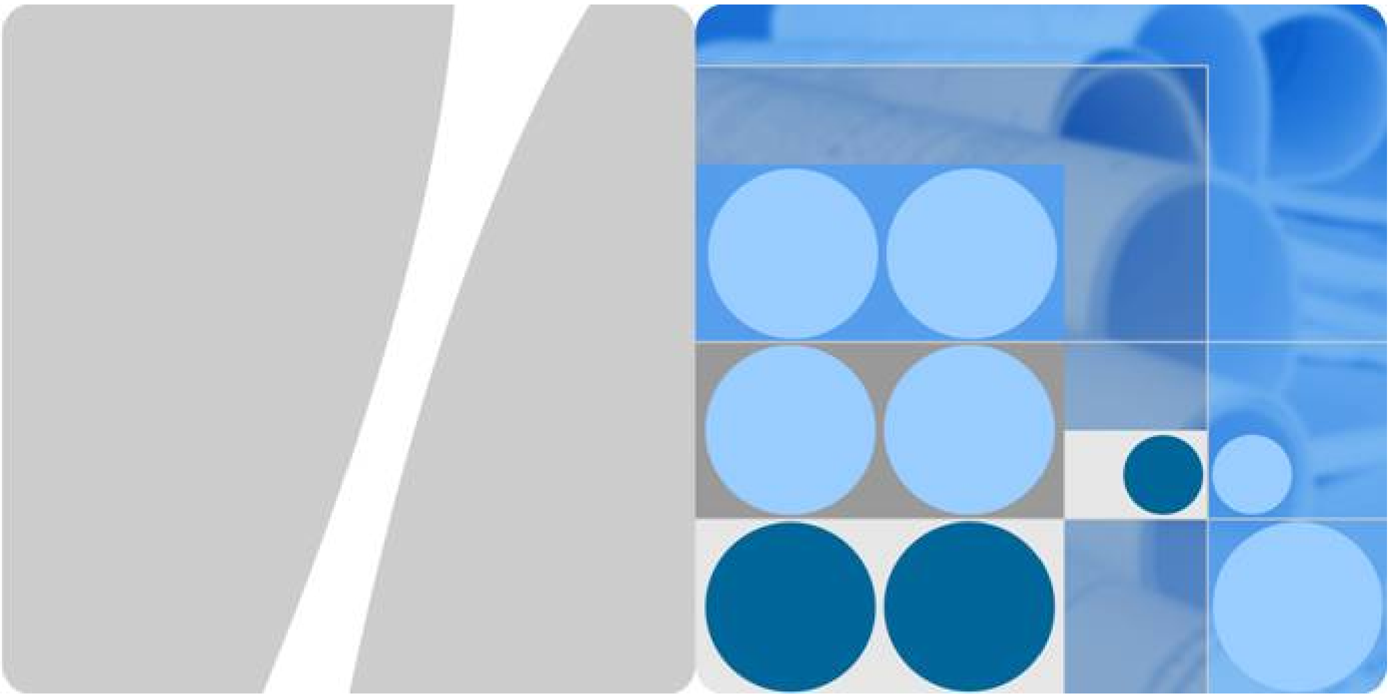

FabricInsight

Product White Paper

Issue 01

Date 2018-12-12

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

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.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://e.huawei.com

Contents

1 Product Overview

Challenges to Traditional Data Center O&M

With the acceleration of digital transformation in the industry, more and more services and applications are deployed in data centers. In addition, the development of software technologies such as Big Data, machine learning, distribution, and servitization accelerates the pace of digital transformation in the industry. Cloudification of enterprise data centers becomes increasingly urgent, and cloud computing is becoming the basic capability of each industry. It is an urgent task for enterprises to quickly build cloud-based data centers that can support future service development. Data center networks, as the cornerstone of constructing cloud data centers, are facing great challenges. Traditional data center networks can hardly be cloudified. To handle this problem, the SDN is developed.

In the SDN era, computing resource pooling, storage resource pooling, network resource pooling, and network and service automation bring convenience to users but great challenges to network O&M. Compared with traditional network O&M, network O&M in the SDN era features in the following: proactive, real-time, and large-scale.

- **Proactive**: The SDN scenario requires that services can be provisioned quickly and dynamically. For example, if logical networks are created and deleted as required, network or service configuration changes frequently. Frequent configuration change increases the fault probability. The O&M system must be able to proactively and intelligently detect these faults, and use big data analysis and experience databases to help users quickly locate and rectify faults.
- **Real-time**: The O&M system can detect microburst exceptions on the network in a timely manner. For example, an enterprise customer complained that its lightweight network had the issue of transient packet loss and suspected that there were millisecond-level traffic bursts. However, these issues cannot be detected in the minute-level SNMP mechanism, let alone be optimized.
- **Large-scale**: Large-scale management has many meanings. On one hand, managed objects are extended from physical devices to virtual machines (VMs) and the NE management scale is increased by dozens of times. On the other hand, the device indicator collection granularity is improved from minutes to milliseconds to meet real-time analysis requirements, and the data volume is increased by nearly 1000 times. For active awareness and troubleshooting of issues, FabricInsight needs to collect and analyze network device indicators, and analyze the actual forwarding service flows. To improve the service flow analysis accuracy, the collection must cover all paths on the entire network.

The traditional O&M management system is challenged by the preceding three features of SDN network O&M. According to a survey conducted by the EMA on over 100 enterprises, about 70% of customers are concerned about whether the existing network O&M system is applicable to the SDN scenarios.

To deal with the O&M challenges (proactive, real-time, and large-scale) in the SDN scenario, the customer needs to change the overall O&M architecture so that the SND network can be easily used. Huawei FabricInsight, an intelligent network analysis platform, overrides the traditional monitoring focusing on resource status, detects fabric status and application behavior in real time, and breaks the boundaries between networks and applications. In addition, FabricInsight analyzes networks from the application perspective, proactively detects network or application issues, and provides automatic troubleshooting capabilities for service connectivity issues, helping users quickly demarcate and rectify faults and ensure continuous and stable running of applications.

Figure 1-1 FabricInsight

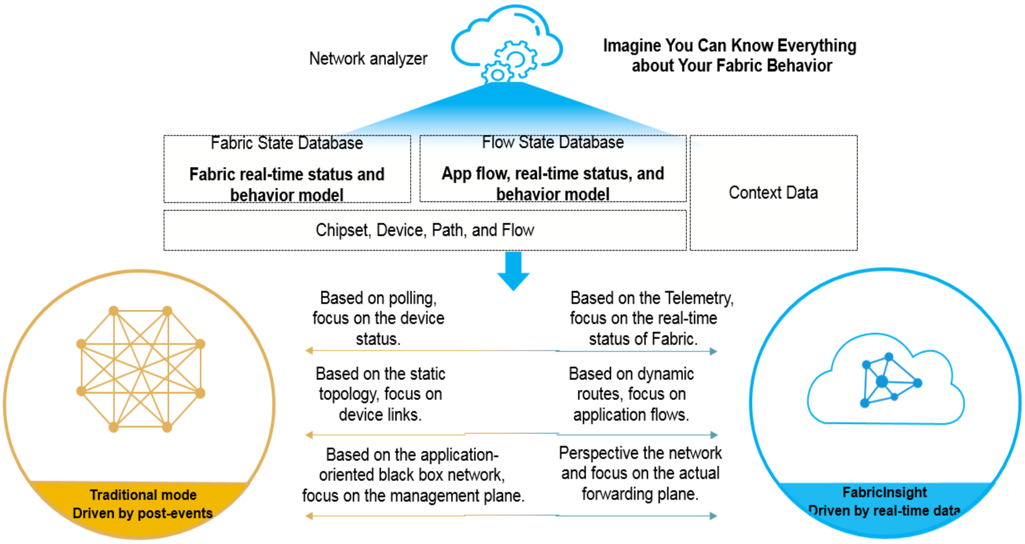
Network analyzer
Imagine You Can Know Everything about Your Fabric Behavior
Fabric State Database
Fabric real-time status and behavior model
Flow State Database
App flow, real-time status, and behavior model
Context Data
Chipset, Device, Path, and Flow
Based on polling, focus on the device status.
Based on the Telemetry, focus on the real-time status of Fabric.
Based on the static topology, focus on device links.
Based on dynamic routes, focus on application flows.
Based on the application-oriented black box network, focus on the management plane.
Perspective the network and focus on the actual forwarding plane.
Traditional mode
Driven by post-events
FabricInsight
Driven by real-time data

The FabricInsight O&M architecture is constructed based on the following points:

- **Visualization: visible and clear**

  The concept of "visible" consists of two aspects: observed objects and real-time observation. Observed objects include physical objects such as devices, interfaces, and links and logical objects such as packet forwarding path, service interaction relationship, and service interaction quality. Real-time observation supports perception of millisecond-level symptoms, for example, identifying microburst traffic congestion on the network. The concept "clear" refers to the observation accuracy. On one hand, a large amount of data needs to be collected. On the other hand, the data must be analyzed in real time to identify abnormal service flows.

- **Automation: proactive analysis and automatic troubleshooting**

  To proactively and intelligently detect issues on the network in a timely manner, the O&M system must be able to analyze massive data and identify abnormal events on the network, for example, service connectivity issues and traffic congestion ports. In addition, the O&M system needs to determine whether to generate issue models and recommend them to users based on machine learning algorithms. For automatic troubleshooting, the O&M system must be able to analyze issue data and learn the issue case library. In addition, the O&M system must be able to orchestrate executable troubleshooting task links for different issue patterns, reducing the time required for issue demarcation and locating from days or hours to minutes.

Solution Design

FabricInsight collects and analyzes the original TCP feature packets forwarded on the network, displays the application interaction relationship and quality, and visualizes the network traffic. In addition, FabricInsight parses packet features, and restores hop-by-hop forwarding paths of packets and forwarding traffic and latency of links to implement association between applications and networks. Then, FabricInsight collects the packet loss, traffic, and configuration of network devices through technologies such as Telemetry and proactively evaluates the network service status based on AI algorithms such as dynamic baseline and Gaussian regression. In this way, FabricInsight can build the multi-layer association analysis capability from service flow to forwarding path to network service, and display application behavior and network quality in a structured manner.

Figure 1-2 FabricInsight solution design

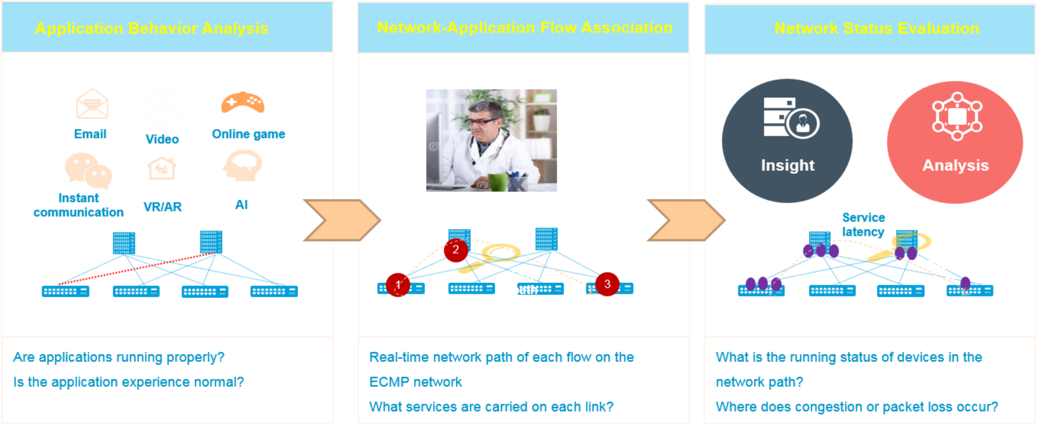
Application Behavior Analysis
Email
Video
Online game
Instant communication
VR/AR
AI
Are applications running properly?
Is the application experience normal?
Network-Application Flow Association
1
2
3
Real-time network path of each flow on the ECMP network
What services are carried on each link?
Network Status Evaluation
Insight
Analysis
Service latency
What is the running status of devices in the network path?
Where does congestion or packet loss occur?

FabricInsight performs big data analysis on collected ERSPAN flows and Telemetry performance metrics through distributed real-time and offline computing. In addition, FabricInsight proactively detects possible issues on the fabric based on AI algorithms such as baseline exception detection and multi-dimension clustering analysis, and intelligently analyzes and identifies whether the network or application has group issues. For service connectivity issues, FabricInsight automatically orchestrates troubleshooting procedures to support one-click automatic troubleshooting. All these help users achieve the proactive and intelligent O&M goal for proactive issue detection and minute-level issue locating and demarcation.

Figure 1-3 Proactive and intelligent O&M

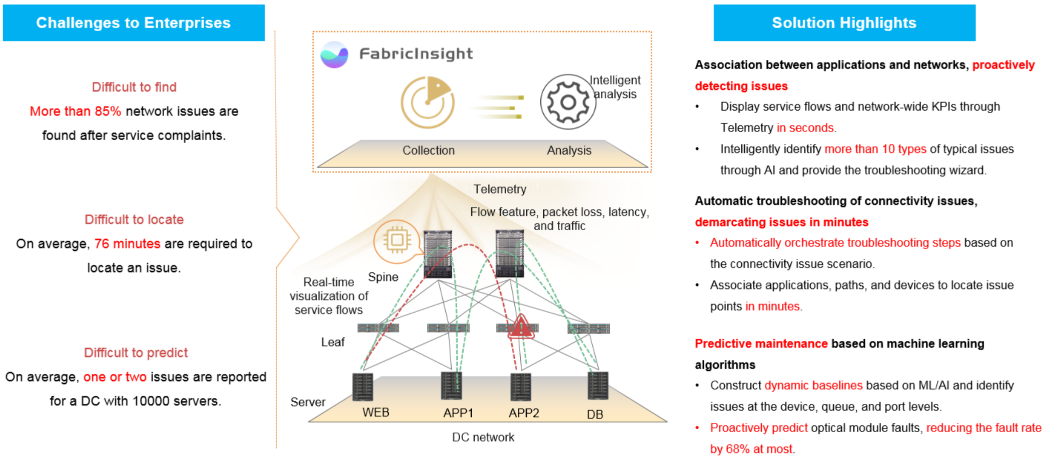

Challenges to Enterprises

Difficult to find

More than 85% network issues are found after service complaints.

Difficult to locate

On average, 76 minutes are required to locate an issue.

Difficult to predict

On average, one or two issues are reported for a DC with 10000 servers.

Solution Highlights

Association between applications and networks, proactively detecting issues

- Display service flows and network-wide KPIs through Telemetry in seconds.
- Intelligently identify more than 10 types of typical issues through AI and provide the troubleshooting wizard.

Automatic troubleshooting of connectivity issues, demarcating issues in minutes

- Automatically orchestrate troubleshooting steps based on the connectivity issue scenario.
- Associate applications, paths, and devices to locate issue points in minutes.

Predictive maintenance based on machine learning algorithms

- Construct dynamic baselines based on ML/AI and identify issues at the device, queue, and port levels.
- Proactively predict optical module faults, reducing the fault rate by 68% at most.

2 Technical Principles

FabricInsight Architecture

Based on Huawei Big Data platform, FabricInsight receives data from network devices in Telemetry mode and uses intelligent algorithms to analyze and display network data. The FabricInsight architecture consists of three parts: network device, FabricInsight collector, and FabricInsight analyzer.

Figure 2-1 FabricInsight architecture

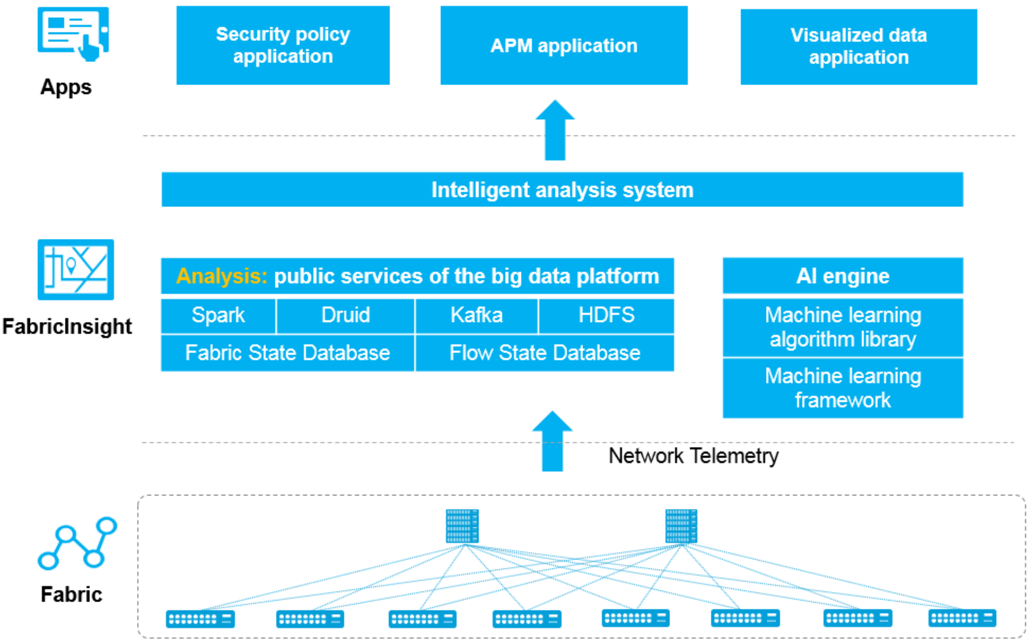
Apps
Security policy application
APM application
Visualized data application
Intelligent analysis system
FabricInsight
Analysis: public services of the big data platform
Spark
Druid
Kafka
HDFS
Fabric State Database
Flow State Database
AI engine
Machine learning algorithm library
Machine learning framework
Network Telemetry
Fabric

Network Devices

Network devices are switches on the data center network, such as the leaf and spine nodes in the figure. Currently, Huawei CE-series switches are supported. For the current FabricInsight version, devices need to report two types of data in Telemetry mode: TCP packets mirrored based on ERSPAN and performance metrics such as interface traffic reported based on the Google Remote Procedure Call Protocol (GRPC).

- ERSPAN mirrored packets: The forwarding chip on the switch identifies TCP SYN, FIN, and RST packets on the network and mirrors the packets to the FabricInsight collector through the ERSPAN protocol.
- GRPC performance metrics: Devices are connected as GRPC clients. Users can configure the Telemetry sampling function for a device using commands. The device then proactively establishes a GRPC connection with the target collector and sends data to the collector. The current version supports the following sampling metrics: CPU and memory usage at the device and board levels; number of sent and received bytes, number of discarded sent and received packets, and number of sent and received error packets at the interface level; number of congested bytes at the queue level; packet loss behavior data.

FabricInsight Collector

The FabricInsight collector collects data reported by switches in Telemetry mode, including TCP packets mirrored based on ERSPAN and performance metrics reported based on GRPC. For mirrored TCP packets, the collector adds timestamps to the packets, and packs and sends the packets to the analyzer for analysis. To improve the packet processing efficiency, the collector is implemented based on Intel Data Plane Development Kit (Intel DPDK). Therefore, the collector needs to support the DPDK network adapter. The Intel 82599 10GE network adapter is recommended.

FabricInsight Analyzer

The FabricInsight analyzer uses the microservice architecture. Each service is deployed in multi-instance mode, which features high reliability and scalability. You can expand the service capacity by expanding instance nodes. The FabricInsight analyzer cluster receives data from the collector, including TCP packets and performance metrics. The analyzer cleans different types of data using related cleaning logic, for example, calculating the forwarding path, forwarding latency, and link latency of packets. In addition, the analyzer analyzes application interaction relationships, associates applications with network paths, establishes dynamic baselines for some performance metrics based on the AI algorithms, detects exceptions, predicts the fault probability of optical modules. The analyzer can collect statistics on and analyze these data and display the analysis result.

TCP Traffic Collection

FabricInsight uses the remote flow mirroring capability of the switch to configure traffic classification on the switch to match TCP packets. Then, FabricInsight sends the packets to the monitoring device (FabricInsight collector) through the ERSPAN protocol.

Figure 2-2 TCP traffic mirroring

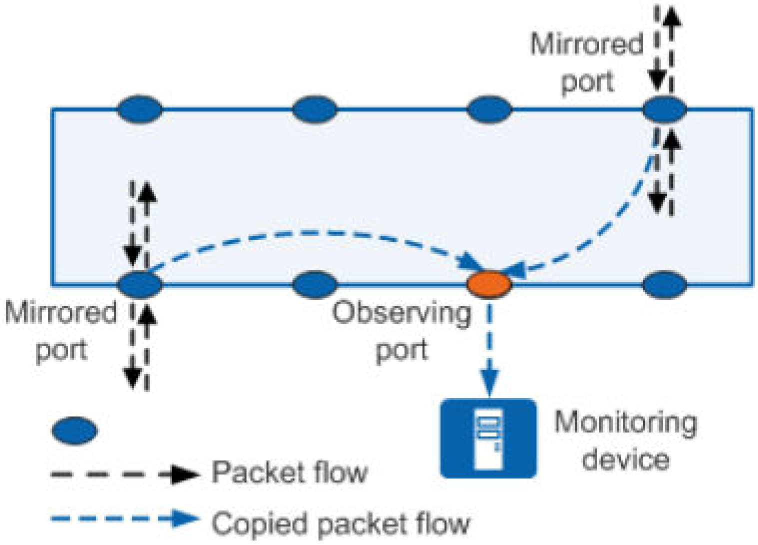

As shown in the following figure, assume that two VMs communicate with each other crossing leaf nodes. The red dotted lines indicate the packet routes. The remote mirroring in the inbound direction is enabled for each switch on the packet transmission route. If the packet passes through three hops from Leaf to Spine to Leaf, each of the three switches mirrors the packet to the FabricInsight collector once. The FabricInsight analyzer uses algorithms to restore the packet transmission route and perform related statistics and analysis.

Figure 2-3 FabricInsight mirrored packet collection

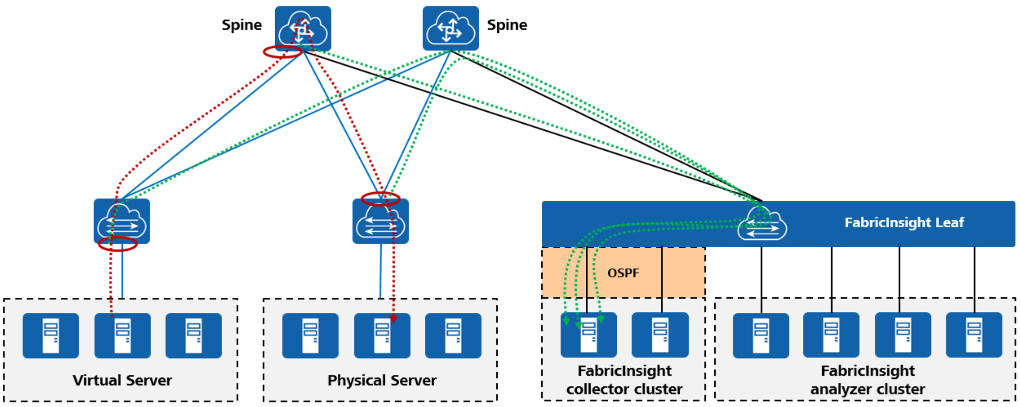
Spine
Spine
FabricInsight Leaf
OSPF
Virtual Server
Physical Server
FabricInsight collector cluster
FabricInsight analyzer cluster

In the TCP protocol, three handshakes are required for setting up a TCP connection and four handshakes are required for tearing down the connection. To monitor TCP connection setup and teardown between applications on the network, FabricInsight needs to mirror the SYN, FIN, and RST packets in the TCP protocol to the FabricInsight collector.

To enable the ERSPAN remote flow mirroring function on the switch, you need to install the ERSPAN plug-in on the switch. For details about how to install the ERSPAN plug-in, see the related manuals of the CE-series switches. After installing the ERSPAN plug-in, you need to complete related configurations on the device and enable the flow mirroring function. For device models that support the ERSPAN enhanced feature, you need to enable the ERSPAN enhanced feature when configuring flow mirroring. The configuration commands may vary depending on the device model and version. For details, see the configuration guide of the CE-series switch.

TCP Flow Analysis

Currently, TCP flow analysis supports the following capabilities:

Capability	Description
Packet forwarding path restoration	Restore the actual forwarding path of packets on the network based on the collected TCP packets. ● ERSPAN Type2 packets: The mirrored packets do not contain the forwarding port. In this case, the calculated transmission path contains the devices of each hop of the packet but the specific ports cannot be identified. ● ERSPAN Type3 packets: The mirrored packets contain the inbound forwarding port. In this case, the port of each-hop device that the packet passes through can be calculated based on the physical link data. The prerequisite is that each-hop device in the packet forwarding path must support the ERSPAN enhanced feature and have the feature enabled.
Packet forwarding latency calculation	The FabricInsight collector adds timestamps to the collected packets and calculates the hop-by-hop transmission latency. The 1588v2 clock synchronization must be configured between collector clusters.
TCP session traffic calculation	FabricInsight calculates the traffic of TCP sessions based on the TCP sequence number of SYN and FIN packets.
Application identification	When processing a reported TCP flow, FabricInsight can identify the application to which the TCP flow belongs based on the source IP address, destination IP address, and destination port number. Information on the application configuration page is entered based on the following hierarchy: application > cluster > network segment. An application can be configured with multiple clusters and a cluster can also be configured with multiple network segments.
TCP exception detection	The following exceptions can be detected: ● TCP signaling packet retransmission ● TCP connection setup failure ● TCP RST packet ● TTL exception ● TCP Flag exception

Telemetry Performance Metric Analysis

FabricInsight uses the Telemetry feature of CE-series switches to collect performance metrics of devices, interfaces, and queues, enabling users to actively monitor and predict network faults. The Telemetry feature uses the GRPC protocol to push data from devices to the

FabricInsight collector. Before using this feature, you need to import the Telemetry license on the device. The following table describes the metric sampling paths supported by FabricInsight of the current version.

Measured Object	Measurement Metric	Supported Device Type	Minimum Sampling Precision (FabricInsight Specifications)
Device/ Board	CPU usage	CE6810EI/CE6810LI/CE6850EI/CE6850HI/CE6850U-HI/CE6851HI/CE6855HI/CE6860EI/CE6870EI/CE6875EI/CE6880EI CE7850EI/CE7855EI CE8850EI/CE8860EI CE12804/CE12808/CE12812/CE12816/CE12804S/CE12808S/CE12804E/CE12808E/CE12816E	1 minute
	Memory usage		1 minute
Interface	Number of received/sent packets, number of received/sent broadcast packets, number of received/sent multicast packets, number of received/sent unicast packets, number of received/sent bytes		1 minute
	Number of discarded received/ sent packets, number of received/sent error packets		1 minute
Queue	Number of Buffer bytes		100 ms
Packet loss behavior	Forwarding packet loss and congestion packet loss	CE6865-48S8CQ-EI/ CE8850-64CQ-EI/CE12800E-X	100 ms

Dynamic AI Baseline Calculation

FabricInsight predicts the baseline for metrics including the device CPU/memory usage and number of interface received/sent packets through AI algorithms such as time sequence data feature decomposition and aperiodic sequence Gaussian fitting algorithms. Compared with the static threshold in the traditional NMS domain, the dynamic baseline is based on the historical data of a period of time and works with the anomaly detection algorithm based on the dynamic baseline to precisely detect metric deterioration on the network in advance.

In this version, FabricInsight creates CPU/memory usage baselines for all connected CE devices, and creates baselines of the number of received/sent packets for interfaces of physical links by default. The details are as follows.

Measured Object	Metric with Default Baseline	Maximum Period of Historical Training Data	Baseline Calculation Period	Baseline Retention Period
Device/ Board	CPU usage	Last 14 days	One day	One month
	Memory usage			
Interface	Number of received/ sent packets	Last 14 days	One day	One month

The dynamic baseline and predicted baseline are calculated once every other day through the Spark Streaming-based offline calculation framework. The granularity of the generated dynamic baseline data is the same as that of the original data. For devices, boards, and interfaces, the minimum data granularity of dynamic baseline data is one minute.

Real-time Exception Detection

Exception detection uses the Spark Streaming-based real-time calculation framework. The granularity of exception detection data is the same as that of original performance metric data. For devices, boards, and interfaces, the minimum granularity of baseline exception data is one minute. By default, FabricInsight performs exception detection calculation on metrics with dynamic baselines.

The core logic of exception detection is implemented by the Python operator. The Spark Streaming framework is used for distributed calculation. The execution logic of this operator is as follows:

- Point-by-point data comparison: Check whether the original data exceeds the baseline by the granularity of period.
- Identification and counting of consecutive out-of-range data: Check whether the out-of-range data is in consecutive periods and record the number of consecutive periods when out-of-range data is generated.
- Alarm suppression and combination: Suppress alarms based on specified rules to prevent excessive redundant baseline data from being generated. By default, a baseline exception is recorded only when the baseline is exceeded for three consecutive periods. In addition, the system automatically combines these out-of-baseline records into one record and the baseline exception record imported into the database contains the start time and end time of the exception.
- Output of the final baseline exception data: Write the calculation result to the storage exception Druid table.

Intelligent Issue Identification and Troubleshooting

FabricInsight performs big data analysis on collected TCP flows and Telemetry performance metrics through real-time and offline computing. In addition, FabricInsight proactively detects possible issues on the fabric based on AI algorithms such as baseline exception detection and multi-dimension clustering analysis, and intelligently analyzes and identifies whether the network or application has group issues. For service connectivity issues, FabricInsight automatically orchestrates troubleshooting procedures to support one-click automatic troubleshooting. All these help users achieve the proactive and intelligent O&M goal for proactive issue detection and minute-level issue locating and demarcation.

Based on the actual O&M scenarios of customers, FabricInsight collects and analyzes the issue case library on live networks of the customers, and summarizes more than 10 typical issue scenarios from the application quality, network service, and security compliance dimensions. In addition, FabricInsight proactively analyzes and identifies issues in different issue scenarios. If an issue is detected, FabricInsight automatically generates an alarm. Users can configure remote alarm notification rules to sense issues in real time. The following table describes the issue categories.

Category	Issue	Description
Application quality	Continuous Service Interruption	Identify continuous service interruption issues.
	Intermittent Service Interruption	Identify intermittent service interruption issues.
	Unreachable Host Port	Identify issues that some ports on the host are unreachable.
	Abnormal Sessions Matched Based on Rules	Allow users to define rules to identify three types of issues: SYN/SYNACK retransmission, SYN-RST, and high latency.
Network service	Insufficient TCAM Resources	Identify issues that the device TCAM resources are insufficient.
	Insufficiency or Sharp Change of FIB Entry Resources	Identify issues that the FIB entry resources are insufficient or change sharply.
	Insufficiency or Sharp Change of ARP Entry Resources	Identify issues that the ARP entry resources are insufficient or change sharply.
	Insufficiency or Sharp Change of MAC Entry Resources	Identify issues that the MAC entry resources are insufficient or change sharply.
Security compliance	Non-compliant Traffic Interaction	Allow users to define rules to identify non-compliant traffic interaction issues on the network.
	Suspicious SYN Flood Attack	Allow users to define SYN flood attack rules to identify suspicious SYN flood attacks on the network.

Category	Issue	Description
	Suspicious Port Scanning Attack	Allow users to define port scanning attack rules to identify suspicious port scanning attacks on the network.

Application Quality Issues

Application quality issues are mainly used to proactively identify applications with abnormal interaction behavior, for example, sessions that fail to set up TCP connections continuously and sessions that are intermittently disconnected repeatedly during connection setup. For these issues, FabricInsight orchestrates related troubleshooting procedures based on different issue patterns and provides the automatic troubleshooting capability. Users can perform one-click troubleshooting on the GUI. FabricInsight analyzes the result of each troubleshooting step and provides the final troubleshooting conclusion. The following uses the Continuous Service Interruption issue as an example to describe the application scenario, issue identification, and troubleshooting:

Issue Application Scenario

Users need to identify sessions (IP address triplet) with continuous TCP connection setup failure on the network, and use the AI clustering algorithm to analyze whether a group issue occurs on the network or application.

Sessions with continuous TCP connection setup failure are as follows:

1. Sessions for which the TCP connection setup never succeeds
2. Sessions for which the TCP connection setup succeeds before but fails continuously later

Issue Identification

FabricInsight calculates sessions with continuous TCP connection setup failure on the network in offline mode based on the Spark Streaming framework. Then, FabricInsight uses dynamic baseline and real-time exception detection technologies to identify time points when the failure times increase sharply and sessions with burst continuous connection setup failure. Finally, FabricInsight analyzes where a group issue occurs on the network or application based on the data.

Automatic Troubleshooting

There are many possible causes for Continuous Service Interruption issues. These issues may be caused by the network or application. Based on the expert experience library and troubleshooting process, FabricInsight summarizes a unified troubleshooting model and provides an automatic troubleshooting framework that can be orchestrated and requires no user perception. Troubleshooting actions cover check on the network and application. Users can perform troubleshooting through one-click, improving the troubleshooting efficiency quickly.

Troubleshooting Object	Possible Cause	Action
Destination host	The destination host is offline or the host system is faulty. As a result, the destination host does not respond.	Check whether the host sends a TCP SYN connection setup request packet.
		Check whether the host sends a TCP SYN ACK connection setup response packet.
		Check whether the host is in normal state.
Destination port	Listening is disabled for the destination port, leading to TCP connection setup failure.	Check whether the host port sends a TCP SYN ACK connection setup response packet.
		Log in to the host and check whether listening is enabled for the destination port.
Service access point	Entries of the service access point are missing, leading to packet forwarding failure.	Check whether the entries of the service access point are complete.
Layer 3 gateway	Layer 3 gateway entries are missing, leading to packet forwarding failure.	Check whether a route to the peer host exists on the Layer 3 gateway.
		Check whether the VNI and tunnel status on the VxLAN Layer 3 gateway are normal.
Firewall	The session is blocked by the security policy configured on the firewall.	Check whether firewalls on the packet forwarding path block the session.
Each hop of device through which the packet is forwarded (on the fabric)	The forwarding link is broken, congested, or faulty.	Check whether the forwarding link is broken (whether the link port is in **Down** state).
		Check whether congestion occurs on the forwarding path. (The Telemetry data reporting function needs to be enabled for the device.)
		Check whether abnormal port behavior occurs on the forwarding path. (The function of reporting port and optical module Telemetry data needs to be enabled for the device.)
		Check whether packet loss matching the session occurs on the device where packets pass through. (TD3 chips support this.)

Network Service Issues

Network service issues are used to proactively identify whether entry usage of the network device forwarding plane on the fabric is abnormal. For example, FIB route forwarding entries are insufficient or change sharply. For such issues, FabricInsight trains the dynamic baseline based on the static threshold or entry usage historical data to proactively identify exceptions in real time. In addition, FabricInsight can display the forwarding entry usage snapshot at the exception time point. For example, if the FIB entry usage is abnormal, FabricInsight allows users to view the resource usage of each VRF instance at the exception time point, enabling users to quickly analyze whether VRF instances with abnormal behavior exist. The following uses the Insufficiency or Sharp Change of ARP Entry Resources issue as an example to describe the application scenario and issue identification:

Issue Application Scenario

Users need to identify devices where ARP entry resources are insufficient or the resource usage changes sharply on the network, which is accurate to the specific board and chip.

Issue Identification

FabricInsight collects the ARP entry resource usage and total number of ARP entry resources of devices through Telnet (STelnet) at an interval of one minute. By comparing the collected data with the static threshold or training the dynamic baseline based on historical entry usage data, FabricInsight proactively identifies issues that resources are insufficient or resource usage changes sharply in real time. In addition, FabricInsight collects the detailed ARP entry resource usage at the exception time point.

Security Compliance

Security compliance issues are used to proactively identify potential SYN flood attacks, port scanning attacks, and non-compliant TCP sessions on the fabric. In attack scenarios, FabricInsight comprehensively analyzes related data and identifies the location of the suspected attack source, for example, the first device that the attack source SYN packet passes through or the real host where the attack source is located. This helps users check whether the attack is initiated from the external network or internal network. For non-compliant TCP sessions, FabricInsight identifies abnormal sessions based on rules configured by users, helping users audit non-complaint traffic. The following uses the Suspicious SYN Flood Attack issue as an example to describe the application scenario and fault identification:

Issue Application Scenario

You need to identify possible TCP SYN flood attacks on the network. The SYN flood attack source usually uses a large number of forged IP addresses to launch attacks. Once an attack occurs, network O&M personnel can hardly trace the attack source based on the forged IP addresses. FabricInsight analyzes the original packets, extracts original attack packets from a large number of packets, and restores the paths of these attack packets. By collecting statistics on the first-hop device of attack packets, FabricInsight can identify the network access location of the attack source, which greatly improves the efficiency for locating the attack source host.

Issue Identification

FabricInsight calculates whether TCP SYN flood attacks exist on the network in real time based on the Spark Streaming framework, and calculates the attack source location based on the actual packet forwarding path.

In addition, FabricInsight calculates the ERSPAN packets in real time and checks whether the destination host meets the SYN flood attack rate threshold. Users can adjust the default threshold on the issue setting page. The threshold conditions are as follows:

1. The TCP half-connection request rate of the destination host reaches a threshold. The TCP half-connection refers to that the destination host responds with a SYN ACK packet after receiving a SYN packet from the source IP address but receives no ACK packet from the source IP address. As a result, the TCP connection cannot be set up successfully. If the destination host has a large number of TCP half-connections, the half-connection queue resources of the TCP protocol stack in the operating system will be used up, and the host cannot respond to other normal session requests.
2. The TCP connection request rate of the destination host reaches a threshold. Normally, the TCP SYN packets received by the host on the fabric are relatively stable. If the number of TCP connection requests received by a host reaches a high threshold at a certain time, the host may suffer from SYN flood attacks.

If either of the preceding conditions is met, a suspicious SYN flood attack is identified. The Spark task calculation period is 10 seconds.

3 Business Values of FabricInsight

Issue Detection, Reducing Service Loss

Quick issue location to recover the services has the highest priority in O&M works. Once the DC is faulty, services will be greatly affected. According to the ITIC survey, business loss in one hour when the services are unavailable exceeds US$100,000 for 98% enterprises. According to DC practices of Huawei IT, the proportion of abnormal flows in a single POD is 3.67% and 70% issues cannot be identified through traditional O&M. These issues include unexpected go-offline of VMs, overhigh TOR load, and non-compliant access. O&M personnel handle these issues reactively, which prolongs the issue location period and undoubtedly brings more trouble to enterprises.

FabricInsight monitors data in real time using the Telemetry technology, and collects actual service flows and network KPI data. Based on services, FabricInsight proactively identifies 80% issues, including more than 10 types of typical issues related to the application quality, network service, and security compliance. In addition, FabricInsight displays active and historical issues in a panoramic manner. FabricInsight also correlatively analyzes applications, network paths, and devices, and provides one-click troubleshooting based on intelligent algorithms and expert experience library. This implements second-level automatic identification and one-click troubleshooting of typical issues, and helps locate issue points quickly.

The following uses the typical DC issue drill as an example to describe how to locate an issue:

- Issue identification: Communicate information layer by layer, specify issue symptoms, and determine information collection items.
- Issue location: View alarms, collect statistics on traffic, and use the **ping** command to locate the issue through traditional O&M.
- Issue isolation: Isolate the issue and solve it gradually.

Figure 3-1 shows the comparison between the time required for manually handling typical issues through traditional means and the time required for handling typical issues through FabricInsight.

Figure 3-1 Comparison between manual issue handling time and FabricInsight issue handling time

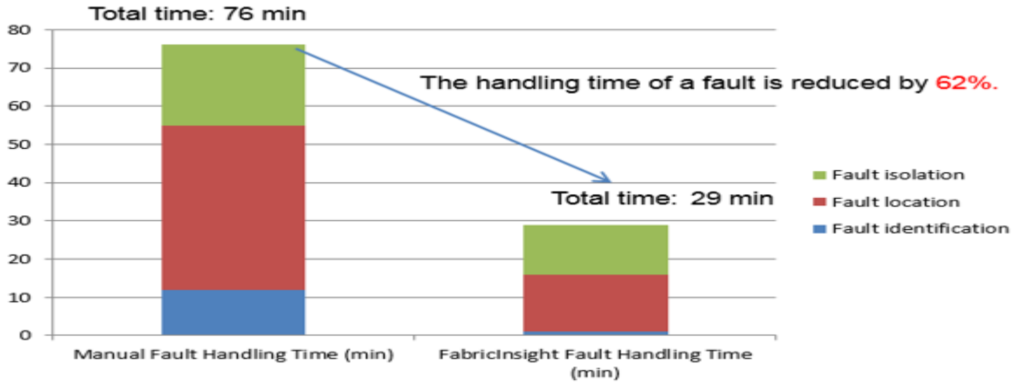

As shown in the figure, FabricInsight can reduce the issue handling time by 62% and reduce the service loss by about US$78,333.

System Down Times	Hourly Loss ($)	Manual Issue Handling Time	FabricInsight Issue Handling Time	Service Loss Reduced ($)
1	100,000	76 minutes	29 minutes	78,333

Insight into Applications and Networks, Simplifying Service Decision-Making and Significantly Reducing OPEX

In traditional O&M, locations are planned before deployment. The SDN enables automatic service deployment. However, where to deploy and what is the path are new O&M challenges. Cloud DC construction breaks the boundaries between IT and CT. During DC integration, migrating multiple traditional networks to the cloud DC networks is challenged by problems such as what are the bearing relationships between networks and services, which services need to be migrated together to prevent cross-DC traffic, how to ensure smooth access policies, and how to guarantee the SLA. Decisions will become a castle in the air without actual data analysis support.

Gaining insight into DC networks and applications, FabricInsight provides the network visualization capability to display the network topology and the application map to display interaction relationships between services based on real service flows, providing data support for service decision-making. In addition, the network administrator can quickly detect non-compliant access and traffic and take related measures, transforming from reactive security isolation to proactive application access sorting.

When releasing or migrating services, the network administrator can sort service relationships based on the network-wide application association map and configure policies correctly, which greatly saves the manpower and reduces the enterprise OPEX. According to the financial industry survey, it takes three months for two senior O&M engineers to sort application access relationships through traditional means during service migration in the scenario where over 200 servers, 300 VMs, and 100 applications are deployed. However, the sorting results may still be incomplete. Through the network-wide application interaction map of FabricInsight, the O&M engineer can analyze data interaction between applications and quickly identify non-compliant traffic. Inter-application interaction relationships can be exported conveniently. One engineer can sort out the application access relationships within seven days in E2E mode, reducing the sorting time by more than 90%.

Accurate Playback of Historical Issues, Proving Innocence of the Network

Based on the SNMP mechanism, traditional O&M has low performance. The issue occurrence time is usually missed due to the minute-level collection period. As a result, it is hard to locate an issue occurred in an application on the network.

Based on the Telemetry subscription and release mechanism, FabricInsight collects real service flows in real time and searches hundreds of billions of data records in seconds to quickly trace network performance issues. In addition, FabricInsight evaluates the service quality of the network in real time. Once the network forwarding quality deteriorates, the administrator can directly locate the specific route and associate the application with the network to quickly locate the issue.

Predictive Maintenance, Effectively Reducing the Network Issue Occurrence Rate

According to Huawei IT survey on network alarm statistics of the data center in a month, about 64% issues are packet loss, optical link, and traffic issues. The mistake rate of the current static threshold alarm detection is about 50% and the time required for troubleshooting an issue is long.

FabricInsight uses machine learning algorithms to train network behavior models based on big data, and displays the dynamic baselines at the device, queue, and port levels. In this way, exceptions are proactively detected when the service level decreases. In addition, FabricInsight can predict optical module fault risks, changing the traditional passive O&M to proactive O&M and enabling the network side to discover potential risks before the service side. According to Huawei IT network practices, the failure rate of the data center network is reduced by up to 68%.

4 Success Story

China Merchants Bank (CMB) is a leading commercial bank in the world. It is accelerating its transformation to "lightweight bank" and "retail finance 3.0", aiming to become a technology-oriented bank. With digital transformation, CMB builds intelligent IT infrastructure focusing on app operations and wants to build a high-intelligent database O&M system, achieving low-cost and large-scale O&M.

Huawei and CMB conduct innovative cooperation. Currently, Huawei FabricInsight has managed the network containing more than 6000 VMs at the Xili data center of CMB. Focusing on applications, FabricInsight proactively manages the network from the service perspective, monitors service flows in real time, opens the network black box of the data center, identifies the real network path of each flow, and visualizes all paths on the network.

According to real cases on the live network of CMB, 300,000 retransmission exceptions occurred at the data center within one hour. FabricInsight can proactively detect issues and locate network issue points in minutes, improving the O&M efficiency by 89% and helping CMB provide 24/7 financial services without interruption.

The cooperation between Huawei and CMB has also brought practical experience for the bank industry to transform to intelligent O&M, building a benchmark for smart finance construction in the industry.