

# Nova Context

## LEVERAGING TOPOLOGY TO CONTEXTUALIZE NETWORK DATA



Unmatched graph-data and semantic modelling that correlates data in operations, business and service assurance systems to increase operational efficiency, avoid outages and reduce downtime by enabling automation.

### KEY FEATURES AND BENEFITS

**Correlated and centralized network, service and customer topology.** Nova Context (previously known as EXFO Ontology) builds accurate and dynamic topologies that connect customers to dependent infrastructure across silos and technologies. Data from any and all potential sources—including EMS, NMS, inventory, CRM, spreadsheets, device configurations—can be continuously reconciled into a comprehensive end-to-end model.

**Faster root cause analysis troubleshooting.** Accelerate root-cause investigations: identify infrastructure shared by multiple resources reporting similar issues to reduce NOC and SOC overheads, ultimately improving customer satisfaction.

**Identifying impact and fall-out of outages or degradations in real time.** Automatically determine the customer-facing impact of network faults. Enable the effective prioritization of repair activities according to business priorities and accurate customer notifications to better manage SLAs and avoid penalties. This visibility is key in avoiding churn.

**Reducing errors related to planned network changes.** Prevent unforeseen outages and cut the time in change approval by detecting change plans where impacts overlap in scope and time. Avoid outages and allow change management optimization, accelerate network management rollout and reduce repeat truck rolls.

**Dynamically linked silo networks for accurately automated provisioning.** Rapid alignment of data between platforms supports automated provisioning flows through APIs. OPEX reduction results from automated service order management and plan-to-build processes.

**Data quality management with consolidated blue print of network to achieve automation.** Flexible modelling architecture enables quick reconciliation between network and inventory, source of truth can be used to compare with additional connected sources. As a result, OPEX savings can be achieved through automated reconciliation processes.

## SEMANTIC GRAPH TECHNOLOGY: HOW DOES IT WORK?

Nova Context uses semantic and graph technology to link together topology across silos and facilitates the linking of different inventory and technology domains. It creates an end-to-end dependency graph that explains all relationships and dependencies that a service or subscriber are subject to across an entire infrastructure.

Nova Context pulls in data from multiple systems without the need for schema mapping and traditional ETL. This saves considerable effort and breaks the fragile close-coupling that complicates traditional integrations. Nova Context then represents that data in a graph database and uses an inference engine powered by semantic rules to classify entities and determine how the different entities depend upon one another.

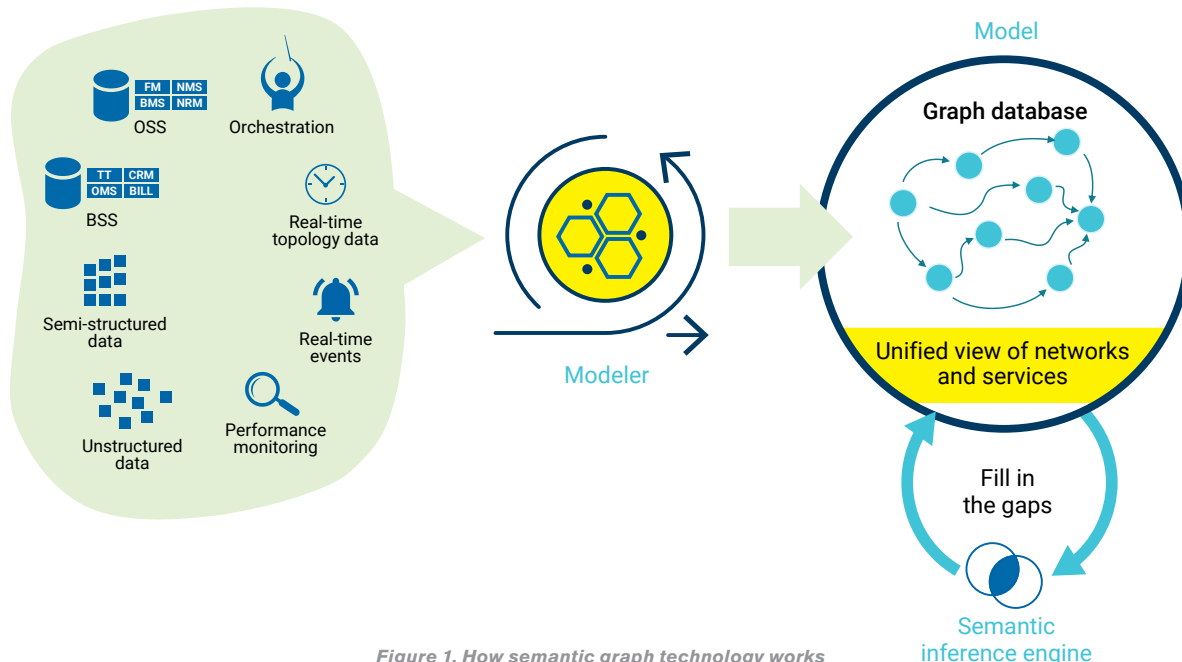


Figure 1. How semantic graph technology works

With the visibility that is achieved, it is possible address critical challenges in a number of use cases for improved network operations.

## AUTOMATE THE IDENTIFICATION OF SERVICE DEGRADATIONS AND ROOT CAUSE TROUBLESHOOTING ACROSS MULTIDOMAIN NETWORKS

Nova Context's automated common cause analysis (CCA) module automates the detection and troubleshooting of distributed service degradations impacting individual fixed customers and mobile cells which are caused by shared infrastructure deep in the network.

It uses pre-existing assurance data sources to identify network entities that are reporting service failure or degradation (*the symptom set*).

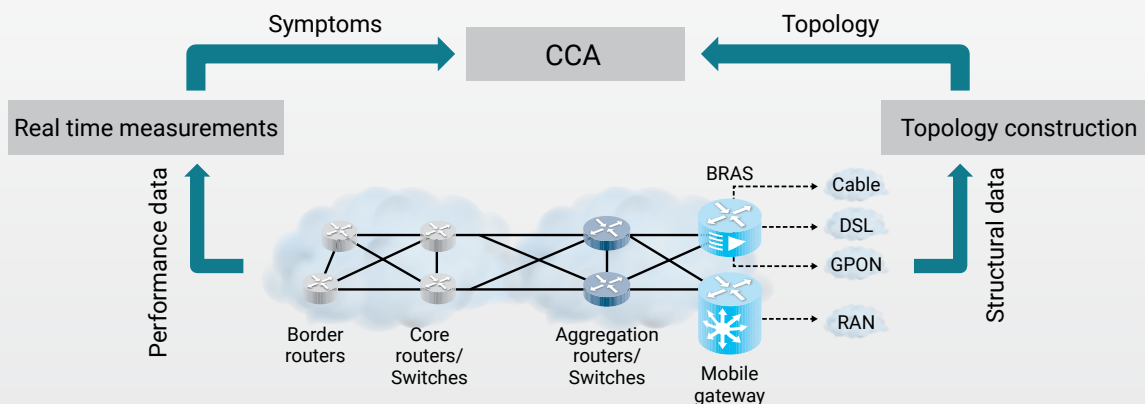


Figure 2. Automated common cause analysis

CCA then uses the network and service topology graph as a dependency function that finds the *potential cause* set of entities that can be associated with the symptoms. CCA searches through the topology to find a set of potential cause entities upon which the symptom set depends. This symptom dependency set is composed of entities that are shared among the symptom entities, i.e., they can be linked via demonstrable topological dependencies.

CCA then reaches out to available assurance data repositories to further analyze and rank the symptom dependency set. An entity in the symptom set that is faulty or performing poorly is a potential root cause candidate for the symptom set. Automating the process of identifying and validating the status of such entities greatly expedites the root cause investigation process.

Where the status of an identified object in the symptom dependency set cannot be established due to lack of supporting data, CCA is able to trigger specific test routines that would generate on-demand assurance data sets if such test infrastructure is in place.

Finally, CCA presents the results of its data analysis to users and external systems, for both validation and escalation. Results can be presented as follows:

- › Interactive dashboards
  - › Allowing the exploration of symptom sets and corresponding symptom dependency sets
  - › Reporting on symptom dependency set entities and their status, helping in assessing whether any of them can be identified as a root cause for the symptom set
  - › Allowing the escalation of detected issues with the symptom dependency set entities
- › Non-interactive ticketing - data enrichment or escalation API

CCA is not domain-specific. It is fully generalized to work with any functional dependencies that are modelled in the topology:

- › Between IT systems and their dependent clients
- › Between network nodes and services and their dependent users
- › Between access resources (e.g., mobile RAN, DSL, FTTH) and the backhaul transport links

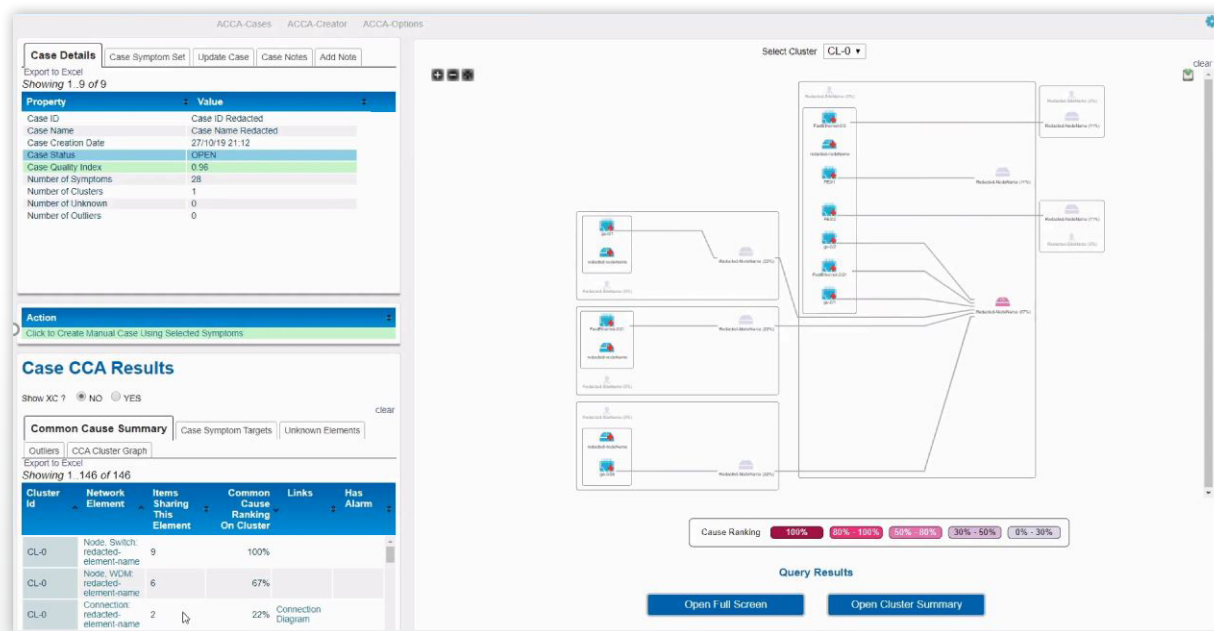


Figure 3. CCA in action.

CCA is much faster than traditional manual approaches of degradation identification, isolation, and analysis; it typically produces an ordered short-list of possible root causes within the minute while validating and eliminating false positives through the analysis of supporting symptom dependency set data, focusing on what really needs attention.

## PATH COMPUTATION IMPROVES SERVICE PROVISIONING

The Nova Context path computation module calculates a set of optimal network pathways according to customer-defined criteria such as latency and equipment preference, and outputs the information necessary for service provisioning. Additional constraints, such as locations or existing paths to avoid, can be passed to the module at runtime to support the calculation of diverse routes or other bespoke customer requirements.

By utilizing Nova Context's ability to easily combine multiple inventory and network data sources into a single reference model, paths spanning multiple networks can be readily calculated to make the best use of existing network capacity. Path computation results can be made available via REST APIs allowing them to be included in provisioning workflows in order to accelerate delivery times.

## FAULT AND PERFORMANCE DEGRADATION ALARM AUTOMATION ACROSS COMPLEX MULTIDOMAIN NETWORKS WITH SERVICE IMPACT ANALYSIS

The module for service impact analysis (SIA) determines the customer and service impacts that arise because of fault and performance degradation alarms across complex multidomain networks. Invented by EXFO, the semantic graph-based approach to SIA is widely deployed in automated or manual versions.

SIA contains a dependency library that is used to extend a Nova Context runtime topology graph with a specialized semantics. Once configured, the newly enhanced topology graph can be interpreted by a specialized and fast SIA rules reasoner, to efficiently calculate the entire set of customer and service impacts arising from alarms. SIA module associates existing alarms with the topology, meaning that they continue to be evaluated until they are cleared, making SIA a true representation of the impact of alarms on the real-time state of the network. The alarm input can be implemented via SIA API or a Kafka message bus. It can also be used in what-if scenario analysis by manually labelling modelled resources as alerted.

The SIA reasoner is not domain-specific. It is fully generalized to work with any entities that are modelled in the customer's topology given that consistent dependency relationships are understood and can be described.

## SIA CALCULATES AND SUMMARIZES CUSTOMER AND SERVICE IMPACTS/IMPEDANCE

- › Dedicated runtime server for optimal process performance
- › Rapid impact analysis
- › Impacts summarized in a single real-time transaction
- › Proactively supports alarm reduction

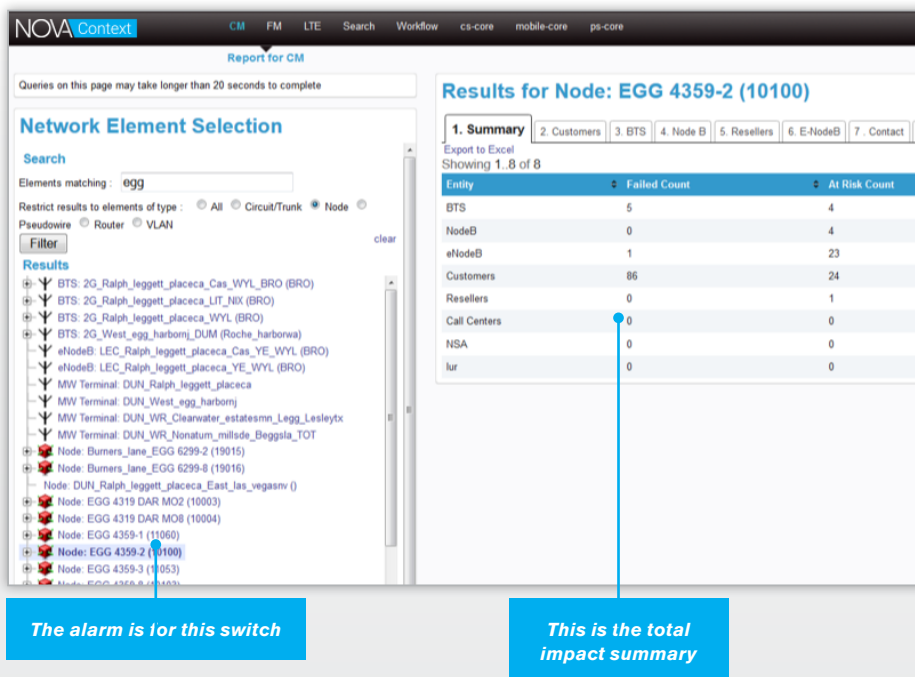


Figure 4. Real-time service impact analysis (RT-SIA)

## VISUALIZING FAULTS AND TICKETS IN TOPOLOGY

- › Customer calls with issues
- › Technician searches Nova Context with customer data (name, account or circuit ID)
- › See real-time customer services and topology with alarms or tickets
- › Pop-up view of alarms/tickets with dynamic information from alarm/ticketing system
- › Supports deep linking into third-party system for more details

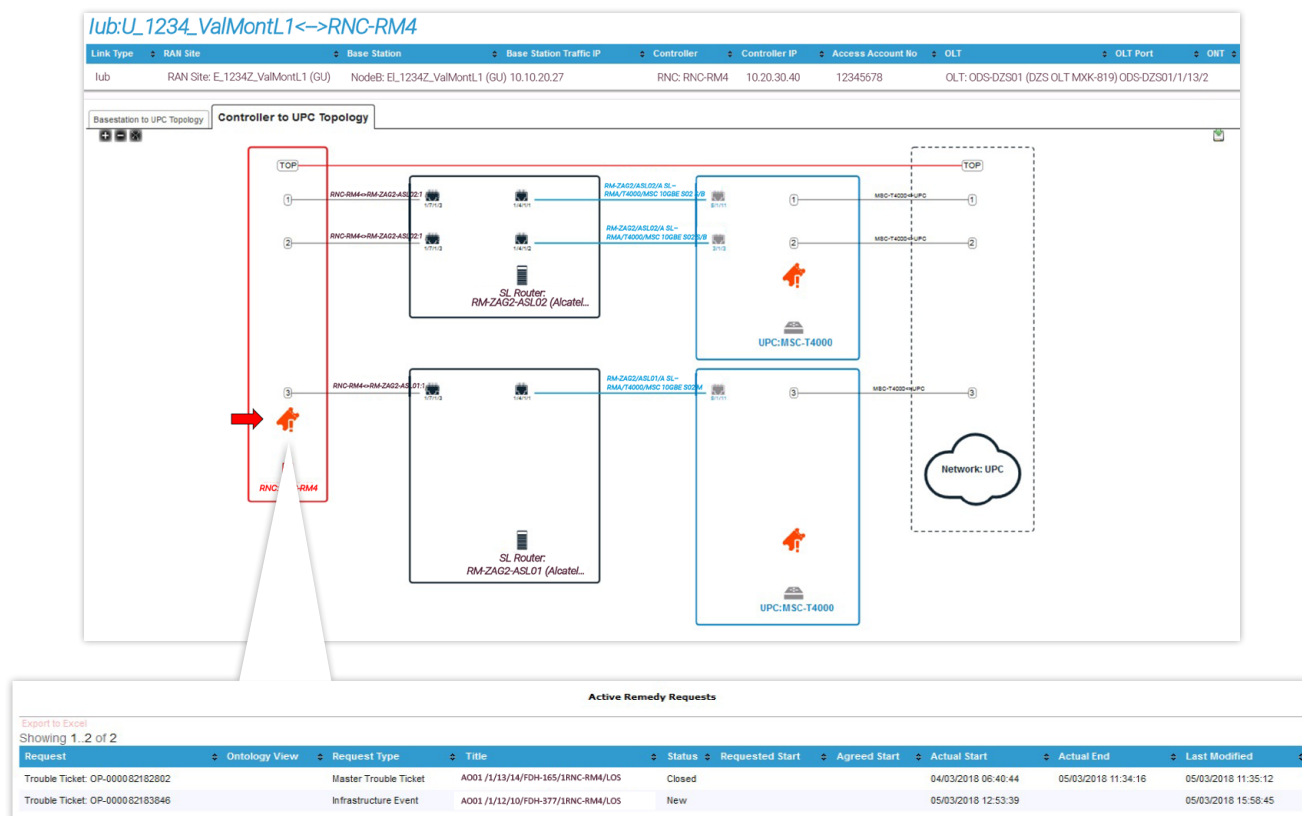


Figure 5. RT-SIA: visualizing faults and TT's in topology

## ANALYZE SETS OF CHANGES FOR UNINTENDED CONSEQUENCES AND UNNECESSARY DUPLICATION WITH CHANGE IMPACT ANALYSIS AND PLANNING

Nova Context enables change management teams to manage change packages through their complete lifecycle, confident in the knowledge that the interactions with other network change plans and the impact on services and customers is being correctly and effectively handled.

The design and feature content of Nova Context's CIAP module is guided by four goals:

### 1. Always plan against accurate network and customer data

Planning changes or upgrades to network infrastructure, as well as disaster planning, are all essential disciplines, but without complete and up-to-date information, decisions may be made that negatively impact on customers or carrying capacity. CIAP is (only) deployable aside a Nova Context Platform which provides dynamic topology dependency models and/or real-time inventory.

## 2. Make the calculation of planned impacts and conflicts easy and obvious, and proactively warn stakeholders if the severity of a pending plan is modified

Dynamic network topology modeling provides a complete understanding of the dependencies within the network to visualize these aggregated changes, and to flag any that cause conflicts.

## 3. Make distributed planning safe

Changes don't happen in isolation; multiple simultaneous planned changes that may impact the same infrastructure components can cause problems and missed opportunities. The change impact analysis and planning product module allows different change planners to submit simultaneous change plans. Impacts are calculated, severities assigned, and central planning supervisors are informed when conflicts occur.

Additionally, if desired, change impact analysis and planning can be integrated with a workflow management system like BMC remedy. This minimizes disruption to existing business processes whilst still providing fully automated change impact analysis.

## 4. Reduce the cost of re-planning

Change impact analysis allows the plan supervisor to investigate the scheduling of plans to determine if cost savings can be made by combining overlapping changes or to see if rescheduling conflicting changes can remove calculated change plan collisions.

This helps to avoid potential outages and service degradations. Centralized coordination of multiple change plans can also help to reduce both cost and the length of disruptions by combining changes that impact the same parts of the network or the same top-level services.

## AUTOMATICALLY FIND COLLISIONS BETWEEN SEPARATE CHANGE PLANS ACROSS ORGANIZATIONS AND NETWORK DOMAINS

- › Global automation for cross domain manual change planning
- › Reduces customer outages
- › Increases service uptime

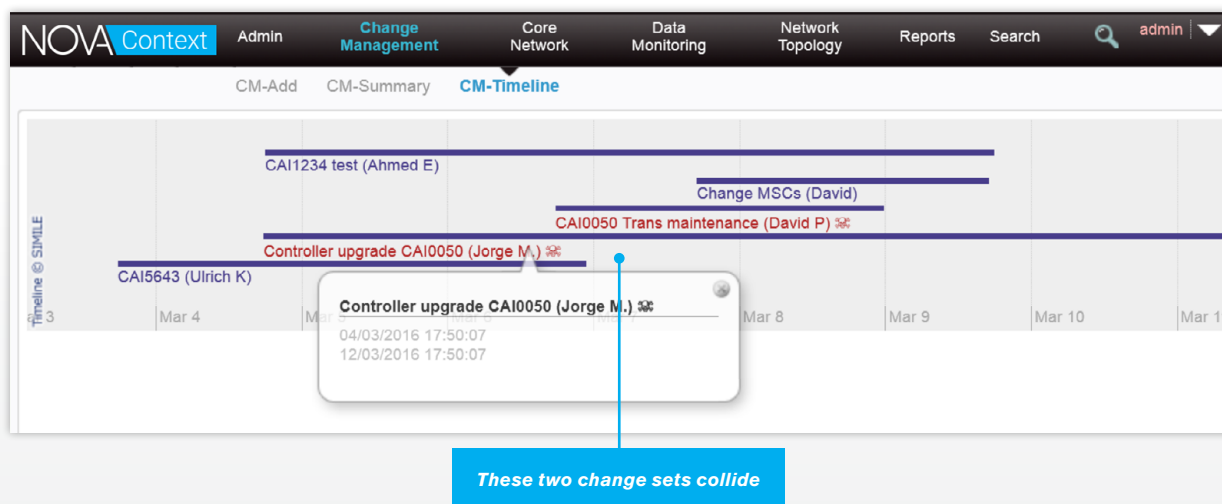


Figure 6. Change impact analysis and planning (CIAP)

## CHANGES SHOWN IN TOPOLOGY VIEWS

- › Technician searches Nova Context with customer data (name, account or circuit ID)
- › See real-time customer services and topology with pending changes
- › Pending changes are shown
- › Pop-up view of changes with dynamic information from salesforce system
- › Supports deep linking into third-party system for more details

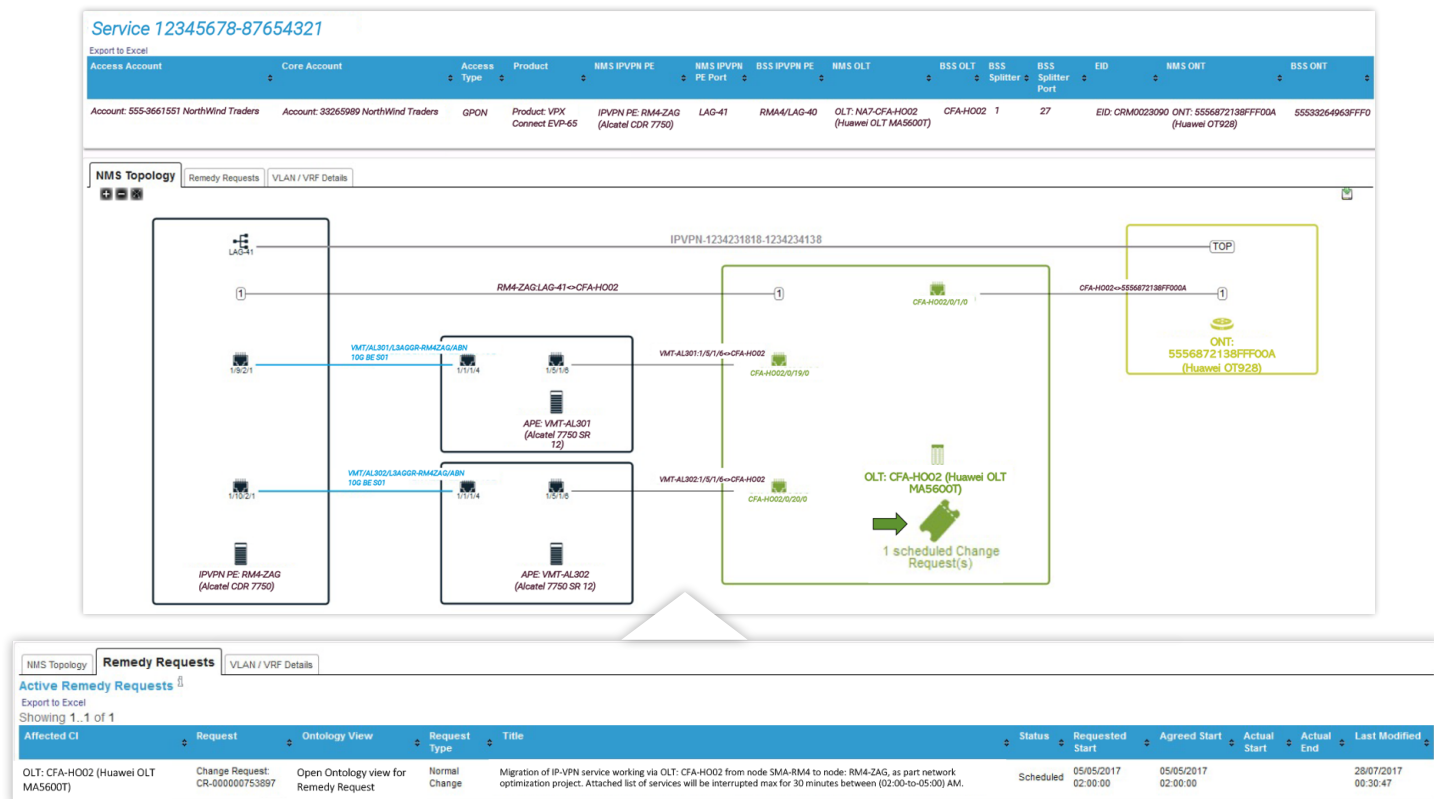


Figure 7. SOC dashboard and drill-down

EXFO headquarters T +1 418 683-0211 Toll-free +1 800 663-3936 (USA and Canada)

EXFO serves over 2000 customers in more than 100 countries. To find your local office contact details, please go to [www.EXFO.com/contact](http://www.EXFO.com/contact).

For the most recent patent marking information, please visit [www.EXFO.com/patent](http://www.EXFO.com/patent). EXFO is certified ISO 9001 and attests to the quality of these products. EXFO has made every effort to ensure that the information contained in this specification sheet is accurate. However, we accept no responsibility for any errors or omissions, and we reserve the right to modify design, characteristics and products at any time without obligation. Units of measurement in this document conform to SI standards and practices. In addition, all of EXFO's manufactured products are compliant with the European Union's WEEE directive. For more information, please visit [www.EXFO.com/recycle](http://www.EXFO.com/recycle). Contact EXFO for prices and availability or to obtain the phone number of your local EXFO distributor.

For the most recent version of this spec sheet, please go to [www.EXFO.com/specs](http://www.EXFO.com/specs).

In case of discrepancy, the web version takes precedence over any printed literature.