Poster: The FAIRer Path Forward: Advancing Research Infrastructures for Unified Reproducibility and Usability

Eric Hauser
Technical University of Munich
Garching near Munich, Germany
hauser@net.in.tum.de

Kilian Warmuth
Technical University of Munich
Garching near Munich, Germany
warmuth@net.in.tum.de

Johannes Späth
Technical University of Munich
Garching near Munich, Germany
spaethj@net.in.tum.de

Georg Carle
Technical University of Munich
Garching near Munich, Germany
carle@net.in.tum.de

Abstract

Artifact evaluation places a significant burden on both authors and reviewers: Authors must package bespoke experimental setups, while committees struggle to validate them. Research Infrastructures (RIs) offer tools to support reproducibility, yet interoperability remains limited due to insufficient standardization efforts.

To address this, we propose new standards that unify RIs through a formal experiment execution format and topology management for networked experiments. In addition, we advance sustainability through integrated energy analysis and modeling. Together, these contributions improve reproducibility, enable cross-RI collaboration, and foster more sustainable experiments. This makes RIs more FAIR—particularly by strengthening interoperability and reusability—while aiding broader adoption across the (network) research community.

CCS Concepts

• **General and reference → Experimentation**; **Measurement**; *Validation*; • **Networks** → *Network experimentation*; *Topology analysis and generation*; *Physical topologies*; • **Social and professional topics** → *Sustainability*.

Keywords

Research Infrastructures; Testbeds; Experiments; Measurements; Reproducibility; RO-Crate; YANG; Network Topology; Topology Management; Energy Monitoring; Sustainability

ACM Reference Format:
Eric Hauser, Kilian Warmuth, Johannes Späth, and Georg Carle. 2025. Poster: The FAIRer Path Forward: Advancing Research Infrastructures for Unified Reproducibility and Usability. In *Proceedings of the 21st International Conference on emerging Networking EXperiments and Technologies (CoNEXT '25), December 1–4, 2025, Hong Kong, Hong Kong*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3765515.3771744

ACM ISBN 979-8-4007-2191-5/2025/12
https://doi.org/10.1145/3765515.3771744

1 Introduction

Research Infrastructures (RIs), i.e., testbeds, provide an essential service to the research community by enabling the design, execution, and evaluation of complex experiments under controlled conditions. A core strength of RIs lies in their emphasis on reproducibility, allowing researchers to validate existing results and build upon prior work. However, despite significant advances in recent years, there remain important challenges to foster usability, sustainability, interoperability, and reproducibility across RIs.

In this work, we present three novel approaches to extend the capabilities of existing RIs:

(1) Standardizing a structure for experiment executions using newly introduced RO-Crate profiles.
(2) Introducing network topology management based on YANG, an IETF data modeling language.
(3) Integrating energy consumption monitoring and analysis into the experimental workflow.

These contributions are motivated by two main goals: First, to propose unified methods addressing a lack of open, widely adopted RI standards, improving convenience and efficiency to encourage broader RI adoption. Second, to strengthen reproducibility, supporting the community in producing more transparent and comparable results for easier evaluation. *This work raises no ethical issues.*

2 Background

RIs such as CloudLab [4], Chameleon [9], and SLICES [5] provide environments where the FAIR principles [12]—Findable, Accessible, Interoperable, and Reusable—can be operationalized. They also enable the automation of these processes at scale. Despite this, achieving reproducibility across different RIs remains challenging due to heterogeneous execution tools and inconsistent measurement practices. In our previous work [8], we demonstrated which elements experiment artifacts should include to improve reproduction across RIs, but interoperability between them remains challenging. To address these challenges, standardized formats are needed to simplify and enhance information exchange between RIs.

3 Implementation

In this section, we outline our contributions. Due to the limited space of this extended abstract, we present a concise overview of our approaches rather than full technical details.

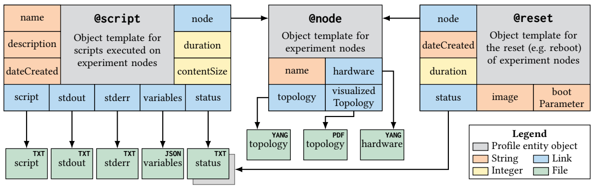

Figure 1: Simplified draft of an RO-Crate profile for RIs.

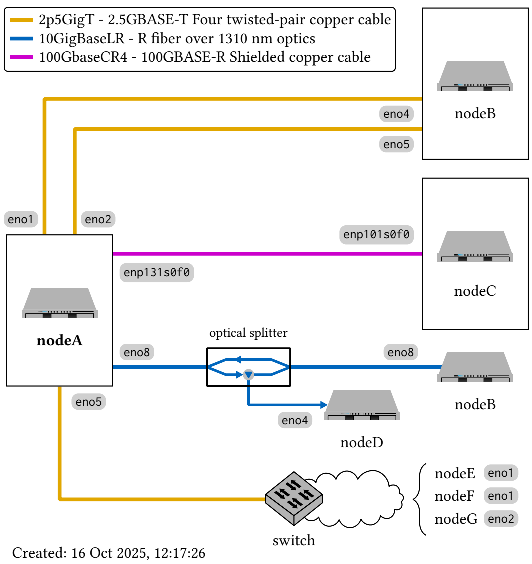

Figure 2: Visualized experiment topology of example nodeA.

3.1 Standardized Experiment Structure with RO-Crate Profiles

RO-Crate [11] is a framework for packaging research data together with its metadata. It builds on JSON for Linked Data (JSON-LD), which enables linking result resources and fosters interoperability. In previous work, we demonstrated how RO-Crate can be used to package artifacts of RIs [7]. RO-Crate encourages using terms from *schema.org* [6] to describe (meta-)data. However, this open vocabulary is often too generic and does not adequately cover the specific workflows of RIs.

The recently released RO-Crate version 1.2 (June 2025) [10] introduces the concept of profiles, enabling domain-specific terminologies tailored to particular use cases. Building on this, we propose a dedicated RO-Crate profile for RIs to more accurately capture their experimental workflows. The profile defines a basic set of unified operations common across RIs and supports the description of key processes such as experiment node resets or script executions. Figure 1 illustrates a simplified draft for these exemplary operations. The draft objects include properties like name or duration, and can link to adjacent files or other objects.

3.2 Network Topology Management with YANG

Experiments in distributed systems and network protocol prototyping often rely on multi-node setups where the network topology

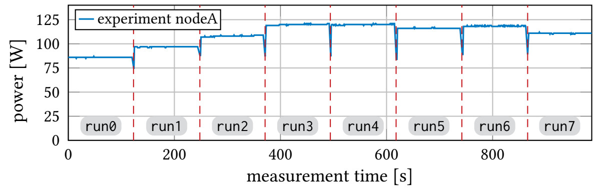

Figure 3: Experiment power consumption over time.

becomes a crucial part of it. Yet, topologies are usually documented only graphically, limiting automated processing and requiring manual work. In multi-user testbeds, frequent changes further underline the need for automated detection, representation, visualization, modification, and reuse—features we summarize under the term *topology management*.

In this work, focusing on the core aspect of representation, we base topology management on the Yet Another Next Generation (YANG) data modeling language. YANG was originally developed for the NETCONF protocol and standardized in RFCs 6020 and 7950 [1, 2]. It allows defining structured data templates, called models, that can be reused and extended through augmentations.

We build on the IETF L2 Topology Model (RFC 8944 [3]), augmenting it to meet the needs of RI-based experiment topologies. Specifically, this includes support for bidirectional links as well as adding link characteristics such as medium type (e.g., fiber) and line rate. In addition, we add support for fiber-optic splitters (i.e., passive optical taps) with their respective split ratios. Figure 2 illustrates an example topology of an experimental node, demonstrating key features supported by our YANG model, including bidirectional links, a splitter configuration, and a switched network segment.

3.3 Energy Monitoring and Modeling

Power consumption and energy efficiency are important, but often overlooked in experimental research. To address this, we integrate automated monitoring of energy-related parameters via metered power outlets, with results directly linked to individual experiment runs through RO-Crate. Figure 3 shows the automated power analysis of an experiment divided into eight runs, each applying a different load and thus consuming different amounts of energy. Beyond analysis, we provide modeling tools in Jupyter Notebooks that help researchers identify phases of high energy consumption, optimize their implementations, and predict the estimated power consumption of different configurations.

4 Conclusion

This work showcased how targeted advancements in standardized experiment execution, topology management, and energy monitoring can make RIs more reproducible, transparent, and convenient to use. Beyond improving experimental workflows, these approaches pave the way for more sustainable and comparable research practices across different testbeds. To test the proposed concepts, we invite the community to seek collaboration with our RI, SLICES-DE (https://slices-de.org/), Germany's contribution to the European SLICES initiative. Looking ahead, we aim to drive the integration of these approaches into future standards and best practices, shaping the next generation of reproducible and interoperable RIs.

Acknowledgments

We received funding from the EU's Horizon 2020 programme as part of the projects GreenDIGIT (101131207) and SLICES-PP (10107977). Moreover, we received funding from the Bavarian Ministry of Economic Affairs, Regional Development and Energy within the project 6G Future Lab Bavaria. Additional funding was received from the Federal Ministry of Research, Technology and Space (BMFTR) under the project 6G-life (16KISK002) as well as from the German Research Foundation (DFG) within the project HyperNIC (CA595/13-1).

References

[1] Martin Björklund. 2010. *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF).* RFC 6020. RFC Editor. http://www.rfc-editor.org/rfc/rfc6020.txt

[2] Martin Björklund. 2016. *The YANG 1.1 Data Modeling Language.* RFC 7950. RFC Editor. http://www.rfc-editor.org/rfc/rfc7950.txt

[3] Jie Dong, Xiugang Wei, Qin Wu, Mohamed Boucadaira, and Anders Liu. 2020. *A YANG Data Model for Layer 2 Network Topologies.* RFC 8944. RFC Editor. http://www.rfc-editor.org/rfc/rfc8944.txt

[4] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, et al. 2019. The Design and Operation of CloudLab. In *2019 USENIX Annual Technical Conference, USENIX ATC 2019, Renton, WA, USA, July 10-12, 2019.* https://www.usenix.org/conference/atc19/presentation/duplyakin

[5] Serge Fdida, Nikos Makris, Thanasis Korakis, Raffaele Bruno, Andrea Passarella, Panayiotis Andreou, Bartosz Belter, Cedric Crettaz, Walid Dabbous, Yuri Demchenko, and Raymond Knopp. 2022. SLICES, a scientific instrument for the networking community. *Comput. Commun.* 193 (2022), 189–203. doi:10.1016/J.COMCOM.2022.07.019

[6] Ramanathan V. Guha, Dan Brickley, and Steve Macbeth. 2016. Schema.org: evolution of structured data on the web. *Commun. ACM* 59, 2 (2016), 44–51. doi:10.1145/2844544

[7] Eric Hauser, Sebastian Gallenmüller, and Georg Carle. 2024. RO-Crate for Testbeds: Automated Packaging of Experimental Results. In *IFIP Networking Conference, IFIP Networking 2024, Thessaloniki, Greece, June 3-6, 2024.* IEEE, 654–659. doi:10.23919/IFIPNETWORKING62109.2024.10619057

[8] Eric Hauser, Kilian Warmuth, Marcel Kempf, Stefan Lachnit, and Georg Carle. 2025. MARTE: Malleable and Automated Reproduction for Testbed-driven Experiments. In *ACM Conference on Reproducibility and Replicability (ACM REP '25).* Vancouver, BC, Canada. doi:10.1145/3736731.3746154

[9] Kate Keahey, Jason Anderson, Zhuo Zhen, et al. 2020. Lessons Learned from the Chameleon Testbed. In *2020 USENIX Annual Technical Conference, USENIX ATC 2020, July 15-17, 2020.* USENIX Association. https://www.usenix.org/conference/atc20/presentation/keahey

[10] Peter Sefton, Eoghan Ó Carragáin, Stian Soiland-Reyes, Oscar Corcho, Daniel Garijo, Raul Palma, Frederik Coppens, Carole Goble, José M. Fernández, Kyle Chard, Jose Manuel Gomez-Perez, Michael R. Crusoe, Ignacio Eguinoa, Nick Juty, Kristi Holmes, Jason A. Clark, Salvador Capella-Gutierrez, Alasdair J. G. Gray, Stuart Owen, Alan R. Williams, Giacomo Tartari, Finn Bacall, Thomas Thelen, Hervé Ménager, Laura Rodríguez-Navas, Paul Walk, brandon whitehead, Mark Wilkinson, Paul Groth, Erich Bremer, Leyla Jael Castro, Karl Sebby, Alexander Kanitz, Ana Trisovic, Gavin Kennedy, Mark Graves, Jasper Koehorst, Simone Leo, Marc Portier, Paul Brack, Milan Ojsteršek, Bert Droesbeke, Chenxu Niu, Kosuke Tanabe, Tomasz Miksa, Marco La Rosa, Cedric Decruw, Andreas Czerniak, Jeremy Jay, Sergio Serra, Ronald Siebes, Shaun de Witt, Shady El Damaty, Douglas Lowe, Xuanqi Li, Sveinung Gundersen, Muhammad Radifar, Rudolf Wittner, Oliver Woolland, Paul De Geest, Douglas Fils, Florian Wetzels, Raül Sirvent, Abigail Miller, Jake Emerson, Davide Fucci, Bruno P. Kinoshita, Maciek Bąk, Jens Hollunder, Martin Weise, Vartika Bisht, Toshiyuki Nishiyama Hiraki, Bram Ulrichts, Michael Falk, Eli Chadwick, Daniel Bauer, James Love, Eleni Adamidi, Josh Moore, Lars Schöbitz, Andreas Meier, Juan Fuentes, Edan Bainglass, and Balázs E. Pataki. 2025. RO-Crate Metadata Specification 1.2.0. doi:10.5281/zenodo.13751027

[11] Stian Soiland-Reyes, Peter Sefton, Mercè Crosas, et al. 2022. Packaging research artefacts with RO-Crate. *Data Sci.* 5, 2 (2022). doi:10.3233/DS-210053

[12] Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data* 3, 1 (2016).