

A Novel Malicious Intent Detection Approach in Intent-Based Enterprise Networks

Federica de Trizio, Giancarlo Sciddurlo, Dominga Rutigliano, Giuseppe Piro, Gennaro Boggia
Dept. of Electrical and Information Engineering - Politecnico di Bari, Bari, Italy

Email: {name.surname}@poliba.it

CNIT, Consorzio Nazionale Interuniversitario per le Telecomunicazioni

Abstract—In an era dominated by digitalization in corporate strategies, the effectiveness of network infrastructures constitutes the foundation upon which enterprises build their businesses. Within this context, Intent-Based Networking (IBN) emerges as a networking paradigm that addresses the need for autonomous and intelligent systems, shifting network management from static to dynamic and automated processes. Despite its significant advancements in network automation, IBN introduces new complexities and challenges, particularly concerning the assessment of reliability or trustworthiness. Formulating intent without thorough investigation of users' intentions could potentially lead to deliberate network sabotage, denial of services, data breaches, and privilege escalations, thereby posing significant risks to enterprise networks. To bridge this gap, this study proposes a novel Malicious Intent Detection (MID) module within the IBN framework to construct a security knowledge base for accurately classifying enterprise users' intentions. Specifically, it detects malicious expressions directly during the intent processing stage, thereby preventing the formulation of disruptive network configurations or policies derived from malicious intents. The obtained results demonstrate the effectiveness of the proposed solution, with the ability to detect 94% of malicious intentions and achieve an accuracy of up to 93%.

Keywords—Intent-Based Networking, Malicious Intent Detection, KVI.

I. INTRODUCTION

Enterprise networking has undergone significant expansion in recent years. Over the past two decades, there has been a notable rise in digital-native enterprises, exemplified by companies like Salesforce, Google, Amazon, Uber, eBay, and Airbnb, which have utilized digitization as a disruptive force across various industries [1]. In this context, the evolution of application and network endpoints has rendered traditional manual configurations inadequate. Additionally, significant challenges arise in maintaining consistent device configurations across the network and ensuring alignment with intended service requirements. The automation of operations and services becomes imperative, not only to reduce operational expenses but also to facilitate swift network reconfiguration, ensuring the proper and uninterrupted provision of functionality as intended. Autonomic networks rely on input from operators to provide operational guidance and information regarding the goals, functions, and services the network is intended to support. This input and operational guidance, commonly referred to as intent, is facilitated through the networking paradigm known as Intent-Based Networking (IBN) [2].

IBN plays a crucial role in network automation, enabling proactive and dynamic adaptation to changing conditions. Rather than dealing with low-level configurations, users can express their desired outcomes in terms of high-level objectives in a language understandable to the network, which can be easily converted into specific configurations and policies. Many vendors provide IBN solutions for enterprise networks, including Cisco, VMware, and Juniper Networks, enabling faster and more automated service delivery, network management, and troubleshooting [1]. While IBN functions as a significant means to automate networks, addressing security aspects related to intent formulation remains challenging. The emphasis on security is becoming increasingly prominent for next-generation networks, which must align with the values outlined in the United Nations' Sustainable Development Goals (SDGs), such as reliability and trustworthiness, mapped through proper Key Value Indicators (KVI).

Recent research efforts [3]–[18] focus on developing solutions for processing user intents and extracting relevant information. However, these works do not prioritize the detection of harmful statements that could compromise network functions. As of the time of writing, and to the best of the authors' knowledge, a methodology aimed at detecting malicious intentions in the intent formulation process is still missing. Due to the inherent flexibility in expressing service or resource requests in natural language, users' statements may intentionally contain ambiguities, leading to potential network sabotage, denial of services, data breaches, and privilege escalations. Therefore, assuming intent formulation without further investigation into users' intentions could pose severe risks to the reliability of enterprise networks, as stated in [19].

To bridge this gap, this study introduces a novel approach to detect malicious behavior directly during the intent processing phase. This approach aims to prevent the creation of disruptive network configurations or policies resulting from malicious intents. A Malicious Intent Detection (MID) module, integrated within the IBN framework, is developed to establish a security knowledge base for accurately classifying entities and users' intentions. To achieve this objective, the proposed strategy utilizes the open-source Lumi framework, seamlessly integrated with the Rasa chatbot. This integration facilitates the acquisition, processing, translation, and deployment of intents, ensuring a high level of reliability within the enterprise network. Addressing and simulating a real enterprise scenario,

obtained results testify to the ability of the proposed approach to process and analyze user utterances, detecting malicious intentions with an accuracy reaching up to 93%.

Section II provides a review of the state-of-the-art approaches addressing IBN strategies. Sections III and IV present the enterprise network environment under consideration and delve into the details of the conceived solution. Section V assesses the performance of the proposed approach. Finally, Section VI offers concluding remarks.

II. RELATED WORKS

Supporting the implementation of IBN while promoting the automation of enterprise networks, the current scientific literature focuses on developing strategies to extract information from user expressions and investigate security solutions for intent-based systems.

Initial endeavors [3]–[11] concentrated on processing intents in natural language to extract information about entities, such as network elements, and understanding their context. This enables the detection of how to deploy resources or services in the physical infrastructure. For instance, [4] proposes a system that translates intents into computer programs using Large Language Models (LLMs). This process, facilitated by conversion into execution graphs, generates specific policies for network domains.

The work proposed in [6] introduces Lumi, a Google Dialogflow chatbot-based framework that leverages machine learning and operator feedback to validate and refine translated intents, ensuring alignment with the operator’s objectives. The authors of [7] develop a user-friendly system for enterprise networks called Rasa, using the ETS-Chatbot framework, designed to be accessible even to users with limited networking expertise. User authentication is performed as a procedural step before capturing and elaborating on the intent.

Currently, security-oriented research explores the feasibility of expressing security policies through intents [12]–[18]. For example, [14] defines an approach to ensure the provision of security slices in intent-based systems. By leveraging slice templates, four security types can be requested through intents with a GUI, which can then be translated and used for generating security policies. The authors of [16] and [13] introduce Virtual Network Function (VNF) placement solutions to deploy security intents by identifying providers that satisfy both security and performance requirements.

In [15], a data protection intent framework is proposed. Data services submit data protection intents to the framework, which are deployed by Software-Defined Networking (SDN) controllers in the data plane’s devices through flow rules. In [17], a blockchain-based system is introduced to create security policies from intents recorded as transactions, utilizing the inherent transparency and immutability of blockchains to prevent intruders from formulating harmful demands.

Despite these efforts, current research primarily targets the translation of trusted and authorized security demands into security policies, assuming the data handled and processed is reliable and untampered. The current state of the art lacks

comprehensive investigations into intentional malicious intent detection, particularly within an enterprise context. Existing intent profiling and translation approaches fall short in assessing the potential harm posed to the network by user requests, as they overlook mechanisms for controlling intentions to mitigate the risk of malicious intent injection. Addressing these challenges, this work emphasizes the detection of malicious intents during the intent processing stage to prevent the exploitation of vulnerabilities or tampering with policies, thereby safeguarding the operations of enterprise networks.

III. THE REFERENCE IBN ENTERPRISE ENVIRONMENT

Fig. 1 illustrates the enterprise network environment under consideration, where users express their requests to access enterprise resources and services. Specifically, they articulate their desired outcomes through natural language expressions, employing various methods such as Graphical User Interfaces (GUIs), chatbots, templates, APIs, or voice assistants. To facilitate network automation, the user request is handled by the IBN structure, which comprises four different modules.

Upon submission of a request, the first module, the Information Extraction module, is responsible for extracting entities and contextual information using a Natural Language Processing (NLP) framework to parse and process the user’s requests. Secondly, the Intent Assembly module translates the extracted natural language information into an IBN language. This language acts as a programming language that bridges the gap between human and machine readability [2]. The proposed work considers the Network Intent Language (Nile) as the IBN language. Nile converts the data parsed in the Information Extraction module by formulating the intent with a name and a set of tags defining the intent scope. Once an intent is generated, the network solicits validation from the user through the Intent Confirmation module, either confirming or refining the request. Finally, in the Intent Deployment module, the intent is compiled into network configuration commands or policies to be applied in the physical infrastructure. To this end, Network Service Descriptors (NSDs) can be employed to map the information contained in the Nile intent to a JSON or TOSCA template, which defines all requirements to fulfill the user request. Policy abstraction languages like Merlin and OpenConfig can also be employed. These languages map the Nile features to those provided by the selected language and, through program partitioning, generate code to implement the policy, subsequently compiled into corresponding YANG data models or OpenFlow rules.

IV. THE PROPOSED SOLUTION

The strategy proposed herein aims to extend the Information Extraction module of the IBN framework to support the translation of intents into network policies. This extension is achieved by incorporating the novel MID module, which investigates the nature of requests. This work utilizes Lumi¹ as the IBN framework to support the acquisition and processing

¹<https://lumichatbot.github.io/>




Fig. 1. The considered enterprise network environment.

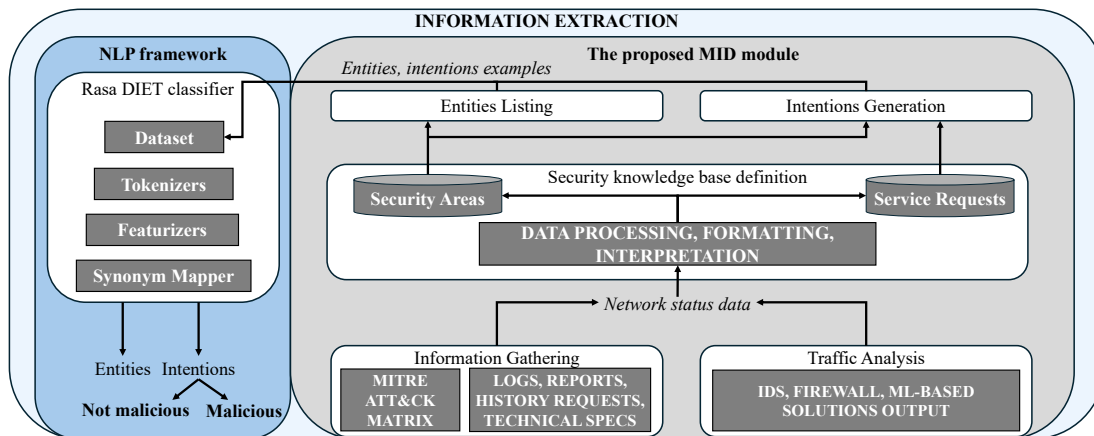


Fig. 2. The Malicious Intent Detection module extension.

of intents, integrating a Rasa-based chatbot to interface with users.

A. Rasa-based Information Extraction chatbot

Rasa stands as an open-source framework tailored for the creation of adaptable and scalable chatbots, employing advanced NLP and machine learning techniques [20]. Within its architecture, it seamlessly integrates the capabilities of the Information Extraction module within the IBN structure to interpret and categorize user expressions. This sophisticated design empowers users to artfully articulate their requests through interactions with text-based chatbots. Furthermore, Rasa offers an extensive suite of models and components meticulously crafted to extract pertinent information from user inputs within the chatbot interface. Leveraging these models facilitates the precise definition of entities, synonyms, and intention classes, thereby enabling the extraction and classification of relevant information in natural language from requests. These intention classes delineate potential meanings and objectives behind user expressions, encompassing diverse scenarios such as conversation initiation or termination, request confirmation or rejection, and more. Notably, Rasa adeptly identifies expressions that diverge from the currently defined intention classes, categorizing them under the designated *out-of-scope* classification. The proposed framework in this paper also incorporates a state-of-the-art multi-task model, known as the Dual Intent Entity Transformer (DIET) classifier, which is proficient in both entity recognition and intent classification. Utilizing advanced tokenization techniques, the

DIET classifier enhances the extraction of features essential for precise entity and intent classification. This process relies on the Synonym Mapper component, which facilitates direct associations between diverse expressions regarding operations, locations, services, and constraints, and their corresponding predefined entities or intents. To accurately recognize entities and classify user expressions, the model requires meticulous training with a comprehensive dataset comprising examples of user expressions systematically categorized by intent classes. Examples of user expressions included in our model are: "Set up an anti-phishing filter on the email server," "Assign network administrator rights to the Operations and Maintenance (O&M) group," and "Enable two-factor authentication for all network management accounts."

The adaptability inherent in Rasa, particularly in the customization of its modules, presents a compelling opportunity to expand this pivotal IBN module. Specifically, this potential could be harnessed for deeper exploration into the detection and classification of malicious intentions articulated by users, a critical endeavor given the significant risks they pose to enterprise networks.

B. Malicious Intent Detection module

To detect malicious behaviour in the intent processing stage, this work customizes the Information Extraction module by considering MID module on the side of the NLP framework. Fig. 2 shows the configuration of the module designed from scratch, comprising five functionalities encompassing: (i) Information Gathering, (ii) Traffic Analysis, (iii) Security knowl-

edge base definition, (iv) Entities Listing, and (v) Intentions Generation.

The Information Gathering and Traffic Analysis functions rely on the physical infrastructure of the enterprise network to collect data related to its status. Specifically, the former entails the collection of internal data, encompassing logs, reports, service request histories, and technical specifications sourced from network components, including servers, devices, and virtual machines. Esteemed tools are employed to facilitate the identification of network vulnerabilities by scrutinizing the methodologies employed by intruders in executing attacks. An example of such a tool is the MITRE ATT&CK Enterprise Matrix [21], which serves as a valuable asset in this endeavor, providing enterprises with a structured framework to comprehend, analyze, and address security threats. Each cell within the matrix corresponds to a specific technique associated with an attacker tactic. The matrix employs color-coded indicators to denote both the frequency and severity of each technique. Conversely, the Traffic Analysis function involves the collection of information derived from Intrusion Detection Systems, Intrusion Detection System (IDS) and firewalls, as well as outputs from machine learning-based solutions. These solutions include Naïve Bayes classifiers, decision trees, and Support Vector Machine (SVM), known for their high interpretability and rapid training, and are designed to identify anomalies within network traffic traces and flows.

The Security Knowledge Base definition function stores structured information, which is leveraged to aid the classifier in making informed decisions. It combines data from diverse network resources into a unified format. After data refinement processes such as filtering, normalization, and aggregation, the information is encoded using descriptive formats and standards like XML, JSON, and STIX. These meticulously represent security-related and threat information, facilitating efficient data management and analysis. Besides, data clustering techniques (e.g., k-means, HDBSCAN) are strategically employed to interpret and extract insights from this enriched data pool. This facilitates the assessment and quantification of risks and vulnerabilities, elucidating the mechanisms through which adversaries illicitly gain access to systems. These techniques identify patterns in the collected data and detect new observations in relation to the existing data. Structured data is organized and stored within two distinct databases. One database is designated for delineating security areas alongside their associated threats and targets, while the other database is dedicated to archiving service requests previously initiated by users. This dual-database structure facilitates efficient data organization and retrieval, enabling comprehensive analysis and response to security incidents.

Utilizing the information stored within the Security Areas database and harnessing Named Entity Recognition (NER) algorithms, the Entities Listing function generates entities and synonyms to refine the chatbot configuration files. Given the requisite diversity of user expressions essential for the comprehensive training of the DIET classifier, it is imperative to tailor the chatbot to define a distinct class of legitimate

service requests, denoted as '*build-intention*', while extending the '*out-of-scope-intention*' class to encompass all potentially malicious user expressions.

Through the application of Generative Adversarial Network (GAN) algorithms to the datasets stored within both databases, the Intentions Generation function adeptly generates instances of conceivable malicious requests. This is essential for the proficient training of the DIET classifier within the NLP framework. The data processed in the newly conceived MID module aids in the proper training of the DIET classifier, facilitating thorough investigation of users' intentions. In accordance with the classification of intentions, the chatbot reacts by delivering the predefined response associated with the '*out-of-scope-intention*' category if a request is classified as malicious. This approach effectively blocks and discards malicious attempts, thereby ensuring that only contextual information retrieved from legitimate user expressions is transmitted to the subsequent Intent Assembly module for further processing and intent translation.

V. PERFORMANCE EVALUATION

The performance of the proposed approach in detecting harmful requests within an enterprise network is investigated herein.

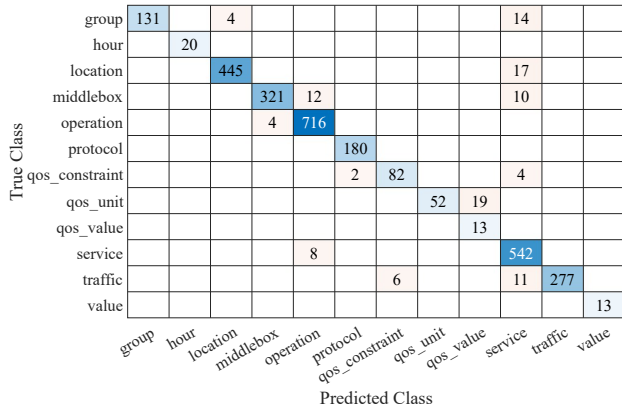
A. Dataset definition for DIET classifier training

A comprehensive definition of the MID module's Security Knowledge Base has been synthesized to elucidate entities and malicious intentions. This encompasses security areas relevant to the physical infrastructure of the enterprise network, including antispam, antivirus, authentication, authorization, access control, malware and data loss prevention, encryption, intrusion detection, network segmentation, patching, hardening, and backup. Each area is correlated with relevant threats and targets, representing entities within the security landscape.

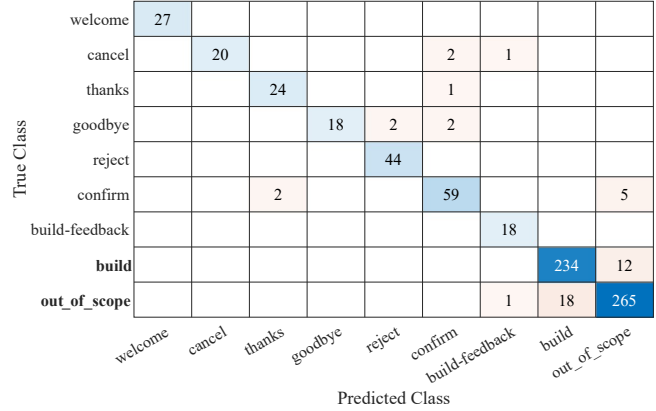
Drawing upon this foundation, 12 types of entities have been delineated, encompassing facets such as traffic, locations, middleboxes, operations, groups, services, Quality of Service (QoS) constraints, QoS values, QoS units, hours, protocols, and values. Examples of these entities, along with potential synonyms, have been cataloged in the `domain.yml` and `lookups.yml` configuration files of the Rasa framework, providing operational guidelines and synonym definitions. Additionally, the `domain.yml` file specifies the entities targeted for detection and classification by the DIET classifier.

Subsequently, potential malicious intentions have been inferred based on delineated threats. These malicious expressions were formulated to serve as training data for the DIET classifier and are stored in the `intents.yml` configuration file of the Rasa framework. In total, 196 malicious requests were defined, complemented by examples from other intention classes, such as initiating, modifying, or concluding a conversation.

Apart from defining the dataset used for training the DIET classifier within the Information Extraction module, this work encompasses 755 user requests sourced from interactions with the enterprise system chatbot. The subsequent section



(a) Entities confusion matrix.



(b) Intentions confusion matrix.

Fig. 3. Confusion matrices of entities and intentions.

harnesses the capabilities of the designed MID module to analyze and classify all requests, aiming to identify and prevent potential malicious ones.

B. Experimental Results

The influence of the MID module on the DIET classifier was examined utilizing k-fold cross-validation. The dataset was initially partitioned into K subsets of comparable sizes, termed folds. Subsequently, the model underwent iterative training K times, with each iteration designating one fold as the validation set and the remaining $K - 1$ folds as the training set. The performance evaluation ensued by averaging the outcomes across all K iterations. For sake of generality, in the proposed work the number of folds is set to 3, in accordance to the limited size of the produced dataset.

To provide a description of the classification efficacy pertaining to both entities and intentions, Fig. 3 illustrates the confusion matrices obtained from the results of the DIET classifier. These matrices detail the count of detected instances, inclusive of false positives and false negatives across all entity types and intention classes, respectively. In this representation, each row corresponds to predicted classes, while each column represents actual entities and user intentions. Specifically, 234 instances of *build* intentions were accurately classified out of 246 instances. Additionally, the DIET classifier successfully identifies 265 instances of malicious intentions out of 284, categorizing them under the *out-of-scope* class. Consequently, only 19 intentions evade malicious classification, eluding the scrutiny of the MID module.

To offer deeper insights, Fig. 4 depicts the distribution of intent prediction confidence score, shedding light on the reliability of the classification outcomes. Ranging from 0 to 1, higher scores indicate a closer alignment with the declared intentions. The reported findings reveal a notable confidence level exceeding 0.98 for accurately classified intentions, with 555 instances correctly identified at a confidence score of 1. In contrast, only 19 intentions received the highest confidence score among the incorrect classifications.

TABLE I
PRECISION, RECALL, F1-SCORE FOR THE INTENTIONS CLASSIFICATION.

Intent	Precision	Recall	F1-score
Welcome	0.96	1	0.98
Cancel	1	0.82	0.90
Thanks	0.92	0.96	0.94
Goodbye	1	0.77	0.87
Reject	0.96	1	0.98
Confirm	0.92	0.86	0.89
Build feedback	0.90	1	0.95
Build	0.92	0.95	0.93
Out-of-scope	0.94	0.95	0.94

The outcomes resulting from the intentions classification was evaluated using precision, recall, and F1-score, as detailed in Table I. Specifically:

- Precision measures the proportion of true positive predictions among all positive predictions made by the model.
- Recall measures the proportion of true positive instances that were correctly identified by the model out of all actual positive instances.
- F1-score measures the harmonic mean of precision and recall, providing an assessment of the overall performance of the model.




Fig. 4. Confidence distribution of predicted intentions.

The *out-of-scope* class consistently demonstrates high-performance metrics, affirming the efficacy of the proposed

MID module in detecting malicious intentions. Specifically, precision stands at 94%, recall at 95%, and the F1-score at 94%. Furthermore, the Rasa framework provides the overall accuracy for classification encompassing both entities and intentions. This accuracy represents the number of correct predictions, including true positives and true negatives, out of the total predictions. Across the entire set of 755 user requests, the observed outcome yields an average accuracy of 93% in correctly classifying intentions. The high accuracy value obtained, particularly across all metrics for the *out-of-scope* class, substantiates the classifier’s capability to detect malicious intentions. Subsequently, these intentions are efficiently intercepted and discarded, ensuring that only legitimate requests proceed to the translation module.

VI. CONCLUSIONS

This paper presented an approach designed to proactively detect malicious intentions at the intent processing stage, thereby preventing the formulation of network configurations or policies disruptive to the enterprise network. Implemented within the open-source IBN framework Lumi, integrated with a Rasa-based chatbot, the proposed approach extended the Information Extraction module with the novel MID module. This enhancement facilitated the creation of a security knowledge base to improve the classification of enterprise users’ intentions. The findings indicate that the conceived module effectively detects malicious intentions, thereby excluding them from the intent translation process and preventing potential reliability issues within the enterprise network. Future research endeavors will focus on implementing scheduled retrains of the DIET classifier, potentially leveraging user feedback to identify optimal times for dataset augmentation. Furthermore, we will explore more complex scenarios to validate the module in diverse environments, additionally investigating security countermeasures within subsequent IBN modules.

ACKNOWLEDGEMENTS

This work was supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, in the context of partnership on “Telecommunications of the Future” (PE00000001 - program “RESTART”, CUP: D93C22000910001), national center on “Sustainable Mobility” (CN00000023 - program “MOST”, CUP: D93C22000410001), and partnership on “Cybersecurity” (PE00000007 - program “SERICS”, CUP: D33C22001300002, project ISP5G+). It was also supported by the PRIN 2022 projects INSPIRE (grant no. 2022BEXMXN_01) and HORUS (grant no. 2022P44KA8) funded by the Italian MUR, by the HORIZON MSCA project BRIDGITISE (grant no. 101119554) and the HORIZON JU SNS project 6G-GOALS (grant no. 101139232), and by “The house of emerging technologies of Matera (CTEMT)” project funded by the Italian MIMIT.

REFERENCES

[1] T. Szigeti, D. Zacks, and F. Matthias, *Cisco Digital Network Architecture: Intent-based Networking for the Enterprise*. Cisco Press, 2018.

[2] A. Leivadeas and M. Falkner, “A Survey on Intent-Based Networking,” *IEEE Communications Surveys & Tutorials*, vol. 25, no. 1, pp. 625–655, 2023.

[3] A. S. Jacobs, R. J. Pfitscher, R. A. Ferreira, and L. Z. Granville, “Refining network intents for self-driving networks,” *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 5, p. 55–63, Jan. 2019.

[4] J. Lin, K. Dzevaroska, A. Tizghadam, and A. Leon-Garcia, “AppleSeed: Intent-Based Multi-Domain Infrastructure Management via Few-Shot Learning,” in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*, 2023, pp. 539–544.

[5] S. Ji, X. Li, P. Zhang, K. Liu, H. Dong, and Y. Zhang, “Intent-Driven QoS Control in Cloud-Network Integration Environment,” in *2022 IEEE International Conference on Satellite Computing (Satellite)*, 2022, pp. 52–53.

[6] A. S. Jacobs, R. J. Pfitscher, R. H. Ribeiro, R. A. Ferreira, L. Z. Granville, W. Willinger, and S. G. Rao, “Hey, Lumi! Using Natural Language for Intent-Based Network Management,” in *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, Jul. 2021, pp. 625–639.

[7] E. El-Rif, A. Leivadeas, and M. Falkner, “Intent Expression Through Natural Language Processing in an Enterprise Network,” in *2023 IEEE 24th International Conference on High Performance Switching and Routing (HPSR)*, 2023, pp. 1–6.

[8] M.-T.-A. Nguyen, S. B. Souihi, H.-A. Tran, and S. Souihi, “When NLP meets SDN : an application to Global Internet eXchange Network,” in *ICC 2022 - IEEE International Conference on Communications*, 2022, pp. 2972–2977.

[9] C. H. Cesila, R. P. Pinto, K. S. Mayer, A. F. Escallón-Portilla, D. A. A. Mello, D. S. Arantes, and C. E. Rothenberg, “Chat-IBN-RASA: Building an Intent Translator for Packet-Optical Networks based on RASA,” in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*, 2023, pp. 534–538.

[10] R. Caldelli, P. Castoldi, M. Gharbaoui, B. Martini, M. Matarazzo, and F. Sciarone, “On helping users in writing network slice intents through NLP and User Profiling,” in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*, 2023, pp. 545–550.

[11] M. Riftadi and F. Kuipers, “P4I/O: Intent-Based Networking with P4,” in *2019 IEEE Conference on Network Softwarization (NetSoft)*, 2019, pp. 438–443.

[12] J. Kim, E. Kim, J. Yang, J. Jeong, H. Kim, S. Hyun, H. Yang, J. Oh, Y. Kim, S. Hares, and L. Dunbar, “IBCS: Intent-Based Cloud Services for Security Applications,” *IEEE Communications Magazine*, vol. 58, no. 4, pp. 45–51, 2020.

[13] E. J. Scheid, C. C. Machado, M. F. Franco, R. L. dos Santos, R. P. Pfitscher, A. E. Schaeffer-Filho, and L. Z. Granville, “INSPIRE: Integrated NFV-based Intent Refinement Environment,” in *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 2017, pp. 186–194.

[14] K. Wang, H. Du, and L. Su, “Intent-Driven Network Slicing Security Provision and Management,” in *2023 IEEE 23rd International Conference on Communication Technology (ICCT)*, 2023, pp. 1318–1324.

[15] B. E. Ujcich and W. H. Sanders, “Data Protection Intents for Software-Defined Networking,” in *2019 IEEE Conference on Network Softwarization (NetSoft)*, 2019, pp. 271–275.

[16] G. Landeau, M. Avgeris, A. Leivadeas, and I. Lambadaris, “Security-Oriented Network Intent Placement using Particle Swarm Optimization,” in *2023 7th Cyber Security in Networking Conference (CSNet)*, 2023, pp. 19–22.

[17] J. J. Diaz Rivera, M. Afaq, and W.-C. Song, “Blockchain and Intent-Based Networking: A Novel Approach to Secure and Accurate Network Policy Implementation,” in *2023 24th Asia-Pacific Network Operations and Management Symposium (APNOMS)*, 2023, pp. 77–82.

[18] N. Herbaut, C. Correa, J. Robin, and R. Mazo, “SDN Intent-based conformance checking: application to security policies,” in *2021 IEEE 7th International Conference on Network Softwarization (NetSoft)*, 2021, pp. 181–185.

[19] ITU-T, “Intent-based network management and orchestration for network slicing in IMT-2020 networks and beyond,” ITU, Recommendation Y.3161, 12 2023.

[20] T. Bocklisch, J. Faulkner, N. Pawlowski, and A. Nichol, “Rasa: Open source language understanding and dialogue management,” *arXiv preprint arXiv:1712.05181*, 2017.

- [21] B. Nour, M. Pourzandi, and M. Debbabi, "A Survey on Threat Hunting in Enterprise Networks," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, pp. 2299–2324, 2023.