IP Network eBook Series

# Cloud Data Center Intelligent O&M Solution

Author : Yanlin Zhang

Information Digitalization and Experience Assurance (IDEA) Department

HUAWEI

# Copyright

| | |
|---|---|
| Author: | Yanlin Zhang |
| Key Contributors: | Le Chen, Fan Zhang, Hua Yang, Chengxia Yao, Lingling Yu, Guixiang Chen, Xuezhao Li |
| Release Date: | 2021-07-19 |
| Issue: | 01 |

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

# Preface

## Author Introduction

**Yanlin Zhang**: joined Huawei in 2020 as a data center (DC) switch documentation engineer and is engaged in developing documentation for data communication products.

## About This Book

This book provides a detailed description of the background, basic architecture, technical implementation, typical cases, and evolution trends of Huawei's Cloud Data Center Network (DCN) Intelligent Operations and Maintenance (O&M) Solution. The information presented in this book is easy to understand and very practical, granting you a good command of this solution.

# Intended Audience

This book is intended for network O&M personnel, mid- and senior-level executives of enterprises, and anyone interested in the Cloud DC Intelligent O&M Solution.

# Table of Contents

# Chapter 1

# Why Do Cloud DCNs Require Intelligent O&M?

**Abstract**

This chapter provides an overview of the challenges facing cloud DCN O&M during service change, routine preventive maintenance inspection (PMI), emergency recovery, and root cause locating.

A DCN is the infrastructure that carries DC services, and consists of two parts: a communication system (built on computing and storage subsystems) and a network management and monitoring system. DCNs allow enterprises or organizations to store, manage, and propagate information conveniently.

With the SDN and cloud computing eras now upon us, pooling computing, storage, and network resources are ever more commonplace, and networks and services are becoming increasingly automated. While this facilitates enterprise digital transformation, it also makes cloud DCNs more complex.

Typically, a cloud DCN consists of multiple systems with different functions. This means that to operate and manage a cloud DCN, network personnel must have a good command of hardware, networks, servers, storage, and security, in addition to being familiar with the ICT resource requirements of different services.

## Challenges to Cloud DCN O&M

Typically, the challenges to cloud DCN O&M fall into one of four categories: service change, routine PMI, emergency recovery, and root cause locating.

### Service Change

- **Difficult to measure resource utilization.** The routine approach to measuring resource utilization during batch delivery of services and modules is to manually log in to the affected devices. This, however, cannot accurately determine whether resource preemption occurs during the delivery, not to mention the impacts of resource preemption on existing services.
- **Difficult to detect network changes.** In reality, it is the SDN controller that is responsible for understanding and translating service needs and automatically delivering network configurations in batches. This, however, makes network configurations hard to be fully grasped by administrators.
- **Complex to roll back configurations.** After troubleshooting, network O&M personnel often forget to roll back configurations, leaving potential risks.
- **Difficult to rapidly perform capacity expansion.** When adding a large number of servers on a network, manual operations will lead to low efficiency, failing to keep pace with service provisioning.

### Routine PMI

- **Unable to perform network-wide evaluation.** Conventional DCNs are operated and maintained based on alarm events. This approach, however, is inapplicable to cloud DCNs, which have a complex network architecture consisting of a physical network (underlay) and a virtual network (overlay).
- **Unable to predict network faults.** In conventional O&M methods, O&M personnel respond reactively to network faults, and cannot predict or prevent faults before they occur. This results in poor service quality and user experience.
- **Unable to comprehensively measure network connectivity.** Currently, network connectivity is typically measured using the **ping** and **traceroute**

commands. However, this approach requires highly experienced O&M personnel to ensure accuracy, and cannot be applied in all scenarios. This makes it especially difficult to measure DCs that feature ultra-large scales and service provisioning virtualization.

### Emergency Recovery

In recent years, large enterprises in the Internet, finance, and other sectors have built DCs to store, analyze, and compute their huge volumes of data. The DCs play a pivotal role in safeguarding enterprise information and promoting their business success. As a result, even the slightest interruption to a DC will cause serious economic loss. As such, it is paramount that DC faults are rectified before they cause any damage to enterprise services.

### Root Cause Locating

To provide high reliability and bandwidth, DCNs are typically designed to forward traffic in the equal cost multi-path (ECMP) mode, in which traffic is steered via the hash algorithm and the traffic forwarding paths increase exponentially as the number of network nodes increases. With so many forwarding paths, however, it is hard for administrators to select the optimal one.

On DCNs, the protocol parsing and traffic forwarding modes of different types of servers and network interface cards (NICs) vary. As well as this, there are also many other types of devices on the network, such as firewalls and load balancers (LBs). These devices may come from different vendors and work in different ways, making it difficult to rapidly locate and rectify faults on DCNs.

Conventionally, fault locating heavily depends on the experience of O&M personnel and is extremely time-consuming. Such an approach is not feasible as DCs continue to grow in scale and configuration complexity. All of this positions intelligent O&M as a must-have for cloud DCN O&M.

# Chapter 2
# Popular Industry Solutions

**Abstract**

This chapter takes a panoramic look at the most popular technologies used within the industry to operate and maintain cloud DCNs.

## Key Expectations

To fully embrace the cloud era, we need to build an all-new system capable of intelligently operating and maintaining SDN DCs to replace the current solutions which favor manual DC O&M based on expert experience. Specifically, this new system must deliver the following:

- **Proactive response**: In some SDN scenarios, such as creating or deleting logical networks on demand, network or service configurations need to be changed frequently, requiring services to be quickly and dynamically provisioned. However, this can often increase the chance of network faults. For this reason, our brand-new O&M system must be highly proactive and intelligent in order to detect such faults before they cause any damage to the network. The system can also draw on big data analytics and expert experience databases to help users rapidly locate and rectify faults, transforming compliant-driven network O&M into service-centric proactive O&M.

- **Real-time performance**: For instance, an enterprise complained about transient packet loss on its lightly loaded network and suspected that there were millisecond-level traffic bursts. These bursts, however, could not be detected via Simple Network Management Protocol (SNMP) — a minute-level traffic detection mechanism. To address this deficiency, the new O&M system must be able to detect microburst exceptions in real time.
- **Large-scale network management**: On today's networks, the O&M system needs to manage not only physical devices but also virtual machines (VMs), involving dozens of times more managed NEs than before. In addition, as real-time analysis becomes a must-have, the O&M system is required to collect device data more frequently — often at intervals of milliseconds instead of minutes. This in turn increases the data volume almost 1000-fold. To keep up, the O&M system must be able to collect, store, analyze, and display masses of data.

# Technology Preferences

To meet these requirements, our new O&M system needs to utilize a number of cutting-edge technologies in terms of both software and hardware to proactively detect network services in real time. This section will detail some popular software and hardware technologies.

Typical examples of software technologies are telemetry and Encapsulated Remote Switched Port Analyzer (ERSPAN).

### Telemetry

Conventional networks are typically operated and maintained through SNMP. This approach gradually becomes ineffective as new DC applications emerge and pose higher requirements on network performance. To elaborate, this approach involves the following drawbacks:

- **Inaccurate data**: Typically, SNMP collects device data (such as interface traffic volumes, link status, memory usage, and CPU usage) every 3 – 5 minutes. This collection interval is too long to ensure high data accuracy.

As shown in Figure 2-1, SNMP indicates that the network is running properly. In reality, however, packet loss has already occurred on the network due to microbursts, and service experience has been affected.

Figure 2-1 Bandwidth utilization presented by SNMP vs. actual bandwidth utilization

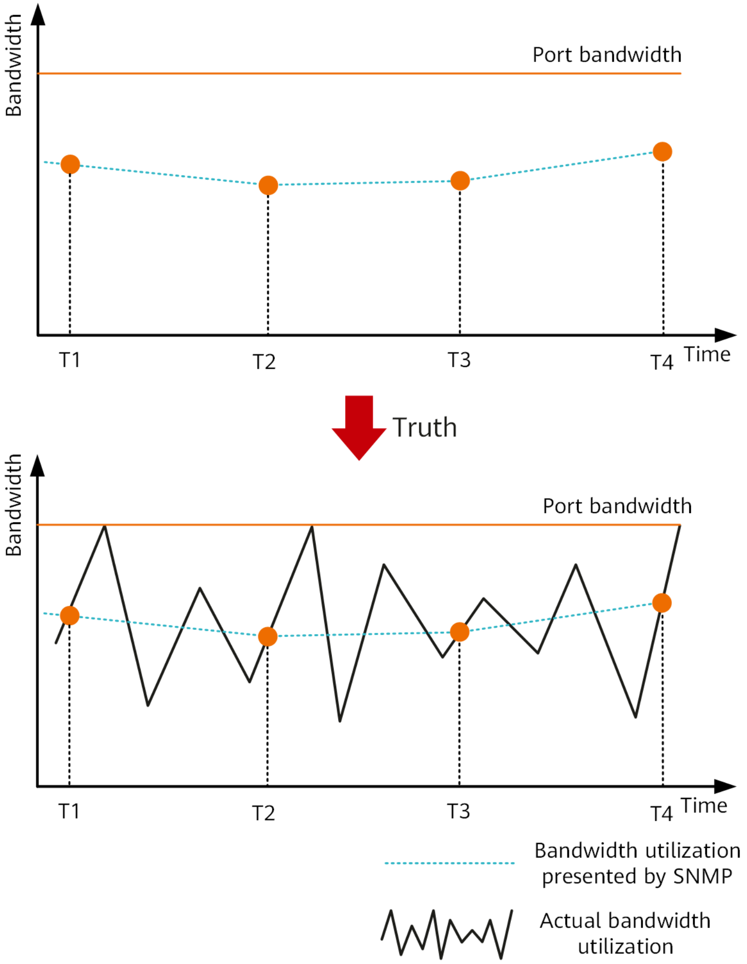

- **Resource-intensive**: In most cases, SNMP is configured to collect device data at shorter intervals, resulting in devices becoming congested as they constantly respond to SNMP requests. The end result is a system overload.
- **Prone to packet loss**: SNMP trap messages are carried over User Datagram Protocol (UDP), which is natively prone to packet loss.

The entire industry is finding ways to overcome these drawbacks, and this is where telemetry technology comes in. Telemetry is ideal for coping with increasingly complex network O&M. Specifically, telemetry codes are embedded into the hardware board chips, so that telemetry can obtain device data from the boards and proactively send it to the O&M system in real time. This ensures that O&M personnel remain aware of the actual status of devices and networks.

### ERSPAN

Most DC applications run over Transmission Control Protocol (TCP). For enhanced operation and maintenance of DCs, O&M personnel must determine whether these applications can access each other, whether user experience deteriorates, and how to rapidly locate the forwarding path of each service flow upon a TCP connection exception.

This is where ERSPAN comes in, as it can collect network-wide TCP packets (such as SYN, FIN, and RST packets) by mirroring them, and send the collected data to the O&M system. In most cases, ERSPAN can collect the following information:

- Forwarding paths of TCP packets
- TCP start and end time
- Transmitted bytes
- Forwarding latency of TCP packets
- Exception information: latency > upper threshold, TCP Flags exception (RST), TCP retransmission, and TTL < 3

The O&M system can use this information to restore the forwarding paths of TCP flows and prompt possible fault points and root causes.

In addition to these software technologies, the industry is also finding ways to advance device hardware, and has adopted a number of innovative technologies, such as server clustering.

### Server Clustering

As DCNs continue to grow in scale, their O&M systems need to collect, analyze, process, store, and display even larger volumes of device data. This requires the systems to:

- Run on dedicated hardware devices, such as servers.
- Deploy servers in clusters. Based on their division of responsibilities, server clusters are further classified into collector clusters and analyzer clusters.
- Be highly reliable and scalable. To achieve this, each server cluster is built on the microservice architecture, and service modules are deployed in multi-instance mode.
- Assign different and clear responsibilities to system components. Specifically, after data subscription, the collection module collects data in seconds and sends it to the high-throughput distributed message system that is responsible for data buffering and distribution. The service module then performs data analysis and calculation based on the AI algorithms and expert experience, and saves the processed data to the distributed data storage system. Finally, detailed data is displayed on the graphical user interface (GUI) for users to check.

# Chapter 3

# Huawei's Cloud DC Intelligent O&M Solution

**Abstract**

This chapter offers a full picture of Huawei's Cloud DC Intelligent O&M Solution.

## Solution Overview

As cloud computing and intelligence technologies gain momentum, DCs continue to grow in scale and networking complexity, challenging enterprise service availability. According to an authoritative survey, a service interruption that lasts for just 1 hour can cause an economic loss of over US$100,000. All of this makes it imperative to implement intelligent DCN O&M.

This is where Huawei's Cloud DC Intelligent O&M Solution comes in. It checks and diagnoses the health of DCs with intelligent fault remediation and comprehensive network health check, two differentiators in this solution.

- **Intelligent fault remediation**: This future-oriented solution can rapidly detect network faults and find their root causes before they cause any damage to the network. Drawing on Huawei's more than 30 years of network O&M expertise and data from more than 7800 live DCs and thousands of fault drills, this solution summarizes 75 types of typical faults in 6 categories. Moreover, equipped with an intelligent knowledge inference engine, the solution can implement "1-3-5" intelligent troubleshooting — detect faults in 1 minute, locate them in 3 minutes, and rectify them in 5 minutes.

- **Comprehensive network health check**: This solution provides a wide range of network health check capabilities — represented by network health management and proactive fault prediction, which comprehensively evaluate network health from five dimensions: device, network, protocol, overlay, and service.

## Solution Architecture

Figure 3-1 shows the architecture of Huawei's Cloud DC Intelligent O&M Solution, which is logically divided into three layers: network layer, control layer, and analysis layer.

- **Network layer**: consists of DCN devices that report mirrored packets, performance data, and logs to the analysis layer for further handling and presentation. The network layer can be considered as a data source of the analysis layer.
- **Control layer**: is built on iMaster NCE-Fabric — an intelligent network management and control system. It interconnects with iMaster NCE-FabricInsight — an intelligent network analysis system — to automatically translate and deliver configurations during O&M.
- **Analysis layer**: is built on iMaster NCE-FabricInsight. It can proactively detect network faults and locate their root causes in minutes.

Figure 3-1 Architecture of Huawei's Cloud DC Intelligent O&M Solution

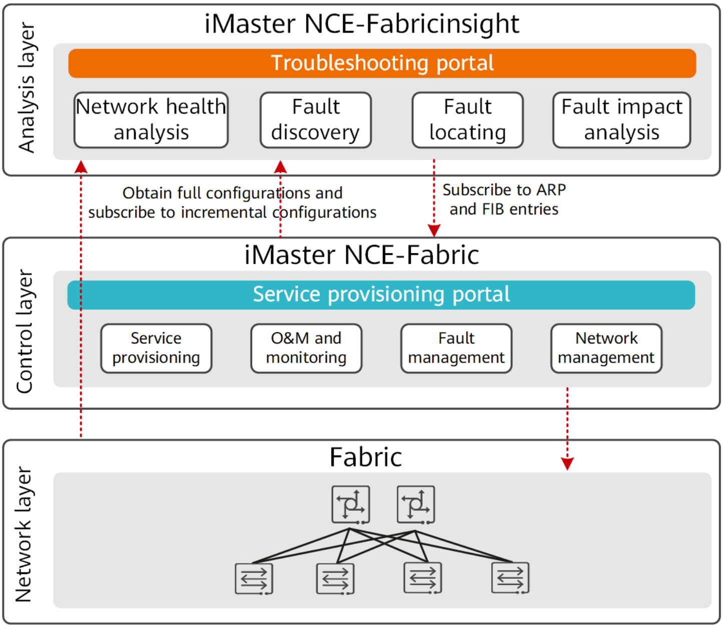

The following part goes into detail about the solution's key components — iMaster NCE-Fabric and iMaster NCE-FabricInsight.

## iMaster NCE-FabricInsight

Based on Huawei's big data platform, FabricInsight can receive data from network devices via telemetry and use intelligent algorithms to analyze and display the received data, as illustrated in Figure 3-2.

iMaster NCE-FabricInsight is built on a microservice architecture in which service modules are deployed in multi-instance mode, thereby delivering high scalability and reliability.

Figure 3-2 Architecture of iMaster NCE-FabricInsight

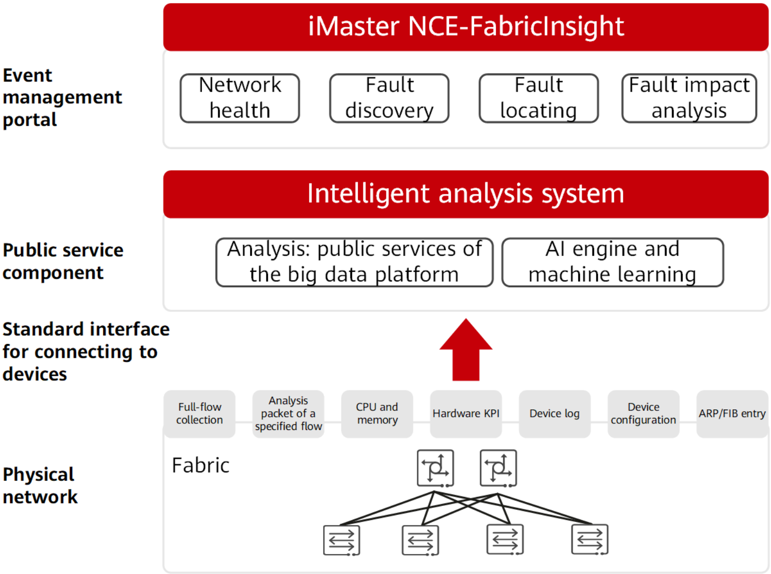

1. **Scalable**: The system capacity can be easily expanded by adding instance nodes. These nodes do not communicate with each other. Instead, the message bus delivers HTTP messages to them for further handling.
2. **Reliable**: The analyzer is connected to the collector in the southbound direction and uses Linux Virtual Servers (LVSs) to deliver system reliability.

The architecture of iMaster NCE-FabricInsight consists of three parts: network device, collector, and analyzer.

- **Collector**

  The collector is responsible for collecting various types of switch data, including ERSPAN mirrored TCP packets, performance metrics reported via Google Remote Procedure Call (gRPC), and FIB/ARP entries. It sends this data to the analyzer for further handling.

- **Analyzer**

The analyzer can comprehensively collect, analyze, and display data. Specifically, it can:

- Parse and compute different types of data. For example, it computes the forwarding path, forwarding latency, and link latency of packets.
- Analyze the relationships between applications that interact with each other and associate applications with network paths.
- Build dynamic baselines for performance metrics using intelligence algorithms to detect and predict network exceptions.

## iMaster NCE-Fabric

To overcome the challenges to DCN O&M, Huawei developed a controller dedicated to DCs — iMaster NCE-Fabric.

Figure 3-3 Architecture of iMaster NCE-Fabric

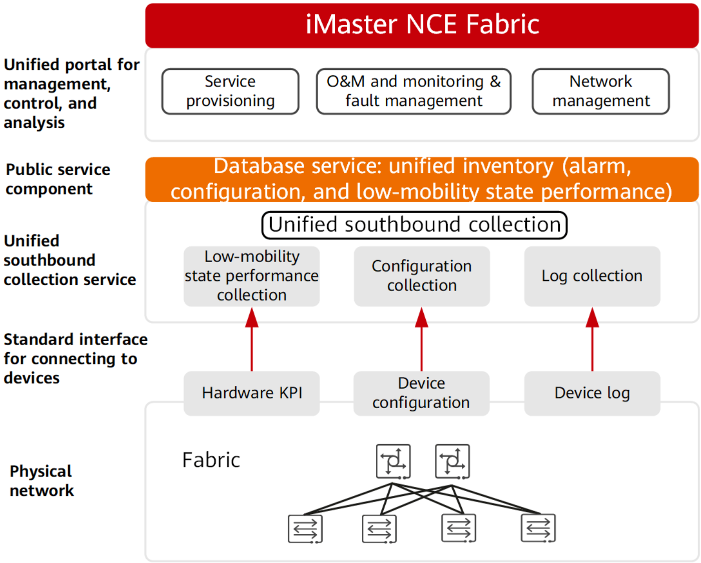
iMaster NCE Fabric
Unified portal for management, control, and analysis
Service provisioning
O&M and monitoring & fault management
Network management
Public service component
Database service: unified inventory (alarm, configuration, and low-mobility state performance)
Unified southbound collection service
Unified southbound collection
Low-mobility state performance collection
Configuration collection
Log collection
Standard interface for connecting to devices
Hardware KPI
Device configuration
Device log
Physical network
Fabric

iMaster NCE-Fabric is a Huawei-developed next-generation SDN controller that is oriented at enterprise and carrier DC markets. As a centralized network control plane, it automatically delivers network configurations to devices, revitalizing service automation. Not only this, this controller provides a lineup of future-proof features — such as cloud-network integration and automated network orchestration — and can provide visualized and fine-tuned network O&M, delivering high reliability and openness.

As shown in Figure 3-3, the controller consists of four parts: a management-control-analysis application layer, public service component layer, unified southbound collection service layer, and southbound standard interface layer. In particular, the first layer is tailored for network O&M and provides three main functions: O&M and monitoring & fault management, network management, and service provisioning.

As mentioned already, the challenges to cloud DCN O&M fall into four categories: service change, routine PMI, emergency recovery, and root cause locating. And the entire industry is finding ways to overcome them. For this reason, the Cloud DC Intelligent O&M Solution makes its debut. The real question is, how does it overcome these challenges? Find out in the following chapters.

Chapter 4

# Solution Highlights for Service Changes

**Abstract**

This chapter begins by describing the differentiators of Huawei's Cloud DC Intelligent O&M Solution in service change scenarios. They include simulation verification, real-time insights into network change differences, full-lifecycle VM tracing, configuration rollback, and server capacity expansion. The chapter then goes on to discuss some typical use cases.

## 4.1 Simulation Verification

### Benefits

As SDNs continue to grow in scale, automated network deployment and multi-dimensional network O&M are gradually becoming essential, bringing unprecedented convenience to users. That said, this has brought up the following question: As the batch delivery of services and modules becomes a new norm in SDN scenarios, how can we ensure sufficient device resources for each service delivery?

Conventionally, to evaluate the impact of network changes, O&M personnel must log in to each device to check the resource utilization and then send the corresponding data to experts for further evaluation. This approach is inefficient and cannot ensure the evaluation accuracy.

This is where the simulation verification function comes in. With this function, iMaster NCE-Fabric can achieve the following:

- Continuously monitor device resources to obtain the resource utilization in real time.
- Use control plane verification (CPV) technology to accurately estimate the network resources required for each service rollout and determine whether the remaining network resources are sufficient.

This minimizes the probability of service rollout failures or other faults due to insufficient resources.

This process is like decorating a house. By referencing a three-dimensional (3D) simulation picture, we can easily figure out whether the decoration meets our expectations and what size of furniture we need to purchase, eliminating the risk of human errors or negligence.

# Implementation

Figure 4-1 shows how iMaster NCE-Fabric performs simulation verification. We first need to cover some key concepts, which will help you understand the entire process.

- **Design state**: in which iMaster NCE-Fabric orchestrates virtual private cloud (VPC) services. In this state, iMaster NCE-Fabric delivers the orchestration results of VPC services to a design-state database. After user confirmation, the orchestration data is submitted to a production-state database and delivered to physical devices.
- **Production state**: in which iMaster NCE-Fabric orchestrates VPC services and delivers the orchestration results to physical devices.

- **CPV:** enables iMaster NCE-Fabric to simulate and verify the results of the service orchestration performed in the design state.

Figure 4-1 Simulation verification process

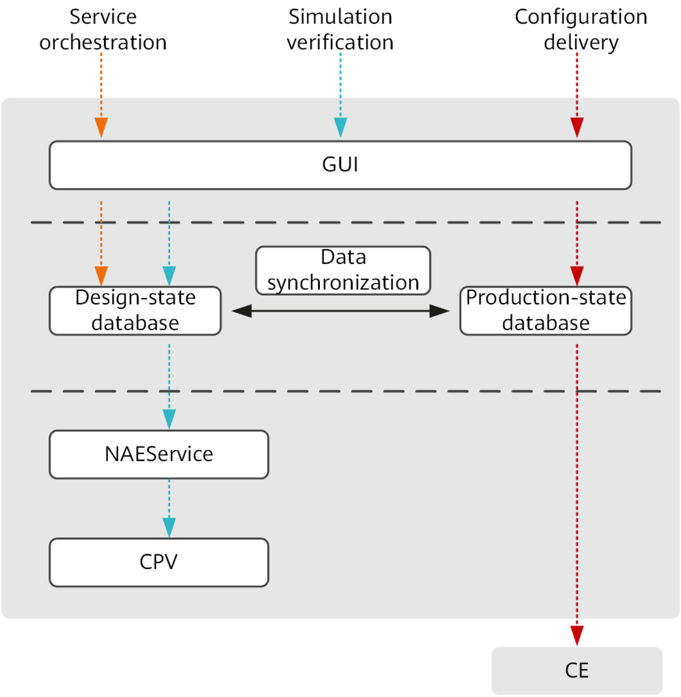

After completing service orchestration in the design state, iMaster NCE-Fabric simulates and verifies the orchestration results and generates a simulation report, by reading which O&M personnel can learn well the resource utilization and configuration changes. Before delivering configurations to network devices in the production state, iMaster NCE-Fabric needs to check network connectivity, evaluate the impact of configuration changes on existing services, and confirm that all these items meet user expectations.

# 4.2 Real-Time Insights into Network Change Differences

## Benefits

We often need to make changes on cloud DCNs when they are running. The issue with this is that service changes most often lead to a heavy workload. For instance, we need to verify whether thousands of configurations have changed, whether tens of thousands of entries on a single device can be learned, and whether the involved applications can properly communicate with each other. This can take us several hours if we use conventional approaches that favor manual operations.

To simplify service changes and free O&M personnel from endless verification, Huawei's iMaster NCE-FabricInsight provides real-time network snapshot. With this function, the controller takes snapshots of device configurations, ARP entries, ND entries, and RIB entries in real time, and compares the snapshots taken before and after changes, enabling any minor configuration difference to be detected immediately.

## Implementation

In line with the preceding description, iMaster NCE-FabricInsight can take snapshots of network-wide or specific device configurations and detect any configuration difference by comparing the snapshots taken before and after changes, enabling O&M personnel to know well about the network status.

As illustrated in Figure 4-2, iMaster NCE-Fabric displays the snapshots of a device's running configurations, ND entries, IPv4 routing entries, ARP entries, and IPv6 routing entries before and after a change.

The controller can also display the time when a snapshot is taken, and the numbers of identical lines, changed lines, added lines, and deleted lines in configuration files before and after a change, as shown in Figure 4-3.

Figure 4-2 Snapshots of network configurations and entries before and after a change

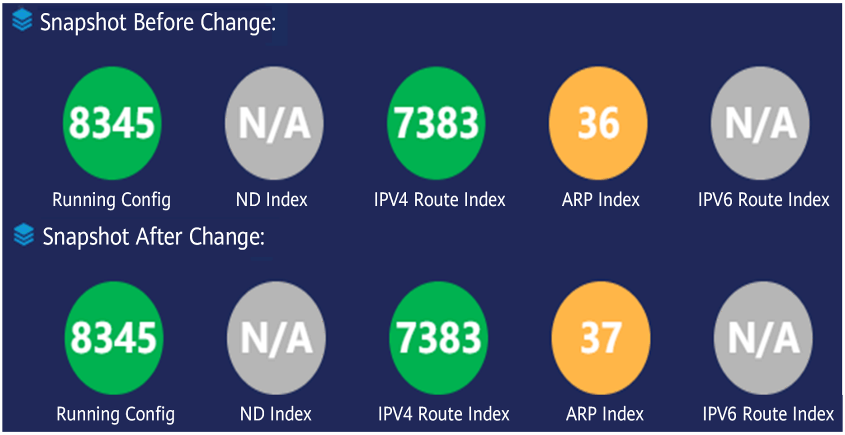

Snapshot Before Change:

8345 Running Config

N/A ND Index

7383 IPV4 Route Index

36 ARP Index

N/A IPV6 Route Index

Snapshot After Change:

8345 Running Config

N/A ND Index

7383 IPV4 Route Index

37 ARP Index

N/A IPV6 Route Index

Figure 4-3 Configuration and entry comparison

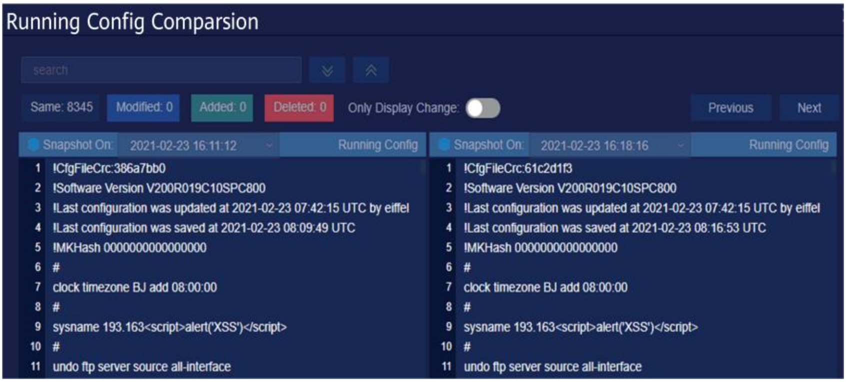

Running Config Comparsion

Snapshot On: 2021-02-23 16:11:12 Running Config

```
!CfgFileCrc:386a7bb0
!Software Version V200R019C10SPC800
!Last configuration was updated at 2021-02-23 07:42:15 UTC by eiffel
!Last configuration was saved at 2021-02-23 08:09:49 UTC
!MKHash 0000000000000000
#
clock timezone BJ add 08:00:00
#
sysname 193.163<script>alert('XSS')</script>
#
undo ftp server source all-interface
```

Snapshot On: 2021-02-23 16:18:16 Running Config

```
!CfgFileCrc:61c2d1f3
!Software Version V200R019C10SPC800
!Last configuration was updated at 2021-02-23 07:42:15 UTC by eiffel
!Last configuration was saved at 2021-02-23 08:16:53 UTC
!MKHash 0000000000000000
#
clock timezone BJ add 08:00:00
#
sysname 193.163<script>alert('XSS')</script>
#
undo ftp server source all-interface
```

# 4.3 Full-Lifecycle VM Tracing

## Benefits

As DCNs move to the cloud, more VMs need to be deployed or migrated, and this can be automatically performed by the virtual machine management (VMM). The drawback lies in the fact that O&M personnel become blind to VM migration, logout, and distribution information, leading to increasingly complex O&M.

To resolve this, iMaster NCE-FabricInsight provides the IP 360 function, which is ideal for implementing full-lifecycle VM management. With this function, FabricInsight can achieve the following:

- Display device and VM statistics in real time, such as the number of online VMs and top N switches accessed by VMs. This information facilitates network resource planning.
- Manage VMs throughout their lifecycle, enabling real-time visualization of VM login, logout, and migration events.
- Provide network-wide IP snapshot analysis to detect all IP address changes and VM exceptions before and after network changes.

## Implementation

IP 360 is further classified into six sub-functional modules: IP subnet, IPv4 distribution, IP address, IP snapshot comparison, subnet import, and VM discovery.

- **IP subnet**: collects statistics on the system subnets planned by users, including the subnet address, name, home fabric, VRF, address space, and online rate, and allows users to edit or delete subnets in batches, add a single subnet, or export all subnet details.
- **IPv4 distribution**: displays the usage and status of each IP address. The IP address status options include: online, transient, offline, excluded, and invalid. On top of that, this module can display the top and bottom 10 subnets ranked by IP address usage.

- **IP address**: displays IP address information about physical machines (PMs) and VMs in a list. This list includes their IP addresses, names, MAC addresses, fabric networks, access devices, gateway interfaces, first discovery time, last discovery time, status (active/inactive), and discovery mode. In addition, IP addresses that frequently migrate are displayed at the top of the list.
- **IP snapshot comparison**: compares snapshots saved periodically and displays the online, offline, migrated, and unchanged IP addresses. We can filter the snapshots by fabric network or IP address.
- **Subnet import**: allows users to import planned subnet information using a template downloaded from the system.
- **VM discovery**: displays all fabric information, including the fabric name, fabric networking type, number of VMs in the fabric, and VM synchronization mode set for the fabric.

Figure 4-4 IP 360 analysis details

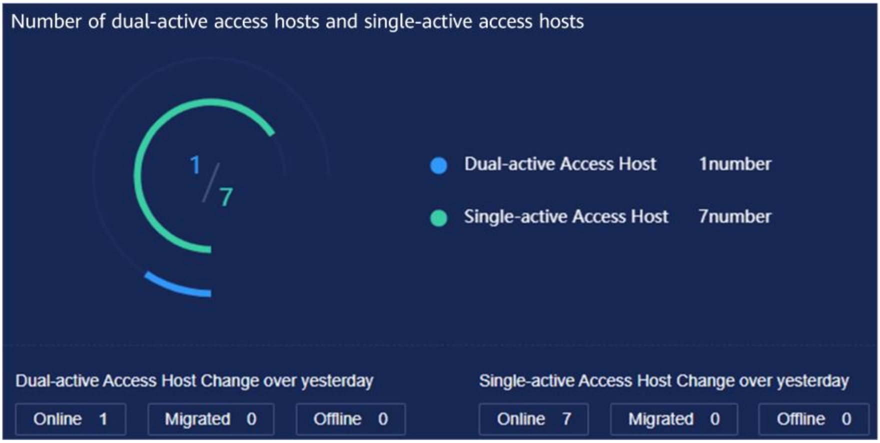
Number of dual-active access hosts and single-active access hosts
1/7
Dual-active Access Host 1number
Single-active Access Host 7number
Dual-active Access Host Change over yesterday
Online 1
Migrated 0
Offline 0
Single-active Access Host Change over yesterday
Online 7
Migrated 0
Offline 0

With IP 360, iMaster NCE-FabricInsight can achieve the following:

- Detect any IP address changes of PMs and VMs.
- Collect ARP/ND entries of network-wide devices via telemetry.
- Collect fabric information via VM discovery.

All of this information enables full-lifecycle VM tracing that spans VM login, migration, and logout.

As illustrated in Figure 4-4, the IP 360 function monitors the IP addresses of PMs and VMs, collects VM statistics (such as the number of online VMs and top 10 switches accessed by VMs), and comprehensively displays the IP address details of network-wide VMs and PMs.

# 4.4 Configuration Rollback

## Benefits

To facilitate troubleshooting, O&M personnel often make temporary configuration changes on switches, such as creating a temporary VLAN, but sometimes forget to roll back these operations after the faults are rectified, potentially creating many risks.

To prevent this, Huawei's iMaster NCE-Fabric provides network-wide and tenant-level configuration rollback, rapidly restoring services when faults occur and minimizing service interruption loss.

## Implementation

### Network-Wide Rollback

In some cases, we need to change network configurations at scale. However, any improper change will lead to severe service interruptions. We therefore need to protect services by rapidly rolling back network configurations. This is where network-wide rollback comes in. With this function, we can rapidly restore network configurations to those at a certain time point, ensuring service continuity and minimizing the impact of changes. As shown in Figure 4-5, this function is applicable to all customer edges (CEs) managed by iMaster NCE-Fabric.

To elaborate, with this function, iMaster NCE-Fabric takes snapshots of network-wide configurations (including all the data in the database and configuration files of CEs) at different time points and rolls back the configurations as required without having to restart the system.

Figure 4-5 Network-wide configuration backup and restoration process

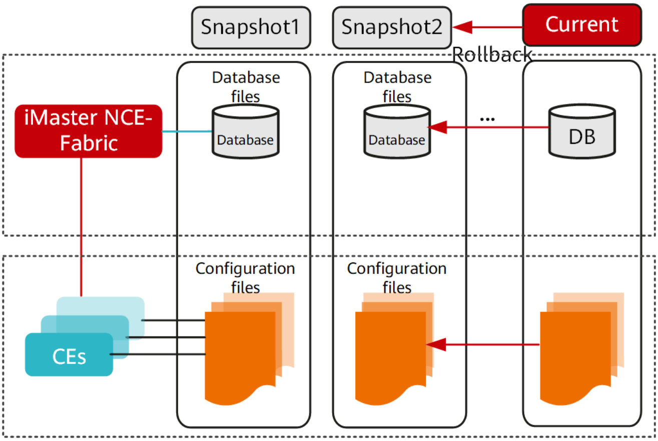

### Tenant-Level Rollback

In tenant scenarios, if a change error or unexpected change occurs, we may need to restore the network configuration of a single tenant to that at a specific time point. Tenant-level rollback is specifically designed for this purpose. With this function, the data of each tenant can be backed up or rolled back without impacting other tenants. It is typically used in computing and hosting scenarios.

Before performing tenant-level rollback, we need to back up and generate snapshots of each tenant's network configurations. As illustrated in Figure 4-6, the network configuration of tenant A is backed up at time points T1 and T2, and Snapshot1 and Snapshot2 are generated, respectively.

Assume we need to roll back the configuration of tenant A to that at T2. As shown in Figure 4-7, the entire rollback process is as follows:

Figure 4-6 Backup process for tenant configurations

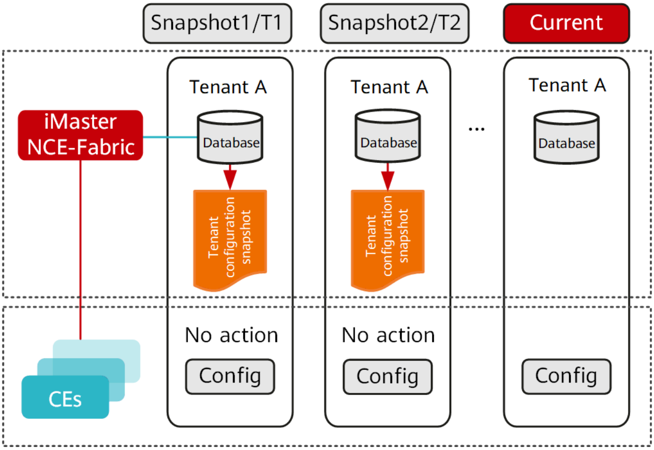

Figure 4-7 Rollback process for tenant configurations

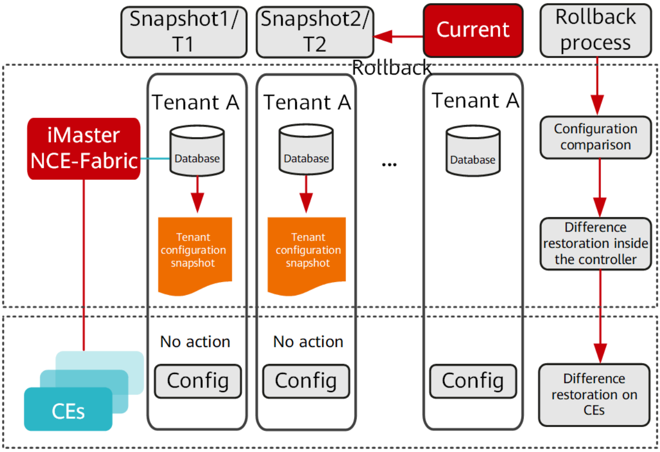

1. **Configuration comparison**: iMaster NCE-Fabric compares the current configuration of tenant A with that at T2 and displays the differences for users to browse.
2. **Difference restoration inside the controller**: If any difference is discovered, iMaster NCE-Fabric invokes the northbound application programming interface (API) to restore the configuration.
3. **Difference restoration on CEs**: iMaster NCE-Fabric performs configuration synchronization and rollback by invoking the southbound device interfaces.

# 4.5 Server Capacity Expansion

## Benefits

Server capacity expansion is a pivotal and routine task on cloud DCNs; however, a heavy workload is required if the conventional approach is used. Specifically, O&M personnel need to plan how the server NICs are connected to switches and set up these connections accordingly to enable the new servers to access the network and go online, as shown in Figure 4-8. This approach is inefficient and time-consuming when numerous servers need to be added, failing to meet customer requirements for rapid service rollout.

Facing this challenge, the entire industry has been seeking an automated and intelligent solution. In response, Huawei's iMaster NCE-Fabric provides intent-driven automated server capacity expansion. With this technique, users only need to enter the parameter settings required for the controller to automatically complete the expansion.

Figure 4-8 Server capacity expansion process in conventional methods

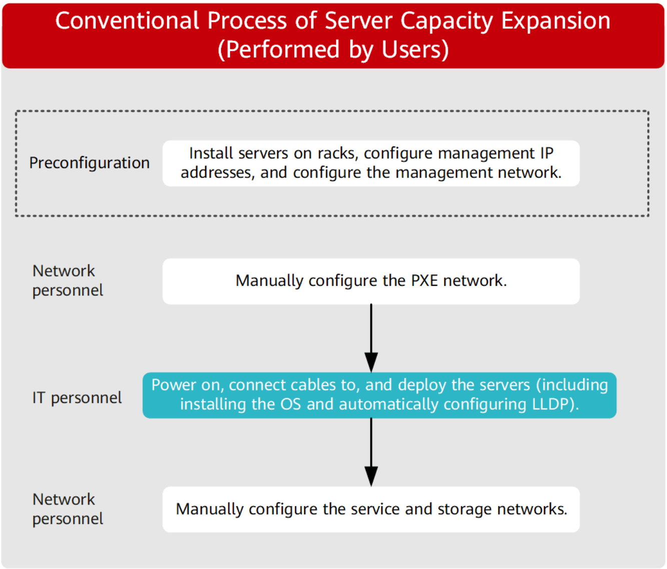
Conventional Process of Server Capacity Expansion (Performed by Users)
Preconfiguration
Install servers on racks, configure management IP addresses, and configure the management network.
Network personnel
Manually configure the PXE network.
IT personnel
Power on, connect cables to, and deploy the servers (including installing the OS and automatically configuring LLDP).
Network personnel
Manually configure the service and storage networks.

## Implementation

Figure 4-9 shows how Huawei's iMaster NCE-Fabric performs server capacity expansion. After completing the pre-configuration, O&M personnel create a server capacity expansion task on the **Server Expansion** tab page of iMaster NCE-Fabric, and specify the VLAN to which servers are to be connected, logical switches, and server NICs. The controller then automatically discovers server links and provisions the network. Users can also enter the common network settings for bringing servers online through one-time orchestration. iMaster NCE-Fabric then automatically identifies the access interfaces of the servers by reading the switches' LLDP information and automatically configures the access interfaces, enabling a plethora of servers to go online rapidly and efficiently.

Figure 4-9 Automated server capacity expansion

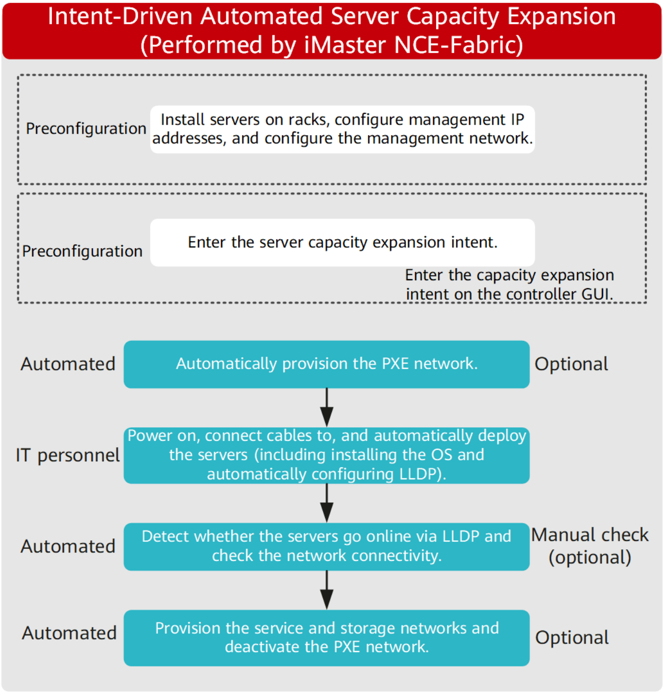
Intent-Driven Automated Server Capacity Expansion (Performed by iMaster NCE-Fabric)

Preconfiguration: Install servers on racks, configure management IP addresses, and configure the management network.

Preconfiguration: Enter the server capacity expansion intent. Enter the capacity expansion intent on the controller GUI.

Automated: Automatically provision the PXE network. Optional

IT personnel: Power on, connect cables to, and automatically deploy the servers (including installing the OS and automatically configuring LLDP).

Automated: Detect whether the servers go online via LLDP and check the network connectivity. Manual check (optional)

Automated: Provision the service and storage networks and deactivate the PXE network. Optional

# 4.6 Service Change Case

## Customer Pain Points

An enterprise has a wide variety of services and needs to change its network every week. During one change, it encountered a service access failure after service migration. It immediately started to troubleshoot the failure and spent over 2 hours rectifying the fault.

# Solution: Network-Wide Configuration Rollback

DCN faults should not cause concerns, but those that persist after long-term troubleshooting are worth worrying about.

Figure 4-10 Network-wide configuration rollback

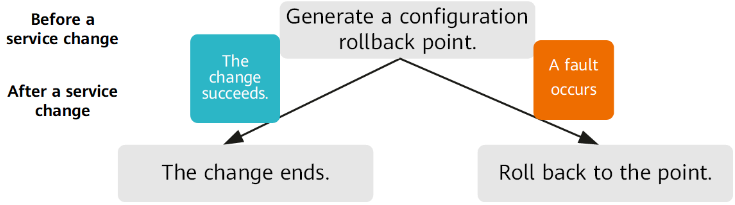

As illustrated in the preceding figure, iMaster NCE-Fabric generates a snapshot — a configuration rollback point — for network-wide configurations before a service change. Assume that a serious network exception occurs after the service change, and device configurations on the entire network need to be rolled back to those at the time point before the change. iMaster NCE-Fabric can complete this complex task in just 10 minutes, compared with several hours if conventional methods are used.

# Chapter 5
# Solution Highlights for Routine PMI

**Abstract**

This chapter describes the highlights of Huawei's Cloud DC Intelligent O&M Solution for coping with routine PMI challenges, followed by an extensive look at use cases across industries. These highlights include comprehensive network health evaluation, proactive device exception prediction, and connectivity check.

## 5.1 Comprehensive Network Health Evaluation

### Benefits

Conventional DCNs are operated and maintained based on alarm events. However, as cloud DCNs continue to grow in scale, so do alarms, gradually going beyond the handling capability of the conventional methods that favor manual operations. One possible compromise is to filter less important alarms and focus only on critical alarms, but this leads to an incomplete grasp of the network health. Not only this, as cutting-edge technologies, such as SDN, gain momentum on DCNs, the DCN architecture becomes increasingly complex. Current DCNs typically consist of a physical (underlay) network and a logical (overlay) network. We cannot operate and maintain such a complex architecture based on alarm events alone.

Facing this, Huawei's iMaster NCE-FabricInsight provides an all-new network health evaluation solution that is ideal for systematic DCN evaluation and monitoring. To elaborate, this solution:

- Models the entire network based on knowledge graphs and builds a five-layer network evaluation system, redefining the conventional network monitoring model, which is built on a single device basis. The layers include device, network, protocol, overlay, and service.
- Adopts telemetry to obtain all-scenario network data in real time, yielding 24/7 intuitive insights into the quality of the entire network.
- Dynamically monitors key network indicators and proactively predicts network capacity and traffic risks.
- Generates network health evaluation reports in real time or periodically, maximizing O&M efficiency and service experience.

# Implementation

The comprehensive network health evaluation function falls into three sub-functions: network-level abstraction and modeling, intelligent network health evaluation, and network fault analysis and remediation.

## Network-Level Abstraction and Modeling

As mentioned above, iMaster NCE-FabricInsight models the entire network based on knowledge graphs and builds a five-layer network evaluation system. As illustrated in Figure 5-1, the network evaluation system consists of five layers — service, overlay, protocol, network, and device layers from top to bottom.

Figure 5-1 Five-layer network evaluation system

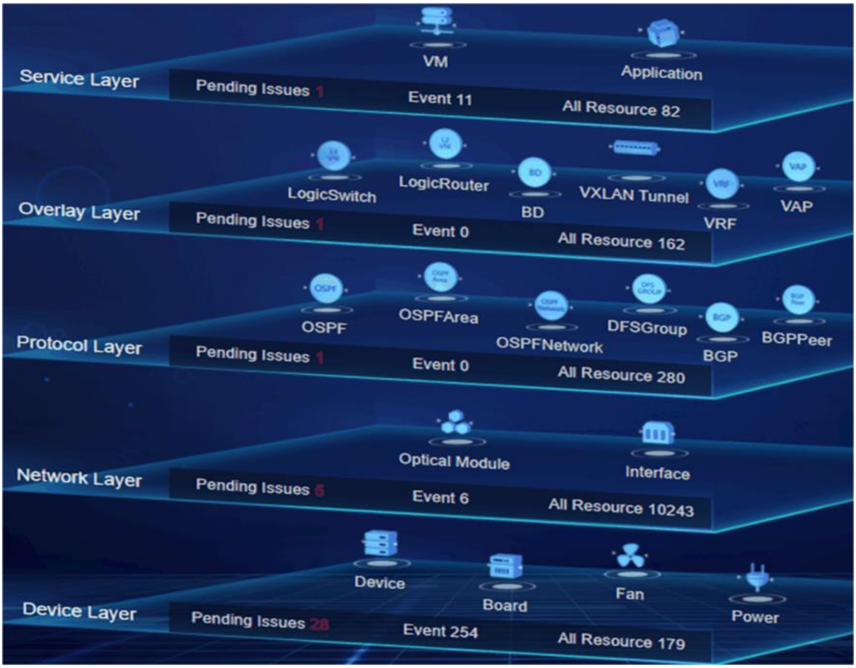

- **Device layer**: is a core element to DC infrastructure. This layer contains all hardware resources, such as boards, power modules, and fan modules on network devices and hosts.
- **Network layer**: connects hardware hosts. This layer includes interconnected interfaces between devices and optical modules.
- **Protocol layer**: provides connectivity and reliability protection for hardware resources. Typical protocols used at this layer include Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and Multichassis Link Aggregation Group (M-LAG).
- **Overlay layer**: houses network services. Take SDN VXLAN as an example. Its VXLAN tunnels, bridge domains (BDs), and virtual access points (VAPs) reside at this layer.

- **Service layer**: carries services once the overlay layer is ready, including VMs on DCNs and applications carried on VMs.

### Intelligent Network Health Evaluation

iMaster NCE-FabricInsight builds a network object model for each layer, comprehensively collects network data, including logs, telemetry KPIs, network device configurations, and service flows between hosts, and uses intelligent analysis algorithms to evaluate the health of each layer. Specifically, this sub-function can accurately identify network issues in the following categories:

- **Performance**: network issues that can deteriorate service quality, such as CPU usage threshold crossing, memory usage threshold crossing, and interface congestion on switches
- **Capacity**: threshold crossing issues for ARP entries, ND entries, and MAC address entries on switches
- **Status**: one-time or repeated exceptions of interface boards or main control boards on switches
- **Policy**: TCP SYN Flood, ARP, and ND attacks
- **Connectivity**: access-side single IP address exceptions, server access exceptions, and TCP interface disablement
- **Intent**: inconsistencies in link interface indicators, routing loops, and routing blackholes

This issue identification is a differentiator that gives O&M personnel a good command over network health. Take optical module management as an example. FabricInsight comprehensively collects optical module data, including the receive and transmit optical power, current, voltage, temperature, and the status of the interfaces where optical modules reside. Not only this, FabricInsight draws on Huawei's years of expertise in the DC domain to identify at-risk optical modules and notifies administrators to replace them before they cause any damage to the network, minimizing service loss.

iMaster NCE-FabricInsight can also periodically generate network health evaluation reports, so that administrators can easily figure out what faults occurred in a specified period, what faults are rectified, and what faults are still ongoing.

### Network Fault Analysis and Remediation

iMaster NCE-FabricInsight can continuously monitor the network health, detect a network fault at the time it occurs, and generate a detailed fault report. This report records the time when the fault occurs, affected objects, the root cause, and possible impacts on existing services, facilitating administrators in further troubleshooting.

# 5.2 Proactive Device Exception Detection

## Benefits

In conventional O&M, we passively respond to network faults and cannot predict or prevent network faults. While this has worked in the past, it cannot meet the increasingly high requirements of users for service quality. For this reason, we need to improve the fault prediction and prevention capabilities so as to minimize the chance of network faults. According to an authoritative survey, about 64% of DC alarms are triggered by packet loss, optical link exceptions, and traffic exceptions, and the routine approach to detecting these alarms is to set static detection thresholds. This, however, leads to many false positives (about 50%) and lots of unnecessary troubleshooting. To improve detection accuracy and efficiency, the industry is looking for ways to detect faults before they cause damage to services.

## Implementation

Huawei's iMaster NCE-FabricInsight innovatively adopts machine learning (ML) algorithms to detect changes in network behavior changes that might indicate oncoming faults. To elaborate, FabricInsight uses a Gaussian process regression algorithm and historical network data to automatically learn and update the dynamic baselines of KPIs for devices, boards, interfaces, and optical modules, so as to intelligently detect network performance exceptions and minimize the impact of faults on services. Furthering to this, FabricInsight can accurately detect sub-healthy optical links and predict traffic volumes on interfaces through a neural network algorithm, detecting potential faults before they can affect services.

# 5.3 Connectivity Check

## Benefits

As DCNs continue to grow in scale and complexity, network services proliferate, making it increasingly difficult to verify services with high efficacy without sacrificing accuracy.

We usually use the **ping** and **traceroute** commands to verify cloud DCN connectivity and traffic forwarding conditions. This, however, is inefficient and can hardly cover all the forwarding paths. For this reason, efficient and automatic service verification becomes a must-have, and this is where single-path and multi-path detection functions come in.

## Implementation

**Single-Path Detection**

Single-path detection enables users to view the physical paths of service flows between two VMs on the fabric network and check whether service flows are interrupted. To be specific, this function can detect the forwarding paths of overlay packets between VMs, physical servers, and devices, as well as between VMs and physical machines (PMs), facilitating fault locating and diagnosis and guaranteeing high network planning accuracy.

Single-path detection includes the following steps:

1. Encapsulate the path detection packet (packet-out message).

   iMaster NCE-Fabric encapsulates a packet-out message based on the 5-tuple information (source IP address, destination IP address, source interface, destination interface, and protocol) set by users and sends the message to the source CE through OpenFlow. The source CE then forwards the packet-out message to devices alongside the path to be detected.

2. Parse the packet-in messages.

   Devices with single-path detection enabled deliver Access Control List (ACL) rules to match packet-out messages in the inbound direction. If a message matches an ACL rule, the device reports a packet-in message to the controller. The controller then parses the message to obtain the device information as well as the inbound and outbound interfaces of the messages, and saves the information to a specified database.

3. Calculate the forwarding path.

   iMaster NCE-Fabric calculates the entire forwarding path based on the saved information.

4. Display the path.

   As shown in Figure 5-2, iMaster NCE-Fabric displays the forwarding path as well as the traveling path of the packet-out message in the network topology.

Figure 5-2 Single-path detection

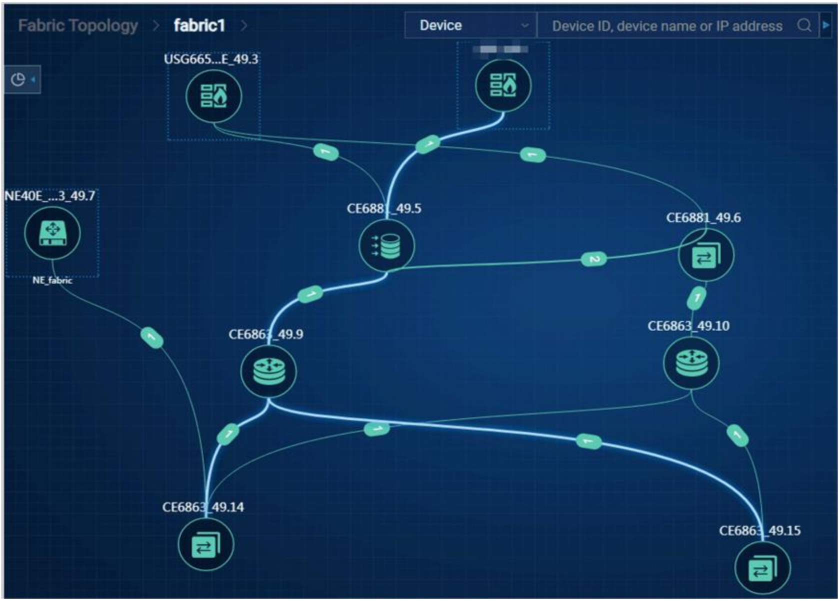

## Multi-Path Detection

Multi-path detection takes a much broader approach, detecting all the possible forwarding paths between two VXLAN tunnel endpoints (VTEPs), and is typically used for network detection and verification. This is done in a similar way to single-path detection. The main difference between them is that users need to specify the number of packet-out messages in the multi-path detection mode. Based on this number, iMaster NCE-Fabric encapsulates packet-out messages with different protocol interface numbers to detect as many forwarding paths as possible.

This function needs to work with switches. To elaborate, iMaster NCE-Fabric sends a packet-out message to a switch. The switch then calculates the outbound interface based on the MAC address table or routing table, and reports the outbound interface, inbound interface, and packet-in message to iMaster NCE-Fabric. The controller then calculates the traffic forwarding paths based on the received information. With these paths, O&M personnel can easily determine the point of failure if a network fault occurs.

# 5.4 Routine PMI Case

## Customer Pain Points

An insurance company operates and maintains a DCN that consists of 19 Performance-Optimized Datacenters (PODs), six of which are SDN networks while the remaining are conventional networks. This DCN houses more than 1000 devices in total. Such a large and complex network already goes beyond the handling capabilities of the company's O&M team. This team cannot monitor and evaluate this network in real time, and still follows their old work habits — manually collecting network KPIs and outputting O&M reports periodically, which is time-consuming and labor-intensive.

## Comprehensive Network Health Evaluation, Simplifying O&M

As illustrated in Figure 5-3, FabricInsight provides comprehensive network health evaluation, which enables 24/7 big data analytics, O&M, and management for DCNs, helping the company rapidly detect network or application issues before they cause any damage to the network.

Figure 5-3 Comprehensive network health evaluation

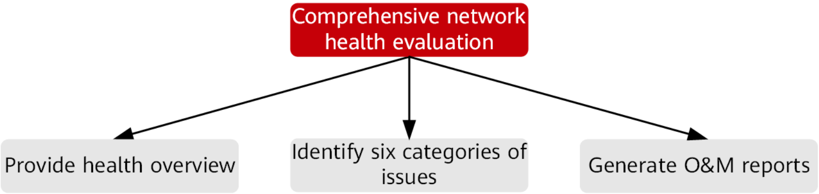
Comprehensive network health evaluation
Provide health overview
Identify six categories of issues
Generate O&M reports

FabricInsight can well meet the insurance company's O&M needs. Specifically, the analyzer can:

- Detect sub-healthy optical modules and provide optimization suggestions to guide the company to replace these optical modules and their matching optical fibers, thereby eliminating potential risks.
- Detect hardware issues, such as insufficient power supply to main control boards caused by misoperation, and help the company rapidly resolve these issues, ensuring continuous and stable network operation.
- Automatically generate O&M reports every week and allow O&M personnel to download these reports with just one click.

# Chapter 6
# Solution Highlights for Emergency Recovery

**Abstract**

This chapter describes the field-proven benefits of Huawei's Cloud DC Intelligent O&M Solution in emergency recovery, including intelligent fault remediation and its use cases across industries.

## 6.1 Intelligent Fault Remediation

### Benefits

In routine maintenance of cloud data center networks, it is important to quickly detect, locate, and rectify network faults. Traditionally, network faults are detected using one of two methods:

**Method 1**: Faults are detected through alarms and logs collected by the Network Management System (NMS) as well as statistical data reported by devices. This method has many issues, such as:

- Inefficient: A delay generally occurs when the NMS collects device data. It also takes some time for administrators to pinpoint data relating to faults such as in alarms on the NMS. In some cases, administrators may not pay enough attention to or promptly handle early fault symptoms.
- Detecting complex faults depends heavily on administrator skills. For this reason, faults are generally determined only after administrators have comprehensively analyzed multiple types of NMS data and metrics.
- There is no effective method for detecting and locating traffic forwarding exceptions caused by device algorithms or underlying chip faults. In this case, vendors often need to dispatch technical professionals to locate faults onsite.

**Method 2**: Faults are detected from user service complaints. During conventional data center network O&M, many faults — such as incorrect device configurations, abnormal forwarding entries, and service exceptions caused by attacks — cannot be detected before they reach users. Despite the logs and statistical data collected by the NMS, these faults are left unnoticed by administrators until users complain about services. Moreover, troubleshooting these faults is time- and labor-consuming.

Faults in cloud data center networks need to be detected, located, and rectified quickly to keep up with the rapid service provisioning and change these networks entail. To this end, the O&M system — in addition to continuing to collect log and alarm information — also needs to collect more data of other kinds, including indicators, resources, table entries, and even session interactions. Furthermore, the O&M system must excel at analyzing, processing, and gaining insightful fault correlations from massive amounts of data. In this way, faults can be located quickly and accurately. The O&M system must also be capable of automatically generating recovery plans for faults that can be recovered or isolated from through parameter settings. These plans can be delivered with one click as needed to quickly recover or isolate faults.

That's where Huawei iMaster NCE-FabricInsight and iMaster NCE-Fabric come in, collaborating to create an intelligent fault remediation system. Specifically, iMaster NCE-FabricInsight draws on a rule engine, intelligence engine, and knowledge graph to perform big data mining and analytics, thereby detecting and locating more than 75 typical data center network faults in minimal time. It can also work with iMaster NCE-Fabric to implement fault recovery or isolation with just one click. iMaster NCE-FabricInsight can present a network and service impact analysis for a given fault. Similarly, iMaster NCE-Fabric can display the network and service impacts that will result before a fault recovery or isolation plan is delivered. This greatly facilitates decision making.

# Implementation

iMaster NCE-FabricInsight uses technologies such as telemetry and NetStream to efficiently monitor cloud data center networks and obtain TCP packets, CPU usage, ARP table entries, Forwarding Information Base (FIB) table entries, and other such information from devices. With this information, iMaster NCE-FabricInsight can detect faults using one of four ways:

- **Detecting faults based on TCP flow exceptions**: iMaster NCE-FabricInsight captures TCP flag packets, analyzes TCP sessions, detects faults from link establishment exceptions and pinpoints faults using the node that reported the exception, ultimately locating root causes using the rule engine.
- **Detecting faults based on alarm logs**: After some faults occur, network devices generate alarms and report such alarms to the NMS or other log collection systems. Conversely, some other faults are triggered upon the receipt of device alarm logs.
- **Detecting faults based on periodic data sampling for network monitoring objects**: In addition to collecting alarm logs from devices, iMaster NCE-FabricInsight periodically samples data from specific objects on devices through Telemetry. This includes statistics about packets sent and received on device interfaces, optical module metrics, and packet loss statistics. By comparing the periodically sampled data of all monitored objects, iMaster NCE-FabricInsight can easily detect exceptions and report faults, for example, faults caused by excessively low transmit/receive power of optical modules.
- **Detecting faults by periodically inspecting network connectivity**: iMaster NCE-FabricInsight detects faults caused by abnormal network reachability by periodically inspecting the connectivity of network monitoring objects. For example, iMaster NCE-FabricInsight easily detects device management channel interruptions.

On top of that, iMaster NCE-FabricInsight innovates with a big data analytics engine to analyze the collected network data, locate faults, and provide root causes. NCE-FabricInsight takes into account the different strengths and weaknesses of the abovementioned fault detection methods, using differentiated fault locating algorithms to improve the accuracy of root cause identification.

After iMaster NCE-FabricInsight detects and locates a fault, iMaster NCE-Fabric determines whether it is feasible to recover the fault through configuration methods. If so, iMaster NCE-Fabric offers a fault recovery plan. When a user selects a fault recovery plan on iMaster NCE-Fabric's fault event management page, it describes the plan and displays the related configuration information to be delivered to the devices. iMaster NCE-Fabric also provides after-the-fact impact analysis of the fault recovery plan, so that users can determine whether the fault recovery plan is suitable or not.

Figure 6-1 Fault isolation plan impact analysis provided by iMaster NCE-Fabric

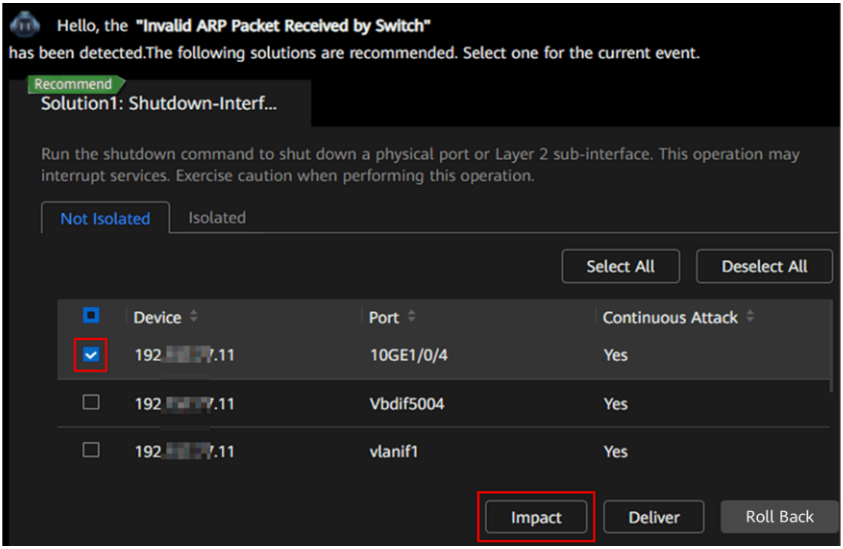

Hello, the "**Invalid ARP Packet Received by Switch**" has been detected.The following solutions are recommended. Select one for the current event.

Solution1: Shutdown-Interf...

Run the shutdown command to shut down a physical port or Layer 2 sub-interface. This operation may interrupt services. Exercise caution when performing this operation.

| ■ | Device | Port | Continuous Attack |
|---|---|---|---|
| ☑ | 192[redacted]7.11 | 10GE1/0/4 | Yes |
| □ | 192[redacted]7.11 | Vbdif5004 | Yes |
| □ | 192[redacted]7.11 | vlanif1 | Yes |

Figure 6-2 Configuration to be delivered for the fault isolation plan

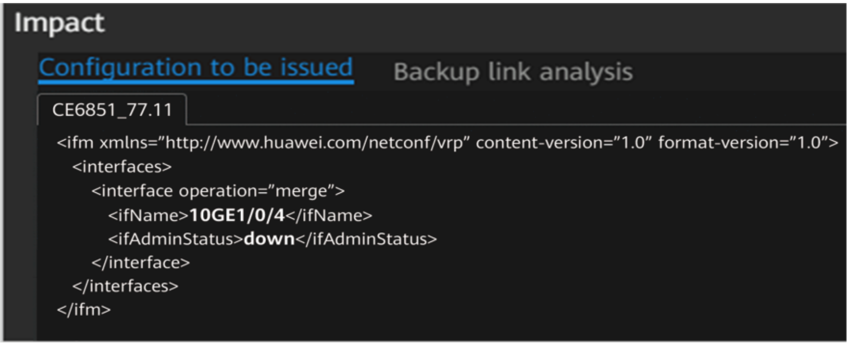

```
<ifm xmlns="http://www.huawei.com/netconf/vrp" content-version="1.0" format-version="1.0">
  <interfaces>
    <interface operation="merge">
      <ifName>10GE1/0/4</ifName>
      <ifAdminStatus>down</ifAdminStatus>
    </interface>
  </interfaces>
</ifm>
```

Figure 6-3 Analysis on remaining backup links after fault isolation

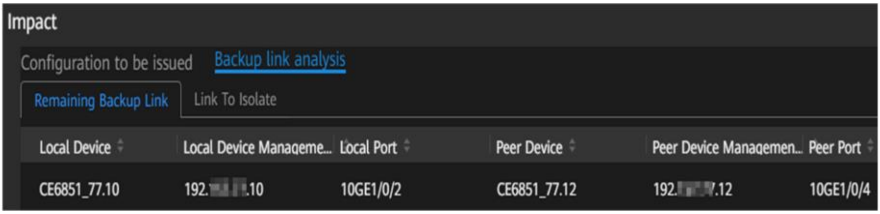

| Local Device | Local Device Manageme... | Local Port | Peer Device | Peer Device Managemen... | Peer Port |
|---|---|---|---|---|---|
| CE6851_77.10 | 192.[redacted].10 | 10GE1/0/2 | CE6851_77.12 | 192.[redacted].12 | 10GE1/0/4 |

As shown in the preceding three figures, iMaster NCE-Fabric provides one or more handling suggestions, covering operational objects and impact analysis. With a full understanding of these suggestions and the network conditions, O&M personnel can click to deliver the fault recovery plan for intelligent fault remediation.

# 6.2 Emergency Recovery Case

## Customer Pain Points

A bank has a large data center, and has already provisioned more than 25 mobile cloud services based on a distributed VXLAN network overlay, 10 software-only cloud services, and 5 traditional Layer 2 cloud services. The bank also rolls out more than 5 new cloud services every year. Facing this, the bank struggles with a heavy O&M workload, let alone the traditional O&M methods, which cannot meet the bank's increasing O&M requirements.

## Intelligent Fault Drills for "1-3-5" Troubleshooting

A fault drill is performed on the live network, where both a traditional NMS and iMaster NCE-FabricInsight are used for troubleshooting. During the fault drill, fault cases are selected at random to test and verify the troubleshooting performance of the traditional NMS and iMaster NCE-FabricInsight. The final results show that iMaster NCE-FabricInsight outperforms traditional NMS in all the eight of the major test categories.

Huawei's Cloud DC Intelligent O&M Solution stands out with features such as big data analysis, intelligent O&M, and real-time insights into fabric status and application behavior. With these traits, this feature-rich solution helps customers quickly detect network and application problems and ensure that applications run stably.

Chapter 7

# Solution Highlights for Root Cause Locating

**Abstract**

This chapter takes a close look at the highlights of root cause locating in Huawei's Cloud DC Intelligent O&M Solution and some of their typical use cases.

## 7.1 Root Cause Locating

### Benefits

To provide high reliability and bandwidth, DCNs are typically designed to forward traffic in the ECMP mode, in which traffic is steered via the hash algorithm and the traffic forwarding paths increase exponentially as the number of network nodes increases. With so many forwarding paths, however, it is hard for administrators to select the optimal one.

On DCNs, the protocol parsing and traffic forwarding modes of different types of servers and NICs vary. Moreover, there are many other types of devices on the network, such as firewalls and LBs, which may come from different vendors and work in different ways, making it difficult to rapidly locate and rectify faults.

In addition, fault locating heavily depends on the experience of O&M personnel and is extremely time-consuming. Such an approach is no longer feasible as DCNs grow in scale and configuration complexity.

In response, Huawei's iMaster NCE-FabricInsight provides a big data analytics engine that can collect and analyze network data so as to rapidly locate the root causes of faults.

# Implementation

iMaster NCE-FabricInsight can locate the root causes of faults based on abnormal TCP flows, alarms, periodically collected data samples of network objects, and network connectivity check results. This significantly improves the accuracy and efficiency of root cause locating.

### Root Cause Locating Based on Abnormal TCP Flows

As illustrated in Figure 7-1, iMaster NCE-FabricInsight monitors the TCP session status in real time by analyzing the TCP packets reported by devices. To locate the root cause of TCP connection exceptions, the analyzer:

- Uses an intelligence engine to analyze and locate the affected TCP flows.
- Adopts a knowledge inference engine to locate the faulty device based on the location where the exception occurs and compare the network data before and at the time point when the exception occurs to figure out the root cause.

Figure 7-1 Logic for locating root causes of traffic exceptions

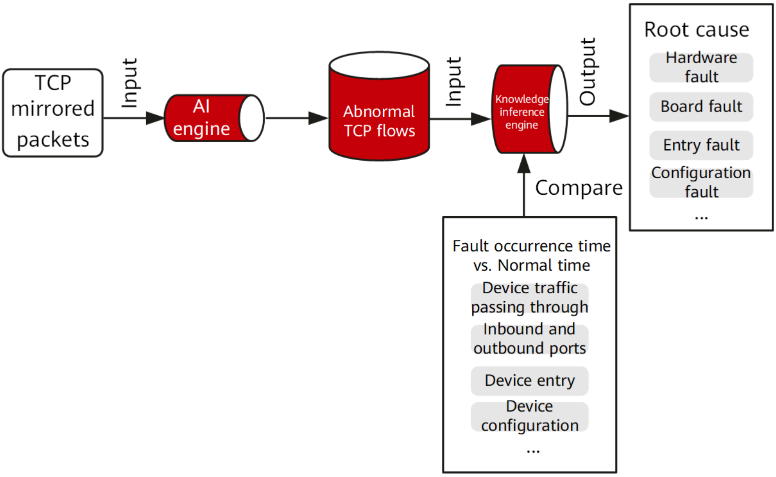

### Root Cause Locating Based on Alarms

Upon receiving an alarm from a network device, iMaster NCE-FabricInsight checks whether the alarm itself reflects the root cause and locates the fault accordingly.

- If the alarm itself reflects the root cause, iMaster NCE-FabricInsight feeds back the root cause according to the alarm. For example, when receiving a device resource alarm, it directly determines that the root cause is resource insufficiency, as this is the most common cause of resource alarms. Typical examples of such alarms include threshold crossing alarms such as for CPU usage, memory usage, and entry.
- If the alarm itself does not reflect the root cause (that is, the fault reflected by the alarm is not the root fault), FabricInsight immediately traces back through the source based on knowledge graphs until the root fault is found, and then feeds back the root cause.

Figure 7-2 Knowledge graph-based inference engine

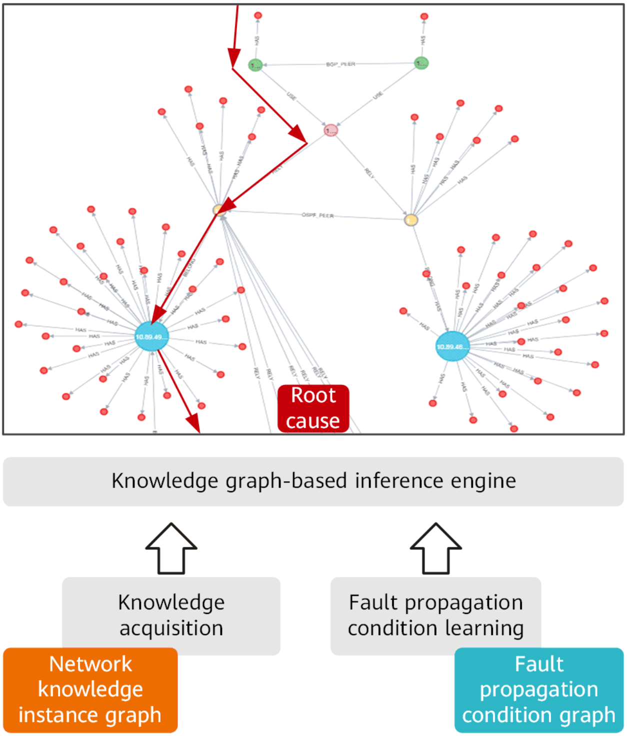
BGP_PEER
USE
RELY
OSPF_PEER
HAS
10.89.49...
10.89.46...
Root cause
Knowledge graph-based inference engine
Knowledge acquisition
Fault propagation condition learning
Network knowledge instance graph
Fault propagation condition graph

The question is how can Huawei's iMaster NCE-FabricInsight find the root cause so easily when the fault directly attributed to an alarm is not the root fault? The answer lies in its knowledge graph-based inference engine. As illustrated in Figure 7-2, this engine builds fault propagation between network objects and models fault knowledge so as to determine the dependencies between network objects.

Take a BGP peer session alarm that is actually caused by a port fault as an example. Figure 7-3 shows how FabricInsight locates the root cause of this alarm. The entire process can fall in five simple steps:

1. iMaster NCE-FabricInsight uses knowledge graphs to find the BGP process bound to the BGP peer, and the OSPF routing processes (OSPF 1, 2, and 3) that bear the BGP process on the underlay network.
2. The analyzer checks the status of OSPF peers associated with the three OSPF routing processes, and finds that the status of OSPF peer 1 associated with OSPF 1 has changed.
3. The analyzer continues to trace back through the source and finds that the status of the L3 port where OSPF peer 1 resides is abnormal.
4. The analyzer digs further, ultimately discovering that a link-down event (Link1 Down) occurred on the L2 port that forwards packets to OSPF peer 1.
5. The analyzer feeds back "Link1 going-Down" as the root cause of this alarm.

Figure 7-3 Locating the root cause of a BGP peer session alarm that is actually caused by a port fault

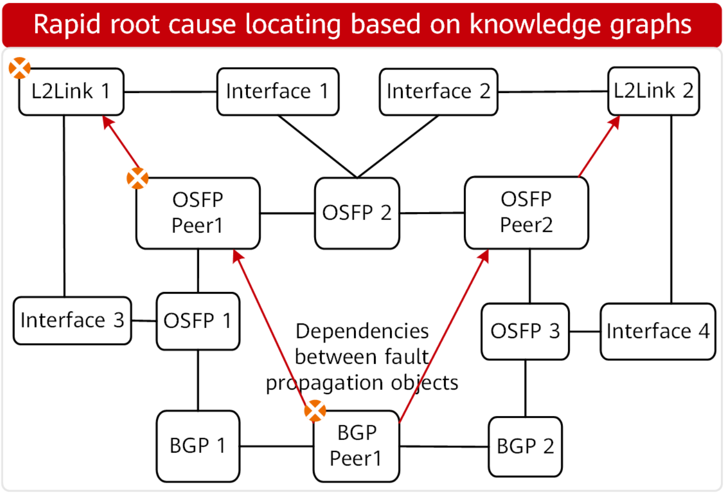

### Root Cause Locating Based on the Periodically Collected Data Samples of Network Objects

iMaster NCE-FabricInsight periodically collects data samples of network objects via telemetry, uses a big data analytics engine to analyze these samples to figure out if there are any exceptions among the network objects, and provides the root cause once an exception is detected.

Take an optical link fault as an example. iMaster NCE-FabricInsight periodically collects data samples of optical modules, including the optical modules' temperature, voltage, current, and optical power. Upon discovering that a parameter value has gone outside of the normal value range, FabricInsight immediately reports an optical link fault, and displays the abnormal parameter and historical data trends of the faulty optical module.

**Root Cause Locating Based on Network Connectivity Check Results**

iMaster NCE-FabricInsight can also locate the root cause via network connectivity check. To elaborate, it checks the connectivity between target objects through ping or OpenFlow, and determines whether a fault has occurred based on the check results.

For instance, iMaster NCE-FabricInsight periodically pings the management IP address of each managed device to figure out if there is a device disconnected from the system. The root cause of any device disconnection can be located easily in this manner.

# 7.2 Root Cause Locating Case

## Customer Pain Points

To support the development of their business, a stock exchange customer requires to build a comprehensive intelligent O&M platform that can perform routine network PMI and rapid fault remediation for its DCNs.

## Rapid Root Cause Locating, Minimizing Service Interruption

Intelligent O&M Solution is ideal for this task. It can rapidly and accurately locate the root cause of faults so as to quickly restore services and minimize service interruption, thereby guaranteeing high service reliability and continuity.

The customer in this case found that its bastion host did not respond and could not be used if no operation was performed on the host for 2 hours.

FabricInsight analyzed the issue, and found that the host received an RST packet 2 hours after it successfully established a link. After checking the forwarding path of the packet, the analyzer discovered that the host traffic traveled through the firewall between south and north, but the firewall did not return any SYN packet after receiving the RST packet.

Through further digging, the analyzer found that a session on the firewall aged out when no operation was performed on the host for a long time.

# Chapter 8
# Evolution Trends of the Intelligent O&M Solution

**Abstract**

This chapter introduces the evolution trends of cloud DCN intelligent O&M.

## Network + IT Converged Intelligent O&M

In conventional O&M, the network and IT teams have distinct responsibilities that are clearly defined. Their responsibilities, however, are increasingly overlapping as DCs move to the cloud. Nowadays, a network fault cannot be located by only one of the network team or IT team. Rather, it is vital that they collaborate. To this end, the intelligent O&M system must support centralized management of network and IT resources.

## O&M + Service Converged System with Full-Lifecycle Management and Control

In the future, different service deployment and O&M functions — such as network PMI, intelligent flow analysis, and automated fault locating — will be integrated as apps in the intelligent O&M system that is converged with the service system. In this manner, the intelligent O&M system can implement end-to-end (E2E) O&M by not only operating and maintaining networks, but also centrally managing service deployment information — for example, reading device configuration information, analyzing configuration changes, and synchronizing PM, VM, and container information. The system can also automatically deliver or delete the configurations of advanced O&M functions without manual intervention.

Such convergence enables full-lifecycle automated service management that spans service rollout, monitoring, upgrade, change, capacity expansion, and going-offline, maximizing service deployment and O&M efficiency. In addition, the logs generated throughout this process can be used as a reference for fault analysis.

## Intelligent Fault Prevention, Detection, and Self-Healing

The entire industry is now attempting to adopt cutting-edge technologies, such as big data correlation analysis and ML, to improve the automation and intelligence levels of the O&M system, with the ultimate goal of providing intelligent network assurance capabilities spanning fault prevention, detection, and self-healing.

- **Proactive Fault Prevention**

  Fault prevention is actually more important than fault remediation, yet it tends to be neglected. Typically, there are two aspects to fault prevention:

  **Zero manual configuration**: As indicated by Huawei's extensive practices, more than 50% of DC faults are caused by manual operation errors, such as incorrect configuration or planning. A minor mistake, however, can cause a serious network accident in DCs, which both store and process masses of data. For this reason, we need to figure out how to implement fully automated network management and control throughout the lifecycle.

**Proactive O&M and analysis**: Drawing on big data technology and a powerful fault pattern library that can be flexibly upgraded, the O&M system can perform correlation analysis across data zones to predict potential risks. Furthering to this, the system can work with the service execution system to rectify network faults before users are aware, preventing services from being affected.

- **Agile Fault Detection**

  Proactive fault prevention cannot prevent 100% of network faults, which is why the O&M system must be able to rapidly discover faults as soon as they occur. To achieve this, we need to build a comprehensive network monitoring system that can detect faults regarding infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS).

- **Intelligent Fault Locating**

  In the cloud era, the distributed and microservice-based software architecture is a common choice. This, however, makes service invoking relationships increasingly complex, hindering fault locating. In fact, the performance of a network O&M system heavily depends on its fault locating efficiency and accuracy. Specifically, a high-quality O&M system should be able to:

  - Invoke data of the monitoring points at each phase of a service flow, thereby tracing the root causes of faults based on the associated alarms and exception information.
  - Draw on a powerful big data-powered expert diagnosis model to determine the root causes and provide optimal troubleshooting methods.

- **Automated Fault Remediation**

  As cloud DCs continue to grow in scale, network faults proliferate. Huawei's years of experience in DC O&M shows that a large cloud DC can encounter more than 1000 faults. If the DC cannot automatically classify or handle so many faults, user experience will be seriously affected. Facing this, we need to find ways to enable the O&M system to rapidly identify common faults and automatically provide and enforce optimal policies to rectify the faults with no manual intervention. In addition, O&M personnel need to be aware of the entire fault handling process so that they can carry out further checks and analysis.

IP Network
E-Books Series

**Contact Us**

networkinfo@huawei.com

**More IP Network eBooks**

https://e.huawei.com/en/solutions/enterprise-networks/ip-ebook