. Device health can influence remote access, and anomalous behavior can trigger session review or temporary isolation. Automated blocking is first limited to high confidence events with low operational consequence. More aggressive response requires process specific tests. Throughout the migration, the enterprise tracks asset coverage, unmanaged communication, privileged accounts, exception age, restoration time, and the percentage of critical paths with tested fallback.

Table 3 Staged migration outcomes

Phase	Primary outcome	Evidence
Discover	Critical assets flows accounts and consequences are known	Validated inventory and data flow map
Contain	Remote access and zones reduce broad reachability	Denied path tests and reviewed exceptions
Identify	People workloads and capable devices use unique credentials	Enrollment renewal revocation and audit tests
Adapt	Context and behavior influence bounded decisions	False action rate operational impact and rollback tests
Sustain	Changes recovery and metrics remain governed	Periodic exercises policy review and restoration evidence

Investment decisions should consider operational simplicity. A platform that requires constant specialist tuning may be unsuitable even if its feature list is extensive. Open standards, exportable configuration, local support, and clear licensing matter. Managed services can supplement scarce expertise, but providers should receive limited access and provide audit evidence. The organization remains accountable for process safety and should be able to operate critical functions when the provider or external connection is unavailable.

8 Verification and Research Agenda

Verification begins with policy tests. For each protected resource, the team defines allowed and denied subjects, devices, actions, and contexts. Automated tests confirm that permitted production flows remain available and prohibited paths fail closed. Timing tests measure added latency and jitter. Failure injection covers loss of identity, policy, logging, and time services. The expected response differs by transaction: established control may continue, a new vendor session may be denied, and an operator may receive an emergency procedure.

Purple team exercises can test realistic attack paths without endangering production. A representative laboratory or digital environment should include the same identity, gateways, policies, and representative devices. Scenarios can begin with a compromised office account, maintenance laptop, IoT sensor, or vendor credential and then test lateral movement. Evidence should show which control interrupted the path, what alerts appeared, and how long containment and recovery required. Successful prevention is useful, but detection and restoration still need testing.

Research should evaluate zero trust in relation to process outcomes. Metrics include unauthorized path reduction, policy decision latency, control traffic impact, time to revoke access, false isolation, exception burden, incident containment, and restoration time. Studies should report architecture, device constraints, traffic types, and safety assumptions. Claims based only on enterprise web applications should not be generalized to deterministic industrial control. Likewise, a laboratory demonstration should not be described as plant validation without field evidence.

An important research gap is policy management at scale for heterogeneous constrained devices. Future work can examine automatically generated candidate policies from observed flows, provided engineers approve them and hidden dependencies are tested. Hardware roots of trust, lightweight attestation, secure gateways, and software defined segmentation also require comparative evaluation. The strongest contribution will be evidence that these techniques reduce attack paths without producing unacceptable outage risk or unmanageable operational complexity.

9 Limitations

This review is narrative and design-oriented rather than a systematic review with a registered protocol; it does not claim exhaustive retrieval of the zero trust and operational-technology security literature or a formal meta-analysis, and it draws heavily on standards and framework documents whose recommendations are necessarily general rather than site specific. The proposed architecture and migration sequence are conceptual and have not been validated on an instrumented manufacturing SME network; the empirical findings attributed to individual cited studies describe those studies' own reported results and should not be read as evidence already obtained for the framework proposed here. Because manufacturing environments differ substantially in legacy equipment, staffing, and process risk, the staged migration is a general sequence that must be adapted, and the review does not quantify the cost or duration of adoption for a specific facility type.

10 Conclusion

Zero trust can improve industrial IoT resilience by replacing broad network trust with explicit identity, resource policy, least privilege, containment, and evidence. Its industrial implementation must preserve safety, availability, and deterministic control. The proposed architecture distributes enforcement across remote access, industrial boundaries, applications, brokers, gateways, and capable endpoints, while using local precomputed policy for time critical flows. Legacy equipment is protected through compensating boundaries and a replacement plan. A staged migration from discovery and remote access control to segmentation, identity, and adaptive assurance is realistic for manufacturing SMEs. The program should be judged through tested attack paths, policy outcomes, service impact, exception burden, and recovery. Zero trust is therefore best understood as an engineering and governance model, not a product deployment.

AUTHOR CONTRIBUTIONS

CRediT: conceptualisation, literature synthesis, framework development, writing, review, and approval of the manuscript.

FUNDING

The authors declare that this research received no funding from any source.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

DATA AVAILABILITY STATEMENT

No original datasets were generated or analyzed in this study, as this article presents a conceptual technical review and design framework.

References

- [1] L. D. Xu, W. He, and S. Li, Internet of Things in industries a survey, *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, pp. 2233-2243, 2014. <https://doi.org/10.1109/TII.2014.2300753>
- [2] A. Sadeghi, C. Wachsmann, and M. Waidner, Security and privacy challenges in industrial Internet of Things, in *Proceedings of the 52nd ACM IEEE Design Automation Conference*, pp. 1-6, 2015. <https://doi.org/10.1145/2744769.2747942>
- [3] H. Boyes, B. Hallaq, J. Cunningham, and T. Watson, The industrial internet of things an analysis framework, *Computers in Industry*, vol. 101, pp. 1-12, 2018. <https://doi.org/10.1016/j.compind.2018.04.015>
- [4] E. Sisinni, A. Saifullah, S. Han, U. Jennehag, and M. Gidlund, Industrial Internet of Things challenges opportunities and directions, *IEEE Transactions on Industrial Informatics*, vol. 14, no. 11, pp. 4724-4734, 2018. <https://doi.org/10.1109/TII.2018.2852491>
- [5] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, Zero Trust Architecture, NIST Special Publication 800-207, 2020. <https://doi.org/10.6028/NIST.SP.800-207>
- [6] K. Stouffer et al., Guide to Operational Technology Security, NIST Special Publication 800-82 Revision 3, 2023. <https://doi.org/10.6028/NIST.SP.800-82r3>
- [7] IEC, Security for industrial automation and control systems Part 3-3 System security requirements and security levels, IEC 62443-3-3, 2013.
- [8] M. Fagan, K. Megas, K. Scarfone, and M. Smith, Foundational Cybersecurity Activities for IoT Device Manufacturers, NISTIR 8259, 2020. <https://doi.org/10.6028/NIST.IR.8259>
- [9] M. Fagan et al., IoT Device Cybersecurity Guidance for the Federal Government Establishing IoT Device Cybersecurity Requirements, NIST Special Publication 800-213, 2021. <https://doi.org/10.6028/NIST.SP.800-213>
- [10] NIST, The Cybersecurity Framework Version 2.0, NIST Cybersecurity White Paper 29, 2024. <https://doi.org/10.6028/NIST.CSWP.29>
- [11] C. Alcaraz and S. Zeadally, Critical infrastructure protection requirements and challenges for the 21st century, *International Journal of Critical Infrastructure Protection*, vol. 8, pp. 53-66, 2015. <https://doi.org/10.1016/j.ijcip.2014.12.002>
- [12] A. Humayed, J. Lin, F. Li, and B. Luo, Cyber physical systems security a survey, *IEEE Internet of Things Journal*, vol. 4, no. 6, pp. 1802-1831, 2017. <https://doi.org/10.1109/JIOT.2017.2703172>
- [13] P. Leitão, A. W. Colombo, and S. Karnouskos, Industrial automation based on cyber physical systems technologies prototype implementations and challenges, *Computers in Industry*, vol. 81, pp. 11-25, 2016. <https://doi.org/10.1016/j.compind.2015.08.004>
- [14] ENISA, Good Practices for Security of Internet of Things in the Context of Smart Manufacturing, European Union Agency for Network and Information Security, 2018. <https://doi.org/10.2824/851384>
- [15] OASIS, MQTT Version 5.0, OASIS Standard, 2019.
- [16] Joint Task Force, Security and Privacy Controls for Information Systems and Organizations, NIST Special Publication 800-53 Revision 5, 2020. <https://doi.org/10.6028/NIST.SP.800-53r5>
- [17] ISO and IEC, Information security cybersecurity and privacy protection Information security management systems Requirements, ISO IEC 27001, 2022.
- [18] N. Neshenko, E. Bou-Harb, J. Crichigno, G. Kaddoum, and N. Ghani, Demystifying IoT security an exhaustive survey on IoT vulnerabilities and a first empirical look on internet scale IoT exploitations, *IEEE Communications Surveys and Tutorials*, vol. 21, no. 3, pp. 2702-2733, 2019. <https://doi.org/10.1109/COMST.2019.2910750>