NORIA UI

Efficient Incident Management on Large-Scale ICT Systems Represented as Knowledge Graphs

GRASEC – ARES 2024

Lionel Tailhardat, Yoan Chabot, Antoine Py, Perrine Guillemette

Orange & EURECOM

Aug 01, 2024

Context & motivations – Efficiency in incident management with KGs

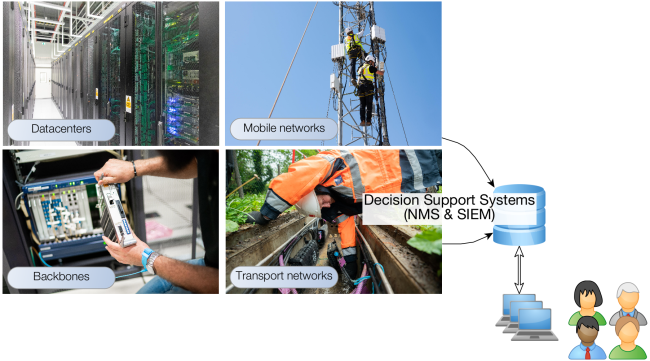

ICT systems Large-scale networks involve combining information from multiple monitoring tools.

Efficiency Handling complexity with AI tools (user and entity behavior analytics, anomaly detection, root cause analysis),

and unified view of ICT systems (breaking down technical silos) using knowledge graphs (KGs).

DSS design Lack of necessary hindsight to leverage KGs within Network Monitoring System (NMS) and Security Information & Event Management (SIEM) solutions.

NetOps/SecOps We must remain aligned with the business requirements in any proposal for the evolution of the information system.

Context & motivations – Efficiency in incident management with KGs

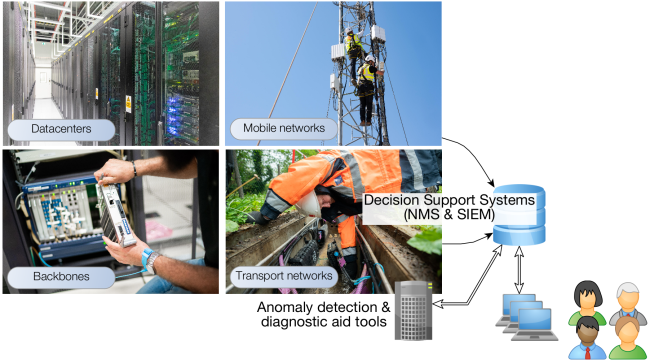

Datacenters
Mobile networks
Backbones
Transport networks
Decision Support Systems (NMS & SIEM)
Anomaly detection & diagnostic aid tools

ICT systems Large-scale networks involve combining information from multiple monitoring tools.

Efficiency Handling complexity with AI tools (user and entity behavior analytics, anomaly detection, root cause analysis), and unified view of ICT systems (breaking down technical silos) using knowledge graphs (KGs).

DSS design Lack of necessary hindsight to leverage KGs within Network Monitoring System (NMS) and Security Information & Event Management (SIEM) solutions.

NetOps/SecOps We must remain aligned with the business requirements in any proposal for the evolution of the information system.

Context & motivations – Efficiency in incident management with KGs

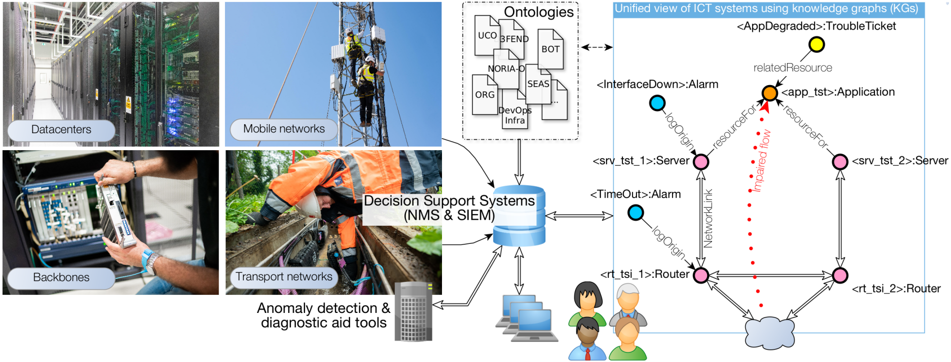

ICT systems Large-scale networks involve combining information from multiple monitoring tools.

Efficiency Handling complexity with AI tools (user and entity behavior analytics, anomaly detection, root cause analysis),

and unified view of ICT systems (breaking down technical silos) using knowledge graphs (KGs).

DSS design Lack of necessary hindsight to leverage KGs within Network Monitoring System (NMS) and Security Information & Event Management (SIEM) solutions.

NetOps/SecOps We must remain aligned with the business requirements in any proposal for the evolution of the information system.

Context & motivations – Efficiency in incident management with KGs

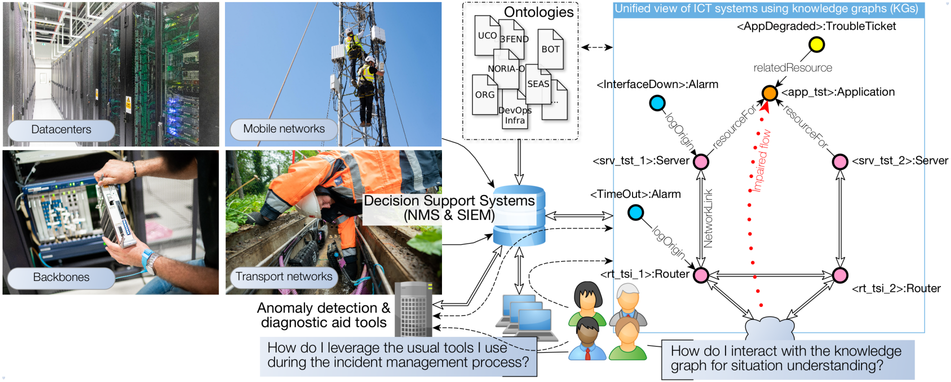

How do I leverage the usual tools I use during the incident management process?

How do I interact with the knowledge graph for situation understanding?

ICT systems Large-scale networks involve combining information from multiple monitoring tools.

Efficiency Handling complexity with AI tools (user and entity behavior analytics, anomaly detection, root cause analysis),

and unified view of ICT systems (breaking down technical silos) using knowledge graphs (KGs).

DSS design Lack of necessary hindsight to leverage KGs within Network Monitoring System (NMS) and Security Information & Event Management (SIEM) solutions.

NetOps/SecOps We must remain aligned with the business requirements in any proposal for the evolution of the information system.

Problem statement – Smooth integration of KGs in NMS/SIEM DSSs

Information accessibility How can we ensure intuitive and efficient exploration of KGs?

i.e. quick and limited access to only relevant information, given that the multiple knowledge facets needed to be represented for situation understanding pose a limit to data browsing without a deep understanding of the ontologies at work, especially when short response time is imposed by Service Level Agreements (SLAs).

Situation understanding How can we leverage a combination of decision support algorithms?

i.e. using different approaches and models to cover a wide range of system behaviors and support efficient decision-making, notably to ensure reducing cognitive load with alarm grouping.

Continuity of operational tasks How to ensure that a novel design best align with well-established incident management & response processes?

Approach – Towards functional specifications for a next-gen NMS/SIEM Decision Support System (DSS)

Design by formalizing DSS requirements in terms of ergonomics and functions (UI/UX), based on requirements from a panel of 16 NetOps/SecOps experts from Orange [1,2].

Implement a client-server software architecture going beyond the simple data exposition from a knowledge graph with:

- Synergistic reasoning [3] for the combination of various diagnostic AI techniques
- Interaction mechanisms for exploratory analysis of multi-layered systems.

Evaluate the solution based on UI/UX evaluation by users in operational situations (incident diagnosis scenario & SUS survey [4]), and performance analysis (platform analytics & telemetry data).

Problem statement – Smooth integration of KGs in NMS/SIEM DSSs

Information accessibility How can we ensure intuitive and efficient exploration of KGs?

i.e. quick and limited access to only relevant information, given that the multiple knowledge facets needed to be represented for situation understanding pose a limit to data browsing without a deep understanding of the ontologies at work, especially when short response time is imposed by Service Level Agreements (SLAs).

Situation understanding How can we leverage a combination of decision support algorithms?

i.e. using different approaches and models to cover a wide range of system behaviors and support efficient decision-making, notably to ensure reducing cognitive load with alarm grouping.

Continuity of operational tasks How to ensure that a novel design best align with well-established incident management & response processes?

Approach – Towards functional specifications for a next-gen NMS/SIEM Decision Support System (DSS)

Design by formalizing DSS requirements in terms of ergonomics and functions (UI/UX), based on requirements from a panel of 16 NetOps/SecOps experts from Orange [1,2].

Implement a client-server software architecture going beyond the simple data exposition from a knowledge graph with,

- Synergistic reasoning [3] for the combination of various diagnostic AI techniques,
- Interaction mechanisms for exploratory analysis of multi-layered systems.

Evaluate the solution based on UI/UX evaluation by users in operational situations (incident diagnosis scenario & SUS survey [4]), and performance analysis (platform analytics & telemetry data).

Problem statement – Smooth integration of KGs in NMS/SIEM DSSs

Information accessibility How can we ensure intuitive and efficient exploration of KGs?

i.e. quick and limited access to only relevant information, given that the multiple knowledge facets needed to be represented for situation understanding pose a limit to data browsing without a deep understanding of the ontologies at work, especially when short response time is imposed by Service Level Agreements (SLAs).

Situation understanding How can we leverage a combination of decision support algorithms?

i.e. using different approaches and models to cover a wide range of system behaviors and support efficient decision-making, notably to ensure reducing cognitive load with alarm grouping.

Continuity of operational tasks How to ensure that a novel design best align with well-established incident management & response processes?

Approach – Towards functional specifications for a next-gen NMS/SIEM Decision Support System (DSS)

Design by formalizing DSS requirements in terms of ergonomics and functions (UI/UX), based on requirements from a panel of 16 NetOps/SecOps experts from Orange [1,2].

Implement a client-server software architecture going beyond the simple data exposition from a knowledge graph with,

- Synergistic reasoning [3] for the combination of various diagnostic AI techniques,
- Interaction mechanisms for exploratory analysis of multi-layered systems.

Evaluate the solution based on UI/UX evaluation by users in operational situations (incident diagnosis scenario & SUS survey [4]), and performance analysis (platform analytics & telemetry data).

Problem statement – Smooth integration of KGs in NMS/SIEM DSSs

Information accessibility How can we ensure intuitive and efficient exploration of KGs?

i.e. quick and limited access to only relevant information, given that the multiple knowledge facets needed to be represented for situation understanding pose a limit to data browsing without a deep understanding of the ontologies at work, especially when short response time is imposed by Service Level Agreements (SLAs).

Situation understanding How can we leverage a combination of decision support algorithms?

i.e. using different approaches and models to cover a wide range of system behaviors and support efficient decision-making, notably to ensure reducing cognitive load with alarm grouping.

Continuity of operational tasks How to ensure that a novel design best align with well-established incident management & response processes?

Approach – Towards functional specifications for a next-gen NMS/SIEM Decision Support System (DSS)

Design by formalizing DSS requirements in terms of ergonomics and functions (UI/UX), based on requirements from a panel of 16 NetOps/SecOps experts from Orange [1,2].

Implement a client-server software architecture going beyond the simple data exposition from a knowledge graph with,

- Synergistic reasoning [3] for the combination of various diagnostic AI techniques,
- Interaction mechanisms for exploratory analysis of multi-layered systems.

Evaluate the solution based on UI/UX evaluation by users in operational situations (incident diagnosis scenario & SUS survey [4]), and performance analysis (platform analytics & telemetry data).

Problem statement – Smooth integration of KGs in NMS/SIEM DSSs

Information accessibility How can we ensure intuitive and efficient exploration of KGs?

i.e. quick and limited access to only relevant information, given that the multiple knowledge facets needed to be represented for situation understanding pose a limit to data browsing without a deep understanding of the ontologies at work, especially when short response time is imposed by Service Level Agreements (SLAs).

Situation understanding How can we leverage a combination of decision support algorithms?

i.e. using different approaches and models to cover a wide range of system behaviors and support efficient decision-making, notably to ensure reducing cognitive load with alarm grouping.

Continuity of operational tasks How to ensure that a novel design best align with well-established incident management & response processes?

Approach – Towards functional specifications for a next-gen NMS/SIEM Decision Support System (DSS)

Design by formalizing DSS requirements in terms of ergonomics and functions (UI/UX), based on requirements from a panel of 16 NetOps/SecOps experts from Orange [1,2].

Implement a client-server software architecture going beyond the simple data exposition from a knowledge graph with,

- Synergistic reasoning [3] for the combination of various diagnostic AI techniques,
- Interaction mechanisms for exploratory analysis of multi-layered systems.

Evaluate the solution based on UI/UX evaluation by users in operational situations (incident diagnosis scenario & SUS survey [4]), and performance analysis (platform analytics & telemetry data).

Experimental design – Evaluation methodology and scenario

Framework Client-server software architecture + extended SUS survey + platform analytics + three step incident diagnosis scenario ...

Evaluation dataset Test data + production network data in a knowledge graph-based platform [5]

Model-based design Query the graph to retrieve anomalies and their context [1]

- k out-of n devices with faults
- User with unusual account rights
- Absence of traffic on an interface supposed to be active

Process mining Align a sequence of entities to activity models, then use this relatedness to guide the repair [6]

- (EnergyLoss)=>(TimeoutAlert)=>(LossOfSignal)
- (LoginFail)=>(LoginFail)=>(LoginFail)

Statistical learning Relate entities based on context similarities, then use this relatedness to alert and guide the repair [1]

- The hidden cause of the trouble ticket on server 1 is a “data leak” attack that started on server 2

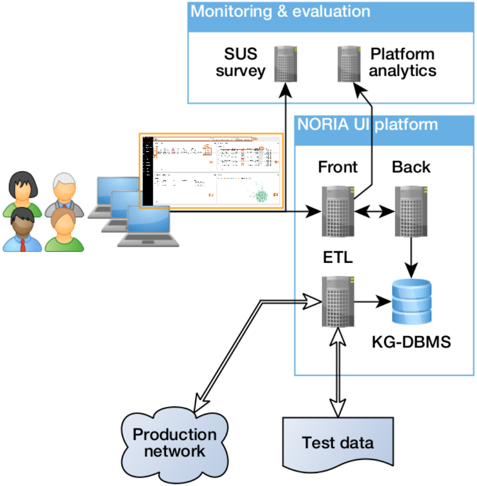

Experimental design – Evaluation methodology and scenario

Framework Client-server software architecture + extended SUS survey + platform analytics + three step incident diagnosis scenario ...

Evaluation dataset Test data + production network data in a knowledge graph-based platform [5]

Model-based design Query the graph to retrieve anomalies and their context [1]

- k out-of n devices with faults
- User with unusual account rights
- Absence of traffic on an interface supposed to be active

Process mining Align a sequence of entities to activity models, then use this relatedness to guide the repair [6]

- (EnergyLoss)=>(TimeoutAlert)=>(LossOfSignal)
- (LoginFail)=>(LoginFail)=>(LoginFail)

Statistical learning Relate entities based on context similarities, then use this relatedness to alert and guide the repair [1]

- The hidden cause of the trouble ticket on server 1 is a “data leak” attack that started on server 2

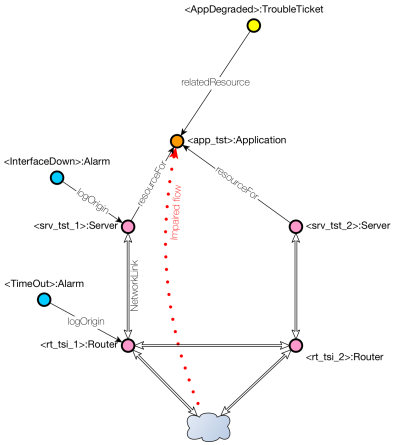

Experimental design – Evaluation methodology and scenario

Framework Client-server software architecture + extended SUS survey + platform analytics + three step incident diagnosis scenario ...

Evaluation dataset Test data + production network data in a knowledge graph-based platform [5]

Model-based design Query the graph to retrieve anomalies and their context [1]

- k out-of n devices with faults
- User with unusual account rights
- Absence of traffic on an interface supposed to be active

Process mining Align a sequence of entities to activity models, then use this relatedness to guide the repair [6]

- (EnergyLoss)=>(TimeoutAlert)=>(LossOfSignal)
- (LoginFail)=>(LoginFail)=>(LoginFail)

Statistical learning Relate entities based on context similarities, then use this relatedness to alert and guide the repair [1]

- The hidden cause of the trouble ticket on server 1 is a “data leak” attack that started on server 2

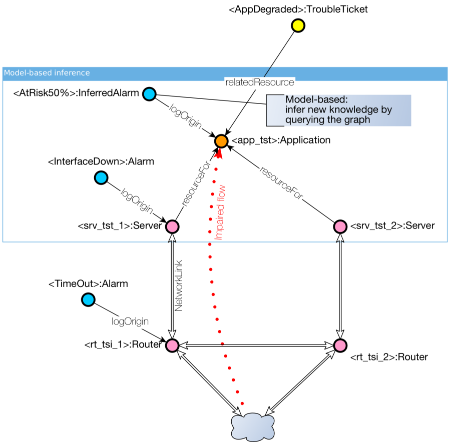

Experimental design – Evaluation methodology and scenario

Framework Client-server software architecture + extended SUS survey + platform analytics + three step incident diagnosis scenario ...

Evaluation dataset Test data + production network data in a knowledge graph-based platform [5]

Model-based design Query the graph to retrieve anomalies and their context [1]

- k out-of n devices with faults
- User with unusual account rights
- Absence of traffic on an interface supposed to be active

Process mining Align a sequence of entities to activity models, then use this relatedness to guide the repair [6]

- (EnergyLoss)=>(TimeoutAlert)=>(LossOfSignal)
- (LoginFail)=>(LoginFail)=>(LoginFail)

Statistical learning Relate entities based on context similarities, then use this relatedness to alert and guide the repair [1]

- The hidden cause of the trouble ticket on server 1 is a “data leak” attack that started on server 2

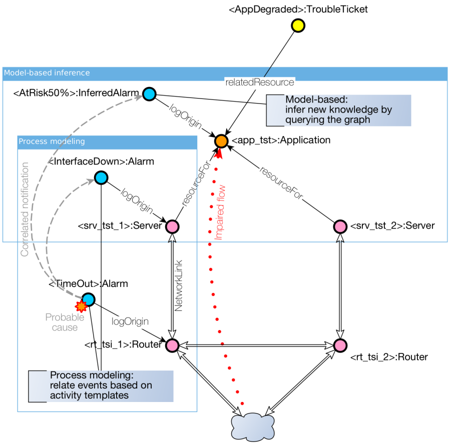

Model-based: infer new knowledge by querying the graph

Process modeling: relate events based on activity templates

Experimental design – Evaluation methodology and scenario

Framework Client-server software architecture + extended SUS survey + platform analytics + three step incident diagnosis scenario ...

Evaluation dataset Test data + production network data in a knowledge graph-based platform [5]

Model-based design Query the graph to retrieve anomalies and their context [1]

- k out-of n devices with faults
- User with unusual account rights
- Absence of traffic on an interface supposed to be active

Process mining Align a sequence of entities to activity models, then use this relatedness to guide the repair [6]

- (EnergyLoss)=>(TimeoutAlert)=>(LossOfSignal)
- (LoginFail)=>(LoginFail)=>(LoginFail)

Statistical learning Relate entities based on context similarities, then use this relatedness to alert and guide the repair [1]

- The hidden cause of the trouble ticket on server 1 is a “data leak” attack that started on server 2

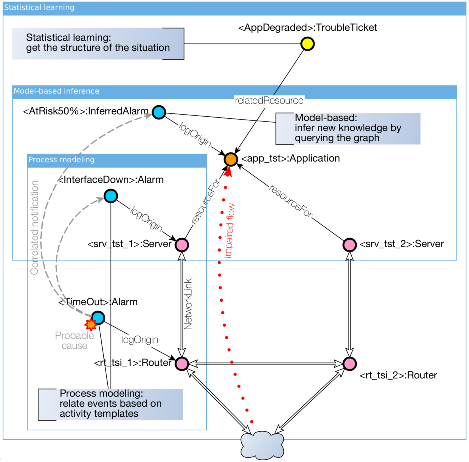

Experimental design – The NORIA UI (key UI/UX features)

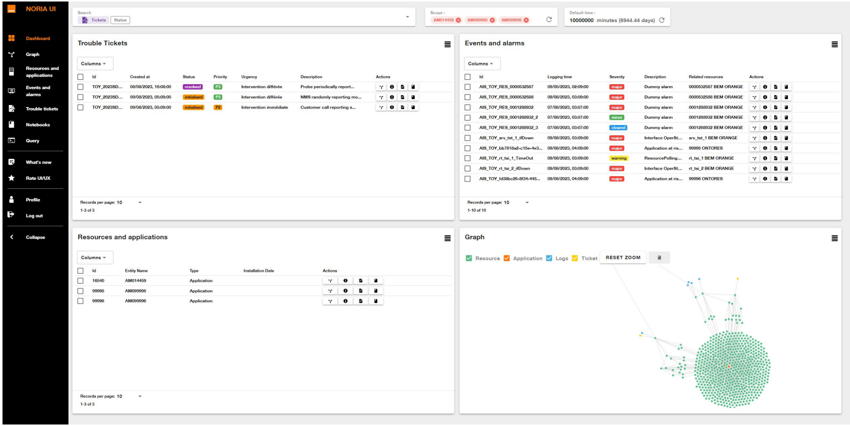

Search: Tickets Status

Scope : AM014459 AM099995 AM099996

Default time : 10000000 minutes (6944.44 days)

Trouble Tickets

	Id	Created at	Status	Priority	Urgency	Description	Actions
□	TOY_2023SD...	08/08/2023, 16:08:00	resolved	P3	Intervention différée	Probe periodically report...	
□	TOY_2023SD...	09/08/2023, 05:09:00	initialised	P3	Intervention différée	NMS randomly reporting mo...	
□	TOY_2023SD...	09/08/2023, 05:09:00	initialised	P2	Intervention immédiate	Customer call reporting s...	

Records per page: 10

1-3 of 3

Events and alarms

	Id	Logging time	Severity	Description	Related resources	Actions
□	AIS_TOY_RES_0000532587	09/05/2023, 02:09:00	major	Dummy alarm	0000532587 BEM ORANGE	
□	AIS_TOY_RES_0000532588	09/08/2023, 03:09:00	major	Dummy alarm	0000532588 BEM ORANGE	
□	AIS_TOY_RES_0001288932	07/08/2023, 03:07:00	major	Dummy alarm	0001288932 BEM ORANGE	
□	AIS_TOY_RES_0001288932_2	07/08/2023, 03:07:00	minor	Dummy alarm	0001288932 BEM ORANGE	
□	AIS_TOY_RES_0001288932_3	07/08/2023, 03:07:00	cleared	Dummy alarm	0001288932 BEM ORANGE	
□	AIS_TOY_srv_tst_1_ifDown	09/08/2023, 03:09:00	major	Interface OperSt...	srv_tst_1 BEM ORANGE	
□	AIS_TOY_bb7818a2-c15e-4e3...	09/08/2023, 04:09:00	major	Application at ris...	99995 ONTORES	
□	AIS_TOY_rt_tsi_1_TimeOut	09/08/2023, 03:09:00	warning	ResourcePolling...	rt_tsi_1 BEM ORANGE	
□	AIS_TOY_rt_tsi_2_ifDown	09/08/2023, 03:09:00	major	Interface OperSt...	rt_tsi_2 BEM ORANGE	
□	AIS_TOY_fd38bc26-8f34-445...	09/08/2023, 04:09:00	major	Application at ris...	99996 ONTORES	

Records per page: 10

1-10 of 10

Resources and applications

	Id	Entity Name	Type	Installation Date	Actions
□	16540	AM014459	Application		
□	99995	AM099995	Application		
□	99996	AM099996	Application		

Records per page: 10

1-3 of 3

Graph

Personas Incident manager, network supervision expert, cybersecurity analyst, system architect.

Principle Providing access to information about the network's life based on four complementary facets derived from the knowledge graph.

1 Cross-consultation of information on network topology, events and alarms,

2 Display of 2D/3D network topology enriched with indicators,

3 Aggregation and analysis of information in a dedicated digital investigation space,

4 Use of analysis and anomaly detection tools,

5 Access to community functions for sharing information between collaborators.

Evaluation – Notes and overall SUS scores by personas

		Q.1	Q.2	Q.3	Q.4	Q.5	Q.6	Q.7	Q.8	Q.9	Q.10	SUS
Persona	N	+	−	+	−	+	−	+	−	+	−	w.Σ
Cybersecurity analyst	2	**10.0**	**0.5**	7.5	4.0	**9.0**	2.0	**9**	2.0	8.5	2.5	78.8
Incident manager	2	**10.0**	**0.5**	**8.0**	8.0	8.5	9.0	7	2.5	**9.5**	2.5	63.1
Network supervision expert	1	**10.0**	2.0	**8.0**	**2.0**	8.0	1.0	8	**1.0**	8.0	**1.0**	**81.3**
System architect	3	7.3	6.7	6.0	4.3	8.0	**0.7**	8	2.7	8.3	4.7	60.8
Average (complete)	8	**9.0**	3.0	7.1	4.9	8.4	3.1	8	**2.3**	8.6	3.1	68.4
System architect (partial)	2	5.5	7.0	3.0	7.5	4.0	5.0	3	7.0	4.0	6.0	21.3
Average (all)	10	**8.3**	3.8	6.3	5.4	7.5	3.5	7	**3.2**	7.7	3.7	59.0

The Q.x columns provide the ratings for SUS questions on a scale of 1 to 10, with the +/− sign indicating whether it is a positive question (the higher the better) or a negative question (the lower the better). The SUS column is the overall SUS score calculated by weighted sum. The values by personas are separated between respondents who completed the test scenario fully and those who completed it partially. The values in bold highlight the highest scores. N stands for the number of respondents.

Q.1 I think that I would like to use this system frequently.

Q.2 I found the system unnecessarily complex.

Q.3 I thought the system was easy to use.

Q.4 I think that I would need the support of a technical person to be able to use the system.

Q.5 I found the various functions in this system were well integrated.

Q.6 I thought there was too much inconsistency in this system.

Q.7 I would imagine that most people would learn to use this system very quickly.

Q.8 I found the system very cumbersome to use.

Q.9 I felt very confident using the system.

Q.10 I needed to learn a lot of things before I could get going with this system.

Evaluation – Notes and overall SUS scores by personas (extended)

		Q.1	Q.2	Q.3	Q.4	Q.5	Q.6	Q.7	Q.8	Q.9	Q.10	SUS
Persona	N	+	−	+	−	+	−	+	−	+	−	w.Σ
Cybersecurity analyst	2	**10.0**	**0.5**	7.5	4.0	**9.0**	2.0	**9**	2.0	8.5	2.5	78.8
Incident manager	2	**10.0**	**0.5**	**8.0**	8.0	8.5	9.0	7	2.5	**9.5**	2.5	63.1
Network supervision expert	1	**10.0**	2.0	**8.0**	**2.0**	8.0	1.0	8	**1.0**	8.0	**1.0**	**81.3**
System architect	3	7.3	6.7	6.0	4.3	8.0	**0.7**	8	2.7	8.3	4.7	60.8
Average (complete)	8	**9.0**	3.0	7.1	4.9	8.4	3.1	8	**2.3**	8.6	3.1	68.4
System architect (partial)	2	5.5	7.0	3.0	7.5	4.0	5.0	3	7.0	4.0	6.0	21.3
Average (all)	10	**8.3**	3.8	6.3	5.4	7.5	3.5	7	**3.2**	7.7	3.7	59.0

Two complementary questions added to the SUS survey in order to refine the evaluation...

- “In a few words, what were you looking for the last time you used NORIA?”
  → dependency visualization; incident correlation; utilization for incident diagnosis; examples of ontology and automated processing.
- “Do you have any feedback, positive/negative opinions about NORIA UI, or any suggestions?”
  → overall relevance of the proposal; some improvements in the sequence of actions; evaluation in further usage contexts; solving minor ergonomic issues.

Evaluation – Notes and overall SUS scores by personas (analysis)

		Q.1	Q.2	Q.3	Q.4	Q.5	Q.6	Q.7	Q.8	Q.9	Q.10	SUS
Persona	N	+	−	+	−	+	−	+	−	+	−	w.Σ
Cybersecurity analyst	2	**10.0**	**0.5**	7.5	4.0	**9.0**	2.0	**9**	2.0	8.5	2.5	78.8
Incident manager	2	**10.0**	**0.5**	**8.0**	8.0	8.5	9.0	7	2.5	**9.5**	2.5	63.1
Network supervision expert	1	**10.0**	2.0	**8.0**	**2.0**	8.0	1.0	8	**1.0**	8.0	**1.0**	**81.3**
System architect	3	7.3	6.7	6.0	4.3	8.0	**0.7**	8	2.7	8.3	4.7	60.8
Average (complete)	8	**9.0**	3.0	7.1	4.9	8.4	3.1	8	**2.3**	8.6	3.1	68.4
System architect (partial)	2	5.5	7.0	3.0	7.5	4.0	5.0	3	7.0	4.0	6.0	21.3
Average (all)	10	**8.3**	3.8	6.3	5.4	7.5	3.5	7	**3.2**	7.7	3.7	59.0

→ **Correlation between the respondents' profiles and their evaluations**, with acceptability level ranging from good (Incident manager, System architect) to high (Network supervision expert, Cybersecurity analyst) for beta testers who completed the test scenario fully, and not acceptable for those who tested it partially (System architect $\times 2$).

→ **Data quality bias in the evaluation scores**, as respondents rated the proposed solution based on the available data (synthetic data + partially consolidated operational data from [5]) and unresolved display bugs at the time of evaluation.

→ **Multi-faceted & homogenized view as beneficial**, notably considering the integration of diagnostic aid tools in the workflow. However additional synthesis/recommendation/interaction capabilities are asked for depending on the user profile, notably of personas (Incident manager and System architect) whose role typically involves using analysis rather than producing it.

→ **Correct UI responsiveness**, as loading and display delay caused by interactions with the graph database and visual rendering in the graph component seem acceptable as it is. This should be monitored under more intensive usage conditions.

Summary & future work

Problem Smooth integration of knowledge graphs in NMS/SIEM solutions towards advanced anomaly detection and decision support in complex ICT systems.

Approach Co-design of a KG-based client-server software architecture, synergistic reasoning, interaction mechanisms for exploratory analysis of multi-layered systems.

Insights Value in providing access to four complementary facets of the network on top of heterogeneous data, value in synergistic reasoning and interaction mechanisms, specific expectations of NetOps/SecOps user profiles in terms of functions and usage actions.

Next Analysis functions (graphical root cause analysis, incident context similarity), XAI features (human-in-the-loop for statistical learning & data quality, reporting on the inference process), platform performance (offloading computation to the backend, improving the data model).

Paper – GRASEC @ ARES 2024

Lionel TAILHARDAT, Yoan CHABOT, Antoine PY, and Perrine GUILLEMETTE.

NORIA UI: Efficient Incident Management on Large-Scale ICT Systems Represented as Knowledge Graphs.

10.1145/3664476.3670438

https://hellofuture.orange.com/en/
noria-network-anomaly-detection-using-knowledge-graphs/

Additional materials

The NORIA UI Web application framework & architecture

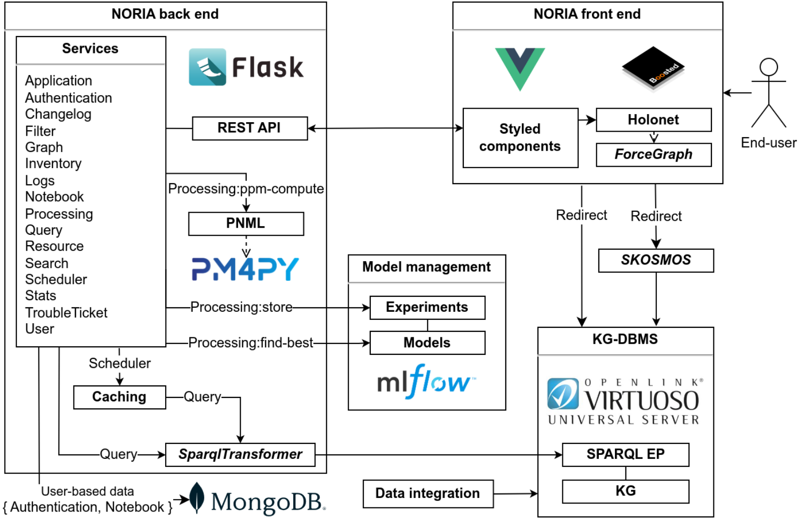

Experimental design – The NORIA UI (dashboard page)

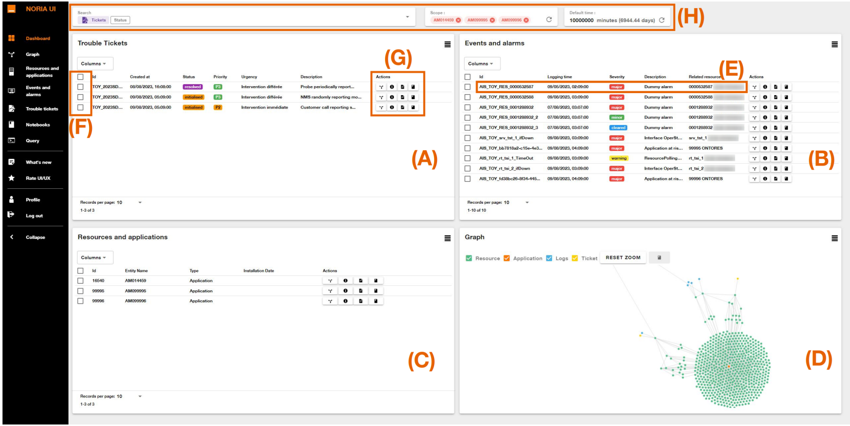

Trouble Tickets

□	Id	Created at	Status	Priority	Urgency	Description	Actions
□	TOY_2023SD...	08/08/2023, 16:08:00	resolved	P3	Intervention différée	Probe periodically report...	
□	TOY_2023SD...	09/08/2023, 05:09:00	initialised	P3	Intervention différée	NMS randomly reporting mo...	
□	TOY_2023SD...	09/08/2023, 05:09:00	initialised	P2	Intervention immédiate	Customer call reporting s...	

Records per page: 10
1-3 of 3

Events and alarms

□	Id	Logging time	Severity	Description	Related resource	Actions
□	AIS_TOY_RES_0000532587	09/05/2023, 02:09:00	major	Dummy alarm	0000532587 [redacted]	
□	AIS_TOY_RES_0000532588	09/08/2023, 03:09:00	major	Dummy alarm	0000532588 [redacted]	
□	AIS_TOY_RES_0001288932	07/08/2023, 03:07:00	major	Dummy alarm	0001288932 [redacted]	
□	AIS_TOY_RES_0001288932_2	07/08/2023, 03:07:00	minor	Dummy alarm	0001288932 [redacted]	
□	AIS_TOY_RES_0001288932_3	07/08/2023, 03:07:00	cleared	Dummy alarm	0001288932 [redacted]	
□	AIS_TOY_srv_tst_1_ifDown	09/08/2023, 03:09:00	major	Interface OperSt...	srv_tst_1 [redacted]	
□	AIS_TOY_bb7818a2-c15e-4e3...	09/08/2023, 04:09:00	major	Application at ris...	99995 ONTORES	
□	AIS_TOY_rt_tsi_1_TimeOut	09/08/2023, 03:09:00	warning	ResourcePolling...	rt_tsi_1 [redacted]	
□	AIS_TOY_rt_tsi_2_ifDown	09/08/2023, 03:09:00	major	Interface OperSt...	rt_tsi_2 [redacted]	
□	AIS_TOY_fd38bc26-8f34-445...	09/08/2023, 04:09:00	major	Application at ris...	99996 ONTORES	

Records per page: 10
1-10 of 10

Resources and applications

□	Id	Entity Name	Type	Installation Date	Actions
□	16540	AM014459	Application		
□	99995	AM099995	Application		
□	99996	AM099996	Application		

Records per page: 10
1-3 of 3

Graph

A Trouble tickets (querying for `noria:TroubleTicket` entities).

B Events and alarms (querying for `noria:EventRecord` entities).

C Resources and applications (querying for `noria:Resource` and `noria:Application` entities).

D The enriched network topology (specific API call + embedded graph viz component).

E Entities are displayed along with their main properties.

F Checkboxes allow for a display pivot that applies to all panels.

G Each entity has contextual commands.

H An input field allows defining a display scope and searching for entities based on their properties.

Diagnosing with NORIA UI – Graphical root cause analysis

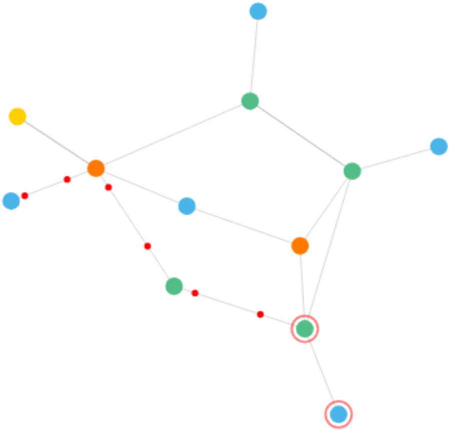

A prototype of the graphical root cause analysis view obtained by projecting the procedural model from the process mining step onto the entities in the NORIA UI notebook. The circled nodes highlight the `noria:Resource` and the `noria:EventRecord` likely responsible for the incident. The dotted lines emphasize the temporal sequence.

Evaluation – Verbatims from the extended SUS survey (1)

“In a few words, what were you looking for the last time you used NORIA?”

Dependency Visualization

Users want to search for equipments and understand their connections; examples or technical diagrams would be helpful during the initial use of the tool to ensure a clear understanding of the displayed graphs; users also want to understand which applications are installed on which servers.

Incident Correlation

Users seek assistance in correlating incidents and events; they suggest that an AI could automate the correlation of events related to a root cause and automatically provide suggestions in the notebook.

Utilization for Incident Diagnosis

Users are interested in exploring how NORIA can assist in incident diagnosis, particularly for transmission supervisors; they want to envision how the tool would be used by supervisors/operators; users note that if the context is saved, it could accelerate the automation of diagnostics.

Ontology and Automated Processing

Users are looking for an illustration of the NORIA-O data model to utilize automatic processing algorithms for detection, diagnosis, and remediation proposals.

Evaluation – Verbatims from the extended SUS survey (2)

“Do you have any feedback, positive/negative opinions about NORIA UI, or any suggestions?”

General feedback The respondents provided feedback on ...

- **Overall relevance of the proposal**, e.g.: “the tool could be very useful for ICT systems supervision to quickly identify the root cause of an incident, calculate incident impact, and analyze incidents retrospectively”,
- **Some improvements in the sequence of actions**, e.g.: “the concept of a notebook to pin relevant elements is interesting, but the manipulations are somewhat tedious” and “the tool appears to be designed as a navigation tool for domain experts, requiring many clicks and not suitable for real-time incident handling”.

Feature requests The respondents requested ...

- **Evaluation in further usage contexts**, e.g.: “a more realistic test scenario would have been helpful to fully grasp the interface and data”,
- **Solving minor ergonomic issues**, e.g.: changes in colors and missing tooltips for increased readability, page and parameter display/refresh problems on a specific Web browser.

Evaluation – Performance

Period 2024-02-10 to 2024-03-10.

Platform Two RHEL7.6 120 GB disk, 16 GB RAM virtual machines on Ericsson HDS 8000 - Intel Xeon DP 2.8 GHz hosts.

Visits 134 visits from authenticated users.

Visit durations Average time on the website of 16 minutes per visitor.

Visit actions Average of 14.3 actions per visitor.

Devices All visitors accessed the UI from a desktop or laptop workstation, with a screen resolution of 2560×1440 for 36% of them, 1920×1080 for 20%, and lower for the rest.

Browsers The UI navigation was done through standard Web browsers, with Chrome accounting for 44%, Microsoft Edge for 29%, and Firefox for 27%.

Page loading The average page loading time observed was 2.89 seconds (2.03 seconds for generating the DOM), with the following average times per specific page (in decreasing order of the top four):

- Graph $= 4.62$ seconds,
- Home page/Dashboard $= 2.97$ seconds,
- Inventory $= 2.79$ seconds,
- Notebook $= 1.94$ seconds.

NORIA-O competency questions driving the NORIA UI design stage

The 26 NORIA-O competency questions, available at https://w3id.org/noria/cqs/, collected during knowledge capture meetings [2].

#	Competency questions
1	Which resource/application/site is concerned by a given incident?
2	What assets are shared by a given asset chain?
3	What logs and alarms are coming from a specified resource?
4	Which metrics are coming from a specified resource?
5	To which event family does this log belong and is this event normal or abnormal?
6	What events are associated with a given event?
7	Which agent/event/resource caused the event under analysis?
8	What do the various fields in the log refer to?
9	Is there any pattern in a given set of logs/alarms?
10	What interventions were carried out on this resource that could have caused the incident?
11	What was the root cause of the incident?
12	Which sequence of events led to the incident?
13	On which resource did this sequence of events take place and in which order?
14	What past incidents are similar to a given incident?
15	What operation plan (automation, operating procedures, etc.) could help us solve the incident?
16	What corrective actions have been carried out so far for a given incident?
17	What is the list of actions taken that led to the resolution of the incident?
18	Given all the corrective actions carried out so far for the incident, what assumptions covered the actions taken?
19	What has been the effect of the corrective actions taken so far for the incident?
20	Given all the corrective actions carried out so far for the incident, what possible actions could we still take?
21	What is the summary of this incident and its resolution?
22	Which agents were involved in the resolution of the incident?
23	What is the financial cost of this incident if it occurs?
24	How long before this incident is resolved?
25	What are the vulnerabilities and the associated risk levels of this infrastructure?
26	What is the most likely sequence of actions that would cause this infrastructure to fail?

Incident diagnosis activity cases driving the NORIA UI design stage

List of use cases from expert panel interviews [1], in simplified form.

#	Description
1	Circumscribe assets and causes search space for multi-applications incident situations
2	Alert on impaired service situations occurring on (distributed) fail-over architectures
3	Assess legitimacy of a given network flow
4	Track single identity from a set of various activity traces
5	Analyze false-positive and recurrent cyber security alerts
6	Analyze compliance of web navigation traces from institutional website

Bibliographical references

[1] Tailhardat, et al. 2023. “Leveraging Knowledge Graphs For Classifying Incident Situations in ICT Systems” (ARES’2023)

[2] Tailhardat, et al. 2024. “NORIA-O: An Ontology for Anomaly Detection and Incident Management in ICT Systems” (ESWC’2024)

[3] Goertzel, et al. 2014. “Engineering General Intelligence, Part 1: A Path to Advanced AGI via Embodied Learning and Cognitive Synergy” (Atlantis Press)

[4] John Brooke, 1995. “SUS: A Quick and Dirty Usability Scale” (Usability Eval. Ind.)

[5] Tailhardat, et al. 2023. “Designing NORIA: a Knowledge Graph-based Platform for Anomaly Detection and Incident Management in ICT Systems” (ESWC’2023)

[6] Tailhardat, et al. 2024. “Graphameleon: Relational Learning and Anomaly Detection on Web Navigation Traces Captured as Knowledge Graphs” (WWW’2024)