




Weaving YANG into a Network Data Fabric

IETF 116, NMRG



YANG and the Data Fabric Paradigm

- Advanced monitoring requires combining YANG data
 - Scattered YANG data sources, both at device and service levels of abstraction
 - Multi-vendor networks, with different YANG data models and telemetry protocols
 - Need for data infrastructure to collect, integrate, and expose all these data from the network
 - Other data sources (think cloud-edge and NFV and...)
- Data Fabric architecture provides a unified view of integrated data
 - Graphs to abstract consumers from the underlying complexities of the data source
 - Standard, secure interfaces for interacting with data



The goal → Implement a data fabric solution integrating YANG monitoring data

Any Standards for Data Fabric? Enter NGSI-LD

ETSI CIM ISG defines the NGSI-LD protocol, which is composed of two parts:


NGSI-LD Information Model

- Based on the labelled property graph (LPG) model
- Semantic annotations based on Semantic Web standards (RDF, OWL)
- Serialized using JSON-LD

NGSI-LD API

- REST-based API
- Context information management
- Queries & subscriptions
- Temporal evolution
- Distributed & federated architectures

Mapping YANG Data Nodes on the NGSI-LD Metamodel



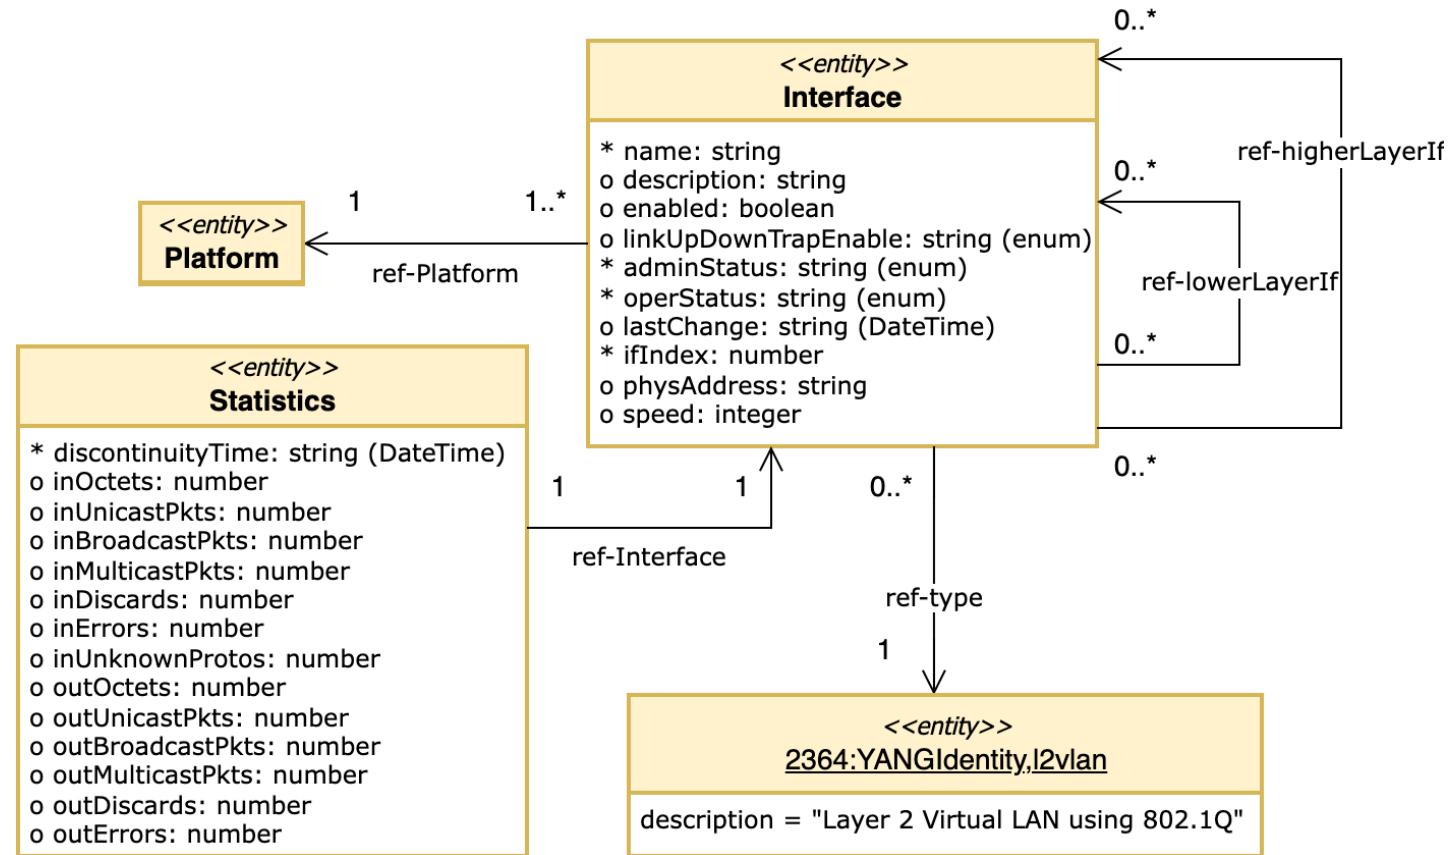
Starting with a Simple Use Case: Interfaces

```

1 module: ietf-interfaces
2   +--rw interfaces
3   |   +--rw interface* [name]
4   |   |   +--rw name string
5   |   |   +--rw description? string
6   |   |   +--rw type identityref
7   |   |   +--rw enabled? boolean
8   |   |   +--rw link-up-down-trap-enable? enumeration {if-mib}?
9   |   |   +--ro admin-status enumeration {if-mib}?
10  |   |   +--ro oper-status enumeration
11  |   |   +--ro last-change? yang:date-and-time
12  |   |   +--ro if-index int32 {if-mib}?
13  |   |   +--ro phys-address? yang:phys-address
14  |   |   +--ro higher-layer-if* interface-ref
15  |   |   +--ro lower-layer-if* interface-ref
16  |   |   +--ro speed? yang:gauge64
17  |   |   +--ro statistics
18  |   |   |   +--ro discontinuity-time yang:date-and-time
19  |   |   |   +--ro in-octets? yang:counter64
20  |   |   |   +--ro in-unicast-pkts? yang:counter64
21  |   |   |   +--ro in-broadcast-pkts? yang:counter64
22  |   |   |   +--ro in-multicast-pkts? yang:counter64
23  |   |   |   +--ro in-discards? yang:counter32
24  |   |   |   +--ro in-errors? yang:counter32
25  |   |   |   +--ro in-unknown-protos? yang:counter32
26  |   |   |   +--ro out-octets? yang:counter64
27  |   |   |   +--ro out-unicast-pkts? yang:counter64
28  |   |   |   +--ro out-broadcast-pkts? yang:counter64
29  |   |   |   +--ro out-multicast-pkts? yang:counter64
30  |   |   |   +--ro out-discards? yang:counter32
31  |   |   |   +--ro out-errors? yang:counter32

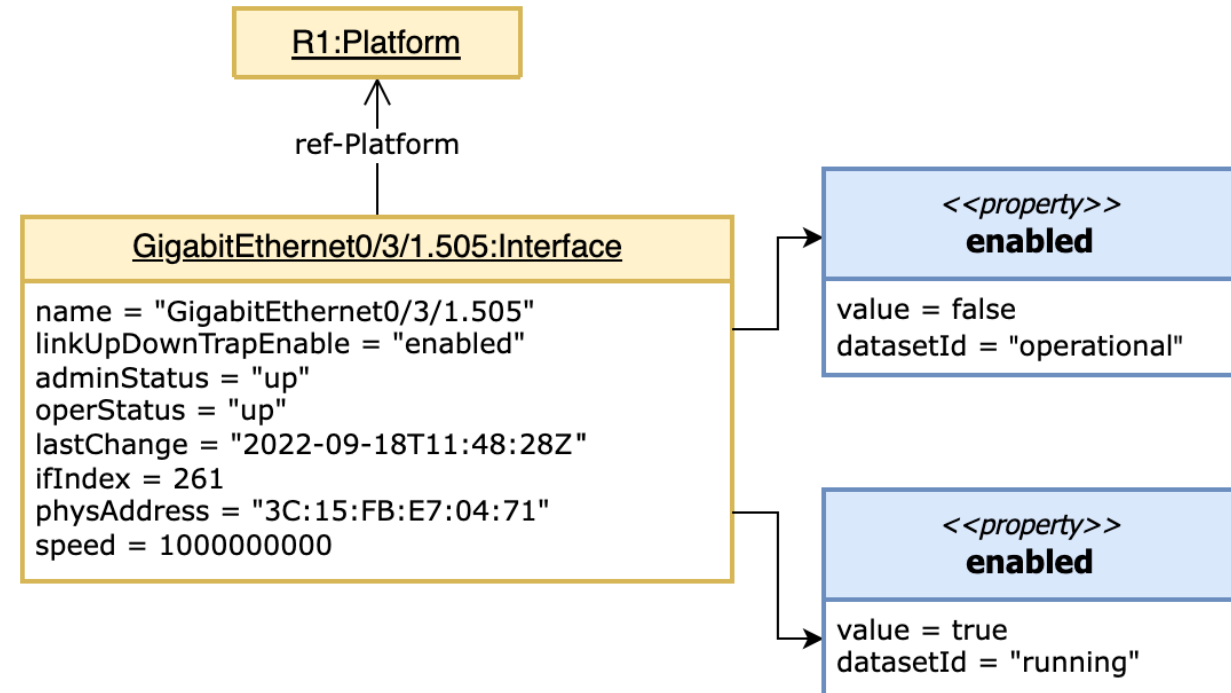
```

Tree representation of ietf-interfaces@2018-02-20.yang



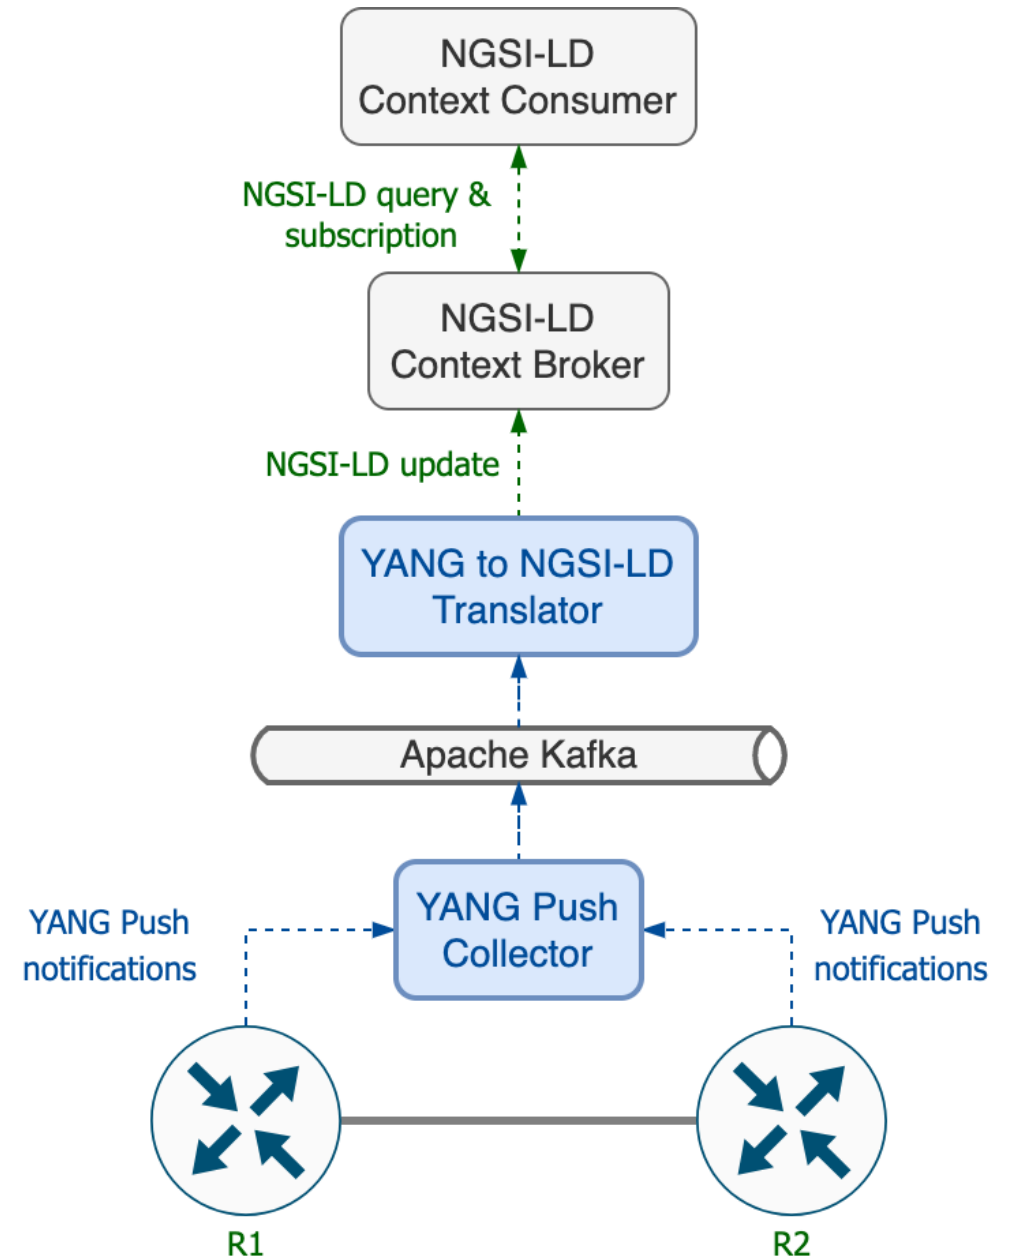
Managing YANG Configuration vs Operational Data

- NGSI-LD defines the **datasetId** feature:
 - Entity attributes can have multiple values provided by different data sources
 - Represented as a sub-property that uniquely identifies the source
- Feature allows for creating graph projections using the same data model
 - Avoid duplication of data
- **Proposal** → Leverage datasetId feature to identify YANG datastores (pre/post-NMDA)
 - Enabler for representing digital twins in the graph: <https://datatracker.ietf.org/doc/draft-irtf-nmrg-network-digital-twin-arch/>




Data Fabric Prototype

- [Containerlab](#) for network virtualization
 - Based on Docker containers
- Testing devices that implement YANG Push and IETF modules (NMDA ideally)
 - Cisco CSR1000v
 - Looking to include more vendors
- Two custom components
 - YANG Push collector, leveraging the [ncclient](#) library
 - YANG to NGSI-LD translator



A Few Research Challenges to Start With

- Address the whole YANG scope
 - Refine mapping rules with more complex YANG data models
 - Connect YANG data from device and service levels → The network map approach
- Mapping rules focused on NGSI-LD information model
 - Could be the basis for mappings to LPG (e.g., Neo4j) and RDF (generation of OWL ontologies from YANG data models)
- Align prototype development with YANG native onboarding
 - <https://github.com/network-analytics/draft-daisy-kafka-yang-integration/blob/main/draft-daisy-kafka-yang-integration-03.md>
 - Data fabric as an enabler of data mesh approach
- Collaboration with ETSI CIM ISG
 - Details of mapping rules in new group report
- Going beyond monitoring data
 - Integration of configuration data → The other half of the closed-loop



Telefónica