



Proposal of a BI/SSBI System for Knowledge Management of the Traffic of a Network Infrastructure – A University of Trás-os-Montes e Alto Douro Case Study

José Bessa¹, Frederico Branco^{1,2(✉)}, António Rio Costa¹,
Ramiro Gonçalves^{1,2}, and Fernando Moreira³

¹ University of Trás-os-Montes e Alto Douro, Vila Real, Portugal
jmiguelbessal6@gmail.com,

{fbranco, acosta, ramiro}@utad.pt

² INESC TEC and UTAD, University of Porto, Porto, Portugal

³ IJP, REMIT, University Portucalense, Porto, Portugal
fmoreira@upt.pt

Abstract. The data volume in organizations has grown at an ever-increasing rate and part of it is associated with the operation of the network infrastructure used to support systems and applications. Given the importance of this infrastructure for organizations and the large amount of data that their operation originates, it is fundamental to manage and monitor it so that it can perform well. The previous concern is transversal to higher education institutions, where the research team assumed as important the development of a BI/SSBI system for the Informatics and Communications Services of the institution where it operates (UTAD), which allows the managing of all data volume; that enables it to be transformed into information and knowledge, which are fundamental resources to support decision-making processes. The purpose of this article is to demonstrate the usefulness of a BI/SSBI system in the described context, therefore a system of this type is presented, along with the adopted technologies, the performed tests and the obtained results.

1 Introduction

Nowadays, Information Systems (IS) are vital for the success and competitiveness of organizations and most of IS do not work in an isolated way [1, 2]. To ensure their integration with other systems, organizations make use of network infrastructures. The use of this type of infrastructure can be found more and more frequently in all sorts of organizations, independent of their business or application area, and it is fundamental to support the proper functioning of systems and applications; consequently, it needs to be monitored and managed [3].

Network infrastructure monitoring is essential for its management processes, as it allows, at an early stage, to identify trends and patterns in data traffic and equipment behavior (e.g. anomalies, performance issues), in order to ensure a reliable, robust, highly productive infrastructure that guarantees quality of service for its users [4].

In order to expedite the treatment of the generated data by the operation of the network infrastructure and to obtain Knowledge and Intelligence, a suitable system type is used, called the Business Intelligence/Self-Service Business Intelligence (BI/SSBI) system. It has a set of mechanisms that help extracting data from the operational servers—their standardization and storage in a centralized repository, their treatment through specific mechanisms and its availability—using mainly visual elements, through the application of Data Visualization (DV) techniques, so that its interpretation become more easier and intuitive [5].

Within the scope of the present work, a BI/SSBI system was conceived for the monitoring and management activities of the network infrastructure applied to the case of a university, in this case the University of Trás-os-Montes and Alto Douro (UTAD). With the purposed system, the Informatics and Communications Services of the institution (SIC-UTAD), which were responsible for this infrastructure, now have a set of views and dashboards for some of the most critical services of this type of infrastructure (RADIUS, DNS and DHCP); this allows them to support decision-making and make strategic decisions.

This article is composed of four sections: after this section, which presents the introduction and the structure of the document, Sect. 2 presents the conceptual approach of its supporting themes. Section 3 presents the case study, its contextualization and the purposed system (technological options, the system architecture and tests/results). The article concludes with all of the considerations extracted from the work presented and perspectives for future work.

2 Conceptual Approach

2.1 Big Data and the Traffic of a Network Infrastructure

Currently, most organizations have network communications infrastructures to support their different activities; actually, many of them can support their business and activities only if they have an infrastructure like that [6].

Given the importance of network infrastructures, it becomes important to manage and monitor them. Monitoring is crucial for management operations, since it allows, in an early stage, to identify trends and patterns in data traffic and in the behavior of network devices (anomalies, performance issues), in order to guarantee reliability, robustness and high productivity of the infrastructure, and also ensure the quality of service for those who use it. Ergo, monitoring is a fundamental element in network management, since it's the management characterized by activities, methods, processes and tools that supports the operation and maintenance of the network infrastructure [7].

Similar to what happens with all other organizational data, data from network infrastructures have grown exponentially and this growth is associated with the Big Data concept. For management processes where decision-making entails major consequences for the organization, it is of all interest to reach a higher level of abstraction, namely Knowledge, which goes through the analysis and extraction of inferences from the available information. However, in order to reach this abstraction level, it is necessary to apply methods and technologies for data treatment and analysis, through the recurrent use of BI/SSBI systems [8].

2.2 BI/SSBI System

Decision-making implies the existence of organizational intelligence, a cognitive state that is achieved after gathering, analyzing and disseminating information, gaining new opportunities and reacting/adapting in a timely manner [9]. However, given that organizations have an increasing data volume, there is a need to use information technologies for their storage, availability and processing. To withdraw organizational intelligence and to respond to this need, the concept of BI emerged [10].

For this investigation, the definition of BI followed was that elaborated by Sharda et al. [11], in which it is stated that BI is an aggregator and umbrella term that combines architectures, data sources, analytical tools, applications and methodologies. It integrates data analysis with decision support systems to provide information to all the people in the organization who need to make tactical and strategic decisions.

Recently, associated with the BI concept, a new tendency or an evolution of this concept has surfaced, called SSBI. This concept focuses on four main objectives, namely [12] quick access to data sources for the creation of reports and analysis, make BI tools easier to use and support data analysis without being a technology expert, simple and customizable interfaces by end users and easy to deploy and manage storage options. Given that both in the concept of BI and SSBI the dashboards production is paramount, the next to be explored is the importance of the DV concept.

2.3 Data Visualization

One of the most important modules of traditional BI systems is the DV component, given their importance for understanding and exploiting data through their visualization. The scientific area of DV bases it's the study of the development of techniques for data visual representation using computerized techniques to simulate scenarios and to detect trends, patterns, deviations, and exceptions that, at first sight, are hidden within the data [13].

One of the main factors for the use growth of this type of technique is related to people being able to identify patterns in visual representations easier than in textual representations, once the visual ones have more appealing properties such as color, size, form, location and orientation. Typically, tools that implement DV techniques apply analytical models and statistical functions, presenting the data graphically through interesting dashboards (static or dynamic) composed by dynamic tables, maps, graphs, diagrams, histograms, among others [14].

3 Proposed System

3.1 Contextualization and Needs

UTAD is a public Portuguese university, created in 1973 and established at Vila Real, and it is one of the main contributor institutions for the promotion and development of the region where it is located, both because of its scientific and technological contribution and for the entrepreneurship that encourage regional growth. Towards the case study of management and monitorization of network infrastructure presented in this

paper, at UTAD, the SIC-UTAD is the service responsible for that. The managing and monitoring activities encompass the responsibility to monitor/manage the network's equipment and the traffic that occurs.

Regarding the decision-making processes related to network infrastructure management, the following information needs (KPI's) are highlighted: which equipment operates by location (building) and who uses it; average delay of equipment response; geographical distribution of equipment (floors/areas); temporal evolution of the average load per Access Point (AP); access to internal and external services of UTAD; configuration problems of equipment; anomalies that affect network performance; security problems; equipment/service programming issues.

3.2 Technological Options

The changes that occur in the technology market are constant, so currently, within BI/SSBI, there is great tools diversity, some commercial/proprietary and others open source, each having unique characteristics to the others, which are interpreted as strengths and weaknesses, advantages and disadvantages, according to the needs that are intended to be met.

During this project, there was a paradigm shift at UTAD with regards to the software use, motivated by strategic and sustainability issues, and the institution preferred to use solutions aligned with the Free Open Source Software (FOSS) philosophy. The new adopted technological solution was the stack developed by Elastic[®].

The choice of this solution was motivated by the availability of capable tools for meeting the main needs of a typical BI/SSBI architecture (reviewed in chapter two); despite the short existence (created in 2010) it has a high degree of maturity, stability and robustness of its tools and plugins; advanced and specialized mechanisms for the logs treatment (e.g. operating system, network, websites, security and user support); a friendly environment for programmers and users; a large and proactive community and multiple contact forms that encourage the creation of synergies for the exchange of information and knowledge.

After the justification of the chosen technology has been given, the Elastic stack (formerly called ELK stack) composition is analyzed, highlighting the existence of three main tools: Elasticsearch, Logstash and Kibana [15].

Elasticsearch is a distributed NoSQL database engine for real-time research and analysis built on Apache Lucene[®], a full-text search engine library. Currently, this library is the one that has the most advanced full-text search engines, with better performance and with more resources, both open source and at a commercial level.

In addition to advanced full-text search engines, Elasticsearch can be described as a distributed database engine that is ideal for real-time analysis with the next abilities: scaling to hundreds of servers and petabytes of data structured and unstructured; horizontal scalability; georeferencing; asynchronous replication; fuzzy searches; configuring automatic statistics groups (ideal for debugging) [16]. As alternatives to Elasticsearch, other NoSQL solutions can be used, such as MongoDB, CouchDB, Redis or Cassandra.

Relative to Logstash, this is the responsible tool for the ETL process; its main tasks are data collection and its normalization through the assignment of a uniform structure

(specified in the mapping), allowing its storage and future use. The configuration of this tool is composed by the definition of inputs (filters and codecs, definition of how data transformation and storage will be done) and outputs (connection establishment where the data resulting from the ETL process will be stored) [17]. As alternatives to Logstash, other ETL/parser solutions can be used, such as Fluentd, Kafka, Flume and Graylog Collector.

The last principal tool of this stack is Kibana, a DV web platform developed in AngularJS and highly configurable JavaScript that is used for the analysis and treatment of large volumes of data from Elasticsearch indexes, quickly detecting patterns and irregularities in them. This tool is based on widgets and offers many possibilities of interactive illustrations to shape the data, such as graphs (e.g. lines, area, bars, circular), maps, diagrams and tables, all being presented while taking into account their timestamp [18]. As alternatives to Kibana, can be used other DV solutions can be used, such as Grafana, Graylog, Datadog and Zabbix.

The main gap that this web platform has is related to security, because it does not have authentication mechanisms, which requires the use of additional plugins and tools that guarantee this property [19]. At the moment, this platform has not yet reached a popularity that lives up to its value, largely because its features can only be applied to the Elasticsearch indexes.

3.3 System Architecture

The BI/SSBI system proposed presents an architecture divided into four levels/layers: Data, ETL, DV and Knowledge/Intelligence, being the latest the highest abstraction level (Fig. 1). From the technological point of view, the adopted solutions to support the proposed architecture fit into the FOSS philosophy. For its implementation, virtual machines in datacenter servers of UTAD were created, all of them being Linux[®] with the Ubuntu 14.04.4 amd64 distribution. As a BI/SSBI solution, the most stable versions of the main Elastic stack tools were used, namely Elasticsearch 2.3 (database engine), Logstash 2.3 (ETL) and Kibana 4.5 (DV).

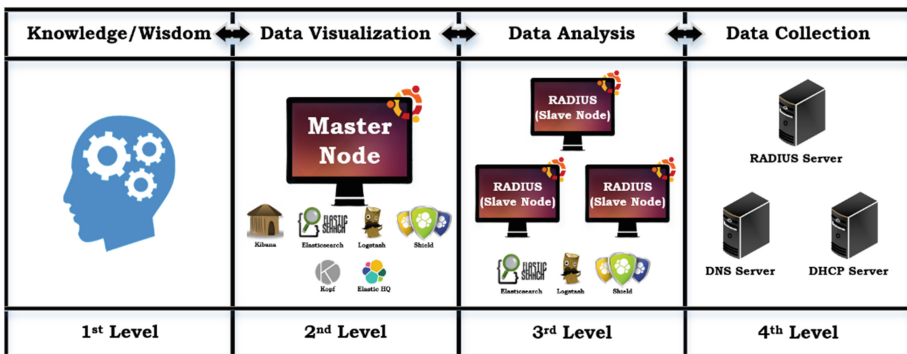


Fig. 1. Composition of the BI/SSBI system architecture proposed for SIC-UTAD.

In what concerns to the lowest level of the architecture, the used data sources were the logs from the services: (1) RADIUS, (2) DHCP and (3) DNS. This architecture has been designed to the data-tier systems (data sources) being plug-and-play, in order to be flexible in their introduction/removal, maintaining the isolation and ensuring that the normal performance of their activities and their relationship with other systems is not compromised.

Above the Data level is the ETL level. At this level, the data is collected from the lower level and a homogenized structure for each type is created, so as to reduce data waste as much as possible, reduce treatment errors (structure inconsistencies, such as more or less fields than what they really are) and streamline their processing and indexing in the data storage structures.

For the ETL level, a virtual machine for each data type was created (maintaining its isolation), an Elasticsearch instance being installed to create the repository where the data will remain and a Logstash instance used to create the responsible parser for data extraction from the data sources, transforming them to assign a uniform structure, and loading into the repository.

The security component is one of the biggest weaknesses of Elasticsearch, so in order to reduce its exposure to vulnerabilities, a Shield instance has been installed on each machine. This plugin allows the protection of Elasticsearch data repositories by implementing mechanisms such as communications encryption (SSL/TLS), user role creation, and IP filtering.

Continually rising in the architecture level, the level of DV is reached where data analysis is being done. It is at this level that views and dashboards are created, which contain critical information for each system that is being monitored and managed. At this level, only one master machine was created and the same technologies presented in the low level were installed, with the addition of Kibana, which is used for data analysis, and the Kopf and ElasticHQ plugins used for the visual management of the data cluster and the used resources of the machines.

Finally, at the highest level of the proposed BI/SSBI system architecture we find Knowledge and Organizational Intelligence. It is also at this level that we find the managers, who are the ones who will have access to all the obtained dashboards and reports with the lower levels implementation.

3.4 Tests and Results

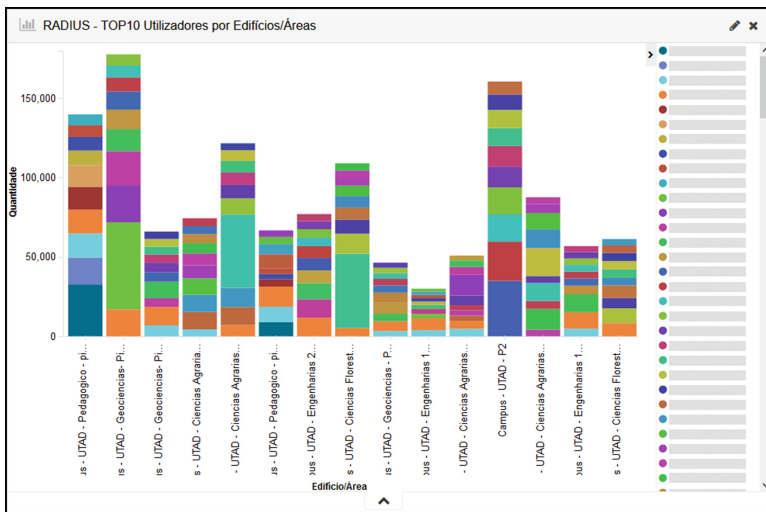
The tests performed focus on the analytical and DV platform (Kibana) of the proposed system and had as objective the construction of visualizations and dashboards that could allow us to conclude whether it would be possible to extract findings that answer to the previously identified management needs.

Regarding to the data that will be exposed, since these are real, and in some cases, may jeopardize the privacy and confidentiality of members of the academic community, the temporal space to which the collected data are associated is variable for the different types. This option was also made because the institution considered that this would be another way to mitigate possible security risks that would disrupt and compromise the normal operation of the concerned services.

3.4.1 RADIUS Service

The data that fed the indicators and dashboards that will be presented next are composed by a sample of 35.517.395 logs from the RADIUS service, produced by 1,419 APs and framed in the time interval of 01/01/2015 to 10/31/2016.

One of the examples of possible indicators to construct this type of service is the one that can be observed through Graphic 1, and which shows the mobility of the users by the buildings/areas of the UTAD campus with the greatest number of authentications (Top 15). This indicator also allows us to identify which are the users that authenticate most often in the APs of the Top 15 buildings/areas. With this indicator we could yet understand if there are users with abnormal behaviors, as a result of attack attempts to the network infrastructure, carried out by themselves or by others who maliciously infected the devices on which they were authenticated (e.g. 4th user in the legend of this graph that appears with a high number of accesses in practically all Top 15 locations).



Graphic 1. Top 10 users in the 15 buildings/areas with the greatest number of authentications.

Regarding to the production of dashboards for the RADIUS service, Fig. 2 presents one related to using the profile of the network infrastructure by a specific student. For this user, the RADIUS service registered 1.365 authentications, scattered over 82 APs of the UTAD campus. Two authentication domains were used, namely “alunos.utad.pt” (VPN) (38) and “utad.eu” (2). The remaining authentications were made using only the username, which concludes that this user has resorted to three different authentication mechanisms.

From the analysis of the temporal evolution of the authentications number, it is possible to conclude that this user did not authenticate in the UTAD network before September 2015, but, from that time on, maintained a constant access, reaching the

maximum of authentications in December 2015 (388) and in 2016 in the months of February (183) and June (72). The buildings/areas of the campus where the network infrastructure is most used (Top 10) were Engineering I (527) and Library (152). In concern to roaming situations, no authentication related to UTAD’s network infrastructure access has been registered through another institution.

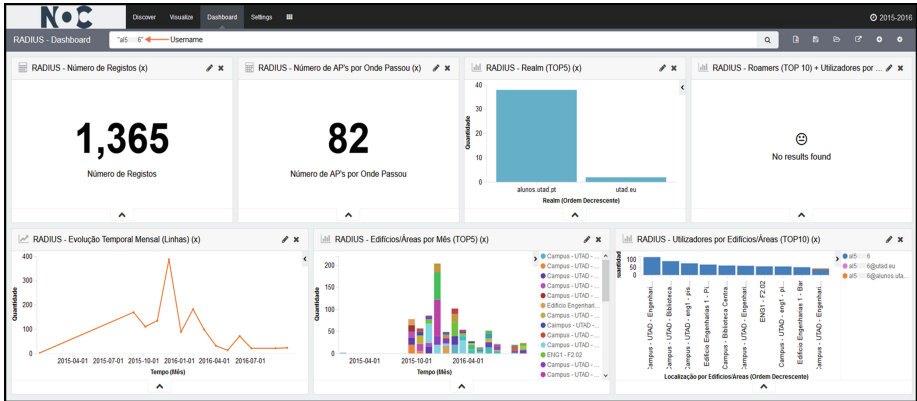


Fig. 2. Dashboard for a user’s authentication profile.

3.4.2 DHCP Service

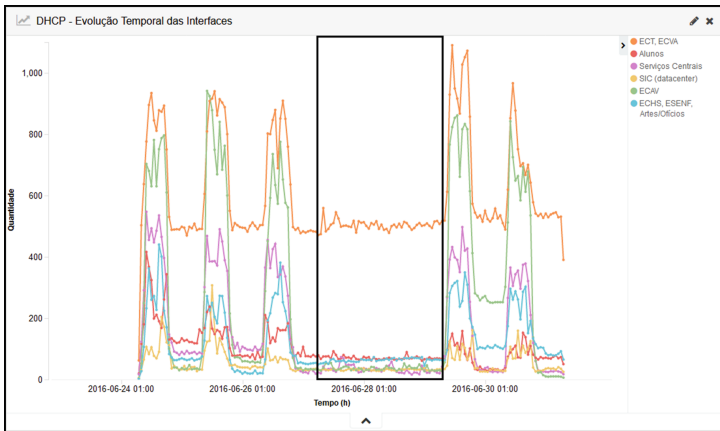
The data that fed the indicators and dashboards that will be presented next are composed by a sample of 368.570 logs from the DHCP service and framed in the time interval of 06/24/2016 to 07/01/2016.

UTAD’s physical network infrastructure is divided into several virtual networks (VLANs), logically independent and accessible through an interface. This division of the network allows the traffic segregation to facilitate infrastructure management and to increase the security of the communications, users and data that goes through there.

Graphic 2 allows us to perceive the temporal evolution (in hours) of the behavior manifested by the DHCP service for each one of the interfaces. The analysis of this graph reveals that the interface that had the highest peak of requests per hour was “ECT, ECVA” with 1,091 (06/29/2016 at 11:00 AM), followed by the descending order of the maximum number of requests per hour: “ECAV” with 942 (06/25/2016 at 10:00 AM), “Serviços Centrais” with 548 (06/24/2016 at 10:00 AM), “ECHS, SENF, Artes/Ofícios” with 441 (06/24/2016 at 3:00 pm), “Students” with 417 (06/24/2016 at 10:00 AM) and lastly, the “SIC (datacenter)” interface, with 308 (06/25/2016 at 12:00).

In general, the behavior presented by this service was the expected one, with a greater number of requests to be made during the day, with a considerable reduction during the night, except for the interface “ECT, ECVA”, which, despite the reduction of requests in the evening, presents very high values in comparison with the other interfaces; this may suggest the existence of clients with configuration problems of this service or the existence of configuration problems of the own VLAN.

Still analyzing the behavior of the DHCP service, it is verified that the number of requests of the interfaces registered in the interval of time between 07:00 AM of 06/27/2016 and 07:00 AM of 06/29/2016 (black rectangle) is out of the behavioral pattern, which may mean that logs have been rejected by the server or, in a more severe scenario, an attack attempt may have occurred on this service, where an external machine attempted to provide DHCP service being passed through the authentic UTAD server that actually provides this service (e.g. man-in-the-middle attack).



Graphic 2. Time evolution of the number of DHCP requests for each interface (VLAN).

The dashboard of Fig. 3 presents indicators related to the DHCP requests associated to an interface (VLAN), which allow the extraction of number of requests, time evolution (realizing if anomalous situations occurred), which are the main IP clients and Media Access Control (MAC) (Top 10), which clients have made more requests, what kind of message is used most often and which are the most recurring messages per IP client. Through the analysis of this dashboard, it is possible to observe that the DHCP service had an expected performance and it is still possible to observe that a specific IP client had an abnormal behavior (40.000 requests, more than any other of Top 10) for the activities it performs (SIC-UTAD support services).

3.4.3 DNS Service

The data that fed the indicators and dashboards that will be presented next are composed by a sample of 22.622.120 logs from the DNS service and framed in the time interval of 11/02/2016 to 11/05/2016.

The evolution of the number of DNS requests is one of the indicators that allow the analysis of the behavioral pattern of the service, to see if it is functioning as expected or if there is any anomaly, Graphic 3 being created for that purpose. The analysis of this graph reveals that the service has an expected behavior, which means there are a greater number of DNS requests during the day in relation to the night period. On the other hand, it is verified the existence of more orders than expected, due to possible

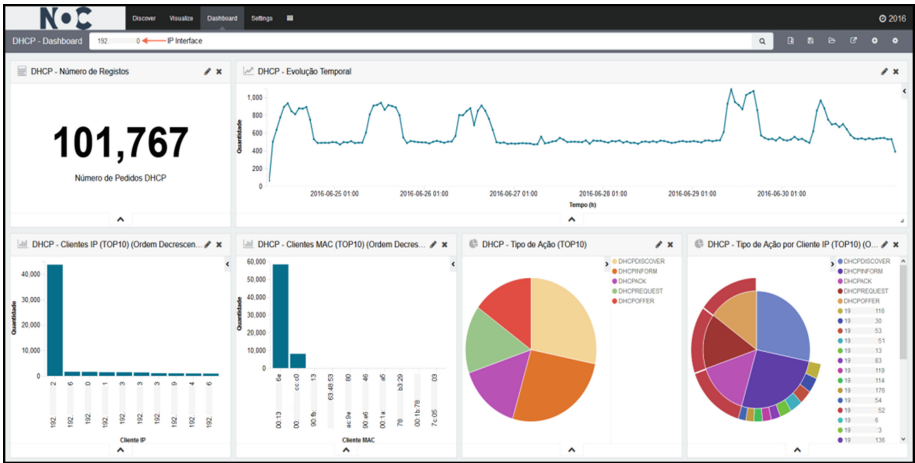



Fig. 3. Dashboard allusive to requests for an interface (VLAN).



Graphic 3. Evolution of the number of DNS requests over time.

configuration errors of customers, reaching by the hour (at the highest peaks) 831.380 (11/03/2016 at 3 pm) and 813.934 (11/04/2016 at 12:00 pm).

To finalize the presentation of the results obtained for the DNS service, a dashboard (Fig. 4) with oriented indicators to the management of www.utad.pt domain is presented.

The analysis of the presented dashboard in Fig. 4 allows us to verify that there were 685.149 requests for the domain “utad.pt”, with an hourly evolution of the number of requests to remain in the order of 10.000. Relative to the destiny subdomains, the main ones are related to the institutional site (mainly redirecting pointers), the Lightweight Directory Access Protocol (LDAP) management platform, the teaching support system (SIDE) and the printer management platforms.

For the considerate domain, “A” was the main registration type with 62% of requests, which means that most of the requests were to use one of the primary functions of the DNS service, the translation of IP addresses into domain names.

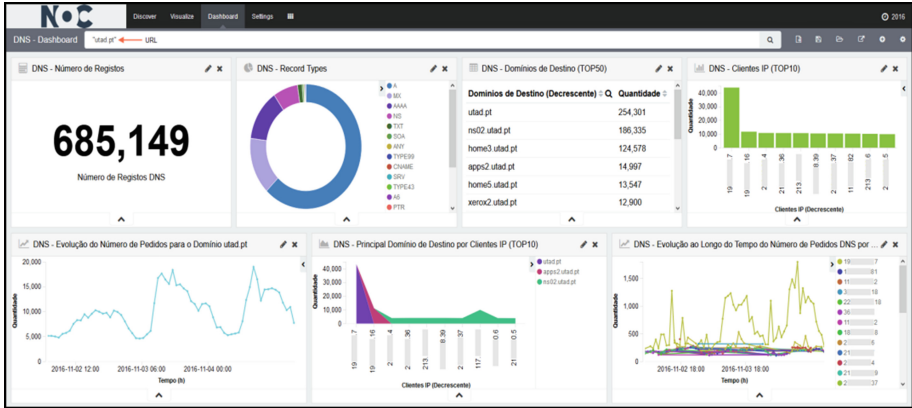


Fig. 4. Dashboard for the management of “utad.pt” domain.

The second registration type with more requests was the “MX” with around 15% (referring to the intention to use e-mail services) and the third type was the “AAAA” type, which proved the existence of IPv6 requests (13%).

For Top 10 IP clients, an average of ~10.500 requests were presented for each one, except for the client whose requests exceeded by more than four times this value. Looking at the evolution of the number of requests by IP clients, it is noticed that all have a value close to the number of requests per hour (between 150 and 300), except for the IP client with the most requests, which reached 1,800 requests (maximum) per hour.

4 Final Considerations

Nowadays, network infrastructures are essential to support the activity of the organizations, and as such, it is important that these are monitored and managed in the best way possible, to guarantee the quality of service for those who use them. This context is also applied to the UTAD and to managing the domains of their IT and Communications Services, taunting thus the need to create a system which could allow the transformation of data in information and the obtaining of knowledge and intelligence about the network infrastructure, using concepts related with BI/SSBI and others referred to in the conceptual approach of this article.

For the implementation of the proposed system, a solution associated with FOSS paradigm was used, the stack developed by Elastic[®]. In what concerns the validation of the proposed system, a battery of tests was made, with the objective of verify if the system could allow the production of results which satisfied the needs of management of the SIC-UTAD related to network traffic and services/equipment that compose the infrastructure itself. The data used for the tests is related to RADIUS, DHCP and DNS services. The incorporation of the proposed system made possible the modernization of the SIC-UTAD in what concerns the adoption of new technologies, allowing the achievement of productivity gains, improving their efficiency, and aiding the mitigation of possible risks and the innovation of decision-making processes.

In terms of future work, the team involved in this research considers it appropriate to use the contributions obtained as a starting point for new projects, such as the addition of new functionalities to the system, such as user creation (aligned to the federation scheme of UTAD), the creation and establishment of thresholds for the alarm component, the addition of new modules for data import/export and creation of a notification module (e.g. e-mail, SMS), and the exploiting of the creation of dashboards, views and reports, from both data sources contemplated until now, as well as new ones that may arise (e.g., SNMP, firewalls).

References

1. Santos, V., Pereira, J., Martins, J., Gonçalves, R., Branco, F.: Creativity as a key ingredient of information systems. In: Mejia, J., Munoz, M., Rocha, Á., Calvo-Manzano, J. (eds.) Trends and Applications in Software Engineering: Proceedings of the 4th International Conference on Software Process Improvement CIMPS'2015, pp. 283–291. Springer, Cham (2016)
2. Gonçalves, R., Martins, J., Branco, F., Perez-Cota, M., Oliveira, A.-Y.M.: Increasing the reach of enterprises through electronic commerce: a focus group study aimed at the cases of Portugal and Spain. *Comput. Sci. Inf. Syst.* **13**, 927–955 (2016)
3. Taylor, M.J., Gresty, D., Askwith, R.: Knowledge for network support. *Inf. Softw. Technol.* **43**, 469–475 (2001)
4. Branco, F., Martins, J., Gonçalves, R., Bessa, J., Costa, A.: A decision support platform for IT infrastructure management: the university of Trás-os-Montes e Alto Douro services of information and communications case study. In: 10th Iberian Conference on Information Systems and Technologies (CISTI), 2015, pp. 1–7. IEEE (2015)
5. Branco, F., Gonçalves, R., Martins, J., Cota, M.P.: Decision support system for the agri-food sector-the sousacamp group case. In: WorldCIST 2015, pp. 553–563 (2015)
6. Sterbenz, J.P.G., Hutchison, D., Çetinkaya, E.K., Jabbar, A., Rohrer, J.P., Schöller, M., Smith, P.: Resilience and survivability in communication networks: strategies, principles, and survey of disciplines. *Comput. Netw.* **54**, 1245–1265 (2010)
7. Lee, S., Levanti, K., Kim, H.S.: Network monitoring: present and future. *Comput. Netw.* **65**, 84–98 (2014)
8. Chen, H., Chiang, R.H.L., Storey, V.C.: Business intelligence and analytics: from big data to big impact. *MIS Q.* **36**, 1165–1188 (2012)
9. Branco, F., Martins, J., Gonçalves, R.: Das Tecnologias e Sistemas de Informação à Proposta Tecnológica de um Sistema de Informação Para a Agroindústria: O Grupo Sousacamp. RISTI - Revista Ibérica de Sistemas e Tecnologias de Informação, pp. 18–32. RISTI (2016)
10. Ramakrishnan, T., Jones, M.C., Sidorova, A.: Factors influencing Business Intelligence (BI) data collection strategies: an empirical investigation. *Decis. Support Syst.* **52**, 486–496 (2012)
11. Sharda, R., Delen, D., Turban, E., King, D.: Business Intelligence: A Managerial Perspective on Analytics. Pearson Education Limited, New York (2015)
12. Stone, M., Woodcock, N.: Interactive, direct and digital marketing: a future that depends on better use of business intelligence. *J. Res. Interact. Mark.* **8**, 4–17 (2014)
13. Jang, Y., Ebert, D.S., Gaither, K.: Time-varying data visualization using functional representations. *IEEE Trans. Vis. Comput. Graph.* **18**, 421–433 (2012)
14. Janvrin, D.J., Raschke, R.L., Dilla, W.N.: Making sense of complex data using interactive data visualization. *J. Account. Educ.* **32**, 31–48 (2014)

15. Chhajed, S.: Learning ELK Stack. Packt Publishing, Birmingham (2015)
16. Kononenko, O., Baysal, O., Holmes, R., Godfrey, M.W.: Mining modern repositories with elasticsearch. In: Proceedings of the 11th Working Conference on Mining Software Repositories, pp. 328–331 (2014)
17. Sanjappa, S., Ahmed, M.: Analysis of Logs by Using Logstash, pp. 579–585. Springer, Singapore (2017)
18. Gupta, Y.: Kibana Essentials. Packt Publishing, Birmingham (2015)
19. Prakash, T., Kakkar, M., Patel, K.: Geo-identification of web users through logs using ELK stack. In: 2016 6th International Conference - Cloud System and Big Data Engineering (Confluence), pp. 606–610 (2016)