

Omada Incident Detection and Health Assessment White Paper

Copyright © 2026 TP-Link Systems Inc. All rights reserved.

No part of this document may be reproduced, copied, or transmitted in any form or by any means without the prior written permission of the company.

Contents

1	Overview	4
1.1	Introduction	4
1.2	User Pain Points	5
1.3	Scope	5
1.4	Purpose of This White Paper	6
2	Key Technologies	6
2.1	Basic Concepts	6
2.1.1	Health	6
2.1.2	Incidents	7
2.1.3	Root Cause Analysis and Suggestions	8
2.1.4	Differences Between Health and Incidents	8
2.2	Health Dimensions and Scoring Criteria	8
2.2.1	Device-Level Health	10
2.2.2	Client-Level Health	16
2.2.3	Service-Level Health	19
2.2.4	Site-Level Health	22
2.3	Incident Categories	22
2.4	Incident Status	25
2.5	Relationship Between Health and Incidents	26
3	Key Technology Principles	28
3.1	Data Foundation and Detection Objects	28
3.2	Health Scoring Mechanism	28
3.2.1	Comprehensive Scoring Model	28
3.2.2	Metric Normalization and Score Band Definitions	29
3.2.3	Incident and Health Score	30
3.2.4	Dimensional Dynamic Weights	30
3.2.5	Smoothing Strategy	31
3.3	Incident Detection Process	31
3.3.1	Perception Layer	32
3.3.2	Preprocessing layer	32

3.3.3 Analysis Layer	33
3.3.4 Storage Layer	33
3.3.5 Presentation Layer	33
3.4 Incident Detection Rules	33
3.4.1 Real-Time Detection Rules	33
3.4.2 Periodic Detection Rules	34
3.5 Incidents and Root Cause Analysis Mechanism	36
3.6 Troubleshooting Workflow and Closed-Loop Verification	37
4 Pro Controller Compatibility	38
4.1 Changes to the Incident Analysis Mechanism	38
4.1.1 Changes to Incident Counting Rules	38
4.1.2 Handling Historical Ongoing Incidents	38
4.2 Changes to Notifications	39
4.2.1 Refactored Notification Categories	39
4.2.2 Merging Account-Level Notification Configurations	40
4.3 Streamlining and Optimization of Incident Settings	40
4.4 Function Support Status Display	41
4.5 Device Detail Health Correlation	41
5 Application Scenarios and Solutions	42
5.1 Multi-Branch Chain Store Networks	42
5.2 High-Density Office/Conference Networks	44
5.3 Campus/Hotel Networks	46
6 Future Outlook	47
6.1 Deeper Determination of Root Causes	47
6.2 From Fault Diagnosis to Fault Prediction	48
6.3 Fully Automated Troubleshooting and Self-Healing	48
7 Appendix	49
7.1 Glossary	49

1 Overview

1.1 Introduction

In enterprise networks, issues perceived by users are usually not caused by a single abnormal metric but by a combination of symptoms. For example, a failed client connection may be related to RF signal quality, authentication, DHCP address assignment, AP load, or the uplink. Likewise, a device operating normally does not necessarily mean the user experience is normal. And a short-lived fluctuation in a metric does not always mean immediate action is required.

Omada Incidents & Health provides Omada users with Health assessment and incident detection capabilities. By leveraging operational data from devices and clients, it helps users understand whether the network is healthy, which objects are at risk, whether an abnormal condition is still ongoing, and where to begin troubleshooting next.

In Omada 6.3, Health and Incidents are not two isolated sets of information. Users can quickly identify risks through Health, then drill down along the incident, root cause, and suggestions. Alternatively, users can start from the Incidents module, filter affected devices and clients based on detected incidents, and then navigate to the corresponding detail pages to view related Health metrics—gradually completing a closed-loop troubleshooting workflow: discover the issue → locate the issue → understand the cause → apply the fix.

This feature is not simply about displaying more charts or logs. Instead, it organizes underlying monitoring data into operations information that is easier to understand:

- **Health** is used to evaluate the overall status of devices, clients, or network experience.
- **Incidents** describe specific issues that are currently occurring or have occurred in the past.
- **Root cause analysis and suggestions** help users understand likely sources of issues and provide next-step troubleshooting directions.

With these three types of information, users can move beyond seeing score changes to understanding the cause and taking action.

1.2 User Pain Points

In daily operations, Omada users' common challenges are not limited to whether there is an alarm but rather the difficulty of mapping alarms and metrics to actual user experience:

- **Traditional operations are device-status-centric.** Even if a device is online and its CPU and memory are normal, client access, roaming, and internet experience may still be impacted. What users truly care about is whether services and experience are affected.
- **Too many metrics, but little experience-level interpretation.** Metrics such as CPU, memory, channel utilization, packet loss, and authentication results are scattered across different pages, and users must rely on experience to judge whether they affect client experience.
- **Lack of guidance after an incident occurs.** Logs only show what happened. Users still need to know the affected objects, likely causes, recommended actions, and whether the issue has been resolved.
- **No summary of overall network Health.** Users need to first understand whether the network is healthy overall, then drill down to devices, clients, Wi-Fi, or specific incidents—instead of inspecting raw data item by item.

Omada Health and Incidents are designed to turn this fragmented information into a more natural troubleshooting path: use Health to assess status, use Incidents to pinpoint issue types and affected objects, and finally use root cause and suggestions to remediate and confirm recovery.

1.3 Scope

This white paper focuses on the capabilities related to Health assessment and incident detection in Omada 6.3, including:

- Site Health
- Device Health
- Client Health
- Wi-Fi Health
- WAN Health

- Incident detection, categorization, severity, status, root causes, and suggestions
- Correlated presentation and troubleshooting entry points between Health and incidents

Note: Supported Health dimensions and incident types vary by object. Refer to the actual Controller UI and device capabilities for details. In version 6.3, device types such as OLT, NVR, and IPC do not support Health assessment or incident detection.

1.4 Purpose of This White Paper

This document is intended for Omada users. It explains the basic principles, usage definitions, and troubleshooting workflow of Health and incident detection to help users answer the following questions:

- What the Health score reflects
- Why the Health score changes
- How incidents are categorized and presented
- What the statuses Ongoing, Unresolved, Resolved, and Ignored mean
- How Health, incidents, root causes, and suggestions relate to each other
- How to proceed with troubleshooting after seeing a degraded Health score or an incident

2 Key Technologies

2.1 Basic Concepts

2.1.1 Health

Health is Omada's composite score for the operational status of a network object. Rather than a simple display of a single metric, it is evaluated across multiple dimensions that reflect device operation and user experience. The supported dimensions may vary by object, and the final result depends on whether valid data is available.

Health status can generally be interpreted as follows:

Status	Score Range (0–10)	Score Range (0–100)	Meaning
Good	8-10	80-100	The current status is good with no obvious risks observed.
Fair	4-7	40-79	Some risk or degraded experience is present and requires attention.
Poor	1-3	0-39	The status is poor, and troubleshooting is recommended as a priority.
No Data	N/A	N/A	There is currently not enough valid data to evaluate the dimension or object, or the object is offline.

On the UI, different Health types use different units: Device Health and Client Health are displayed on a 0–10 scale, while overall Health, such as Site, Wi-Fi, and WAN, is displayed as a percentage (0–100). Regardless of the format, the meaning is the same: a higher score indicates a Healthier object, while a lower score indicates that the user should review the contributing dimensions and related incident information.

2.1.2 Incidents

An incident is a specific issue identified by Omada based on operational data from devices, clients, and network services, and is presented as a structured record. Compared with standard logs, incidents not only record what happened, but also describe which object the issue occurred on, when it occurred, whether it is still ongoing, and where to start troubleshooting. An incident can be triggered by an explicit state change, or by a metric that continuously meets certain conditions over a period of time.

An incident typically includes:

- Incident name, category, and severity.
- The affected object and the potential impact scope.
- First occurrence time, last occurrence time, and current status.

- Possible root causes, related metrics, and recommended actions.

Typical incidents include:

- Client failed to connect
- Client failed to obtain IP
- AP channel utilization too high
- AP wireless interference rate too high
- Switch port detects error packets

2.1.3 Root Cause Analysis and Suggestions

Root cause analysis is intended to explain what may have caused an incident, while suggestions tell users what to check next or how to optimize.

Root cause analysis does not mean the system can always provide a single, definitive cause in every scenario. Network issues are often influenced by multiple factors. Therefore, Omada summarizes issues into directions that users can understand and validate—for example, wireless interference, weak signal, authentication failure, port/link abnormalities, or excessive device resource utilization.

2.1.4 Differences Between Health and Incidents

Health and incidents address different questions:

- **Health** answers: "Is the current status good?"
- **Incidents** answer: "What issue occurred?"
- **Root causes and suggestions** answer: "Why might it have happened, and what should I check next?"

As a result, Health is suited for an overall assessment, while incidents are suited for pinpointing specific problems. In troubleshooting, users typically start by reviewing changes in Health, then examine incidents and the associated root causes and suggestions.

2.2 Health Dimensions and Scoring Criteria

Health for different objects is composed of different dimensions. Omada selects metrics that best reflect the object's operating status and experience quality based on the object type. Each dimension is first scored independently (on either a 0–10 scale or a 0–100

scale) and then aggregated by weight into the Health score at the corresponding level.

Omada's Health system uses a layered design—from individual device/client health at the bottom, to device/service health, and up to overall Site Health—forming a complete assessment framework:

1. Site Health: Provides a global, overall view of Health.

As the highest-level Health indicator, it reflects the operational status of the entire site by aggregating the Health conditions of devices, clients, and network services. This is a weighted composite result, not a simple average.

2. Device Health: Provides device-centric Health insights, focusing primarily on the device itself.

- Includes Health assessment for devices such as Gateways, Switches, and APs.
- Each device type has its own assessment metrics.
- A device's importance in the network affects its weight. Critical devices (such as the Gateway) have a larger impact weight.

3. Client Health: Provides client-centric Health insights, focusing primarily on the client itself.

4. Service Health: Aggregates information from key services to provide service-centric Health insights.

- Mainly covers the Health status of core network services, including WAN and Wi-Fi.
- WAN depends on the Gateway, while Wi-Fi service is evaluated as a Health dimension on Wi-Fi service quality by aggregating parameters from APs and wireless clients.
- The status of these services directly affects the connection quality for devices and clients.

Together, these dimensions form a complete bottom-up Health assessment system. The system makes comprehensive judgments based on network topology, object types, and available valid data. The exact combinations and internal weights may change with versions and product strategy.

2.2.1 Device-Level Health

Device-level health mainly answers: “Are the network devices running stably?” At this level, the focus is on baseline capabilities such as device resources, link status, port quality, and the wireless environment—because problems in these areas often further impact client experience and the overall Site Health.

AP Health

AP Health primarily focuses on device resources and wireless environment quality. In Omada 6.3, APs do not have a standalone Incident dimension. This is because AP-related issues are already reflected in base dimensions such as CPU, channel usage, and channel interference rate, so additional duplicate deductions are avoided.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
CPU Usage	Sustained high CPU indicates heavier processing load on the AP.	Good: $\leq 80\%$ Fair: 80%–90% Poor: $> 90\%$	Resource pressure is determined based on continuous sampling results and device capability.
Memory Usage	The tighter the memory, the higher the risk to device stability.	Good: $\leq 90\%$ Fair: 90%–95% Poor: $> 95\%$	Resource pressure is determined based on continuous sampling results and device capability.
Channel Usage	A busier channel more easily degrades access and throughput	Good: 2.4 GHz $\leq 70\%$, 5 GHz / 6 GHz $\leq 60\%$	Evaluated per enabled band. Actual thresholds are subject to the device model and

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
	experience.	Fair: 2.4 GHz 70%–85%, 5 GHz / 6 GHz 60%–90% Poor: 2.4 GHz > 85%, 5 GHz / 6 GHz > 90%	software version.
Channel Interference Rate	Stronger interference and higher noise floor make connections more unstable.	Interference Rate: Good: 2.4 GHz < 25%; 5 GHz / 6 GHz < 25% Fair: 2.4 GHz 25%–65%; 5 GHz / 6 GHz 25%–53% Poor: 2.4 GHz > 65%; 5 GHz / 6 GHz > 53% Noise Floor: Good: < –80 dBm Fair: –80 to –70 dBm	Interference rate and noise floor are evaluated per band. Actual thresholds are subject to the device model and software version.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
		Poor: > -70 dBm	
Wireless Transmission Quality	More retransmissions and packet loss mean worse wireless efficiency, integrity, and user experience.	Retransmission Rate: Good: < 20% Fair: 20%–48% Poor: > 48% Packet Loss Rate: Good: < 3% Fair: 3%–5% Poor: > 5%	Retransmission rate and packet loss rate are calculated within the detection window and evaluated together. If no valid traffic data is available, No Data is shown.

Note: For the Wireless Transmission Quality dimension, if there is no valid traffic data within the scoring window, the UI displays **No Data**. This is a data condition and does not indicate that the AP or wireless link is abnormal. Common reasons include no clients using the network during the detection window, or no increase in the number of valid traffic packets.

Switch Health

Switch Health mainly reflects device resource status and wired link quality. Compared with APs, Switch Health places more emphasis on wired-side issues such as port packet loss, error frames, and port faults.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
CPU Usage	Sustained high CPU load can affect forwarding and management experience.	Good: $\leq 80\%$ Fair: 80%–90% Poor: $> 90\%$	Resource pressure is determined based on continuous sampling results and device capability.
Memory Usage	Insufficient memory increases operational risk.	Good: $\leq 90\%$ Fair: 90%–95% Poor: $> 95\%$	Resource pressure is determined based on continuous sampling results and device capability.
Temperature	Higher temperatures increase hardware risk.	Good: $-10 \leq T < 93\text{ }^{\circ}\text{C}$ Fair: $T < -10\text{ }^{\circ}\text{C}$ or $93\text{ }^{\circ}\text{C} \leq T < 101\text{ }^{\circ}\text{C}$ Poor: $T \geq 101\text{ }^{\circ}\text{C}$	Determined by the MAC temperature range.
Packet Loss Rate	More port packet loss leads to worse wired experience.	Good: $\leq 0.3\%$ Fair: $> 0.3\%$ and $< 0.7\%$ Poor: $\geq 0.7\%$	Determined by port packet loss within the valid sampling window.
Error Frame	More error frames indicate poorer link	Good: $\leq 0.15\%$	Determined by port error frames within the

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
Rate	quality.	Fair: > 0.15% and $\leq$ 0.3% Poor: > 0.3%	valid sampling window.
Faulty Port	More port faults indicate worse link stability.	Good: linkerrorPort = 0–1 Fair: linkerrorPort = 2–3, and no special low-score rule is triggered Poor: usedPort > 2 and at least half of usedPort are linkerrorPort, or linkerrorPort $\geq$ 4	Determined by a combined evaluation of the number of faulty ports and port usage.
Incidents	Identified access issues, wired-environment issues, device status issues, or security issues.	N/A	Troubleshooting priority is determined by incident severity and incident status.

Gateway Health

Gateway health mainly reflects device resources, egress link quality, and port stability. In 6.3, the criteria are primarily for wired Gateways; wireless Gateways reuse some AP dimensions on this basis.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
CPU Usage	Sustained high CPU utilization affects processing capacity.	Good: CPU utilization $\leq 60\%$ Fair: $> 60\% - 80\%$ Poor: $> 80\%$	Resource pressure is determined based on continuous sampling results and device capability.
Memory Usage	Tighter memory increases stability risk.	Good: $\leq 50\%$ Fair: $> 50\% - 70\%$ Poor: $> 70\%$	Resource pressure is determined based on continuous sampling results and device capability.
Faulty Port	Frequent WAN/LAN port abnormalities indicate unstable links.	Good: linkerrorPort = 0–1, and no special low-score rule is triggered Fair: linkerrorPort = 2–3, and no special low-score rule is triggered Poor: All usedPort are linkerrorPort, or linkerrorPort ≥ 4	Determined based on port status and whether errors remain abnormal over time.
WAN Health	The more unstable the egress link, the worse the site-wide	N/A	Determined by latency, packet loss, and jitter

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
	access experience.		levels.
Incidents	Identified WAN issues, service issues, device status issues, or security issues.	N/A	Troubleshooting priority is determined by incident severity and incident status.

2.2.2 Client-Level Health

Client-level health mainly answers: "After users connect to the network, is the actual experience good?" This level focuses not only on whether the client can successfully join the network but also on factors that strongly correlate with perceived experience, such as RSSI, negotiation rate, roaming, and authentication.

Wireless Client Health

Wireless Client health focuses on whether a client can connect successfully and maintain a good wireless experience. Different from the full incident list on the details page, the Incidents dimension in Health includes only key incidents that participate in scoring, to avoid duplicate deductions with baseline dimensions such as RSSI, negotiation rate, and onboarding time.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
RSSI	Weaker signal means poorer coverage and stability.	Good: ≥ -67 dBm Fair: -68 to -75 dBm	RSSI is evaluated per active band, subject to device capability.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
		Poor: < -75 dBm	
Negotiation Rate	Lower link capability increases the risk of slow internet access.	Good: Meets the good rate baseline for the corresponding PHY and band Fair: Below the good rate baseline but above the minimum expected baseline Poor: Below the minimum expected baseline	Determined against rate baselines based on wireless standard, band, and spatial-stream capability.
Onboarding Time	The longer it takes from the start of access until normal internet access is available, the worse the experience.	Good: Short, within the corresponding expected baseline Fair: Close to or reaching the expected baseline Poor: Clearly exceeds the expected baseline	Determined by time spent in stages such as association, authentication, DHCP, and DNS.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
SNR	A smaller signal-to-noise ratio indicates poorer link quality.	Good: ≥ 35 dB Fair: 25–34 dB Poor: < 25 dB	SNR is evaluated per active band.
Incidents	Identified key incidents directly affect experience.	N/A	Focuses on the severity, status, and duration of incidents included in Health evaluation.

In Omada 6.3, the only key incidents included in the Wireless Client Incidents dimension are: Client sticky roaming, Client ping-pong roaming, Client Portal authentication failure, and Client Portal authentication too slow.

Issues such as abnormal client SNR, abnormal negotiation rate, access failure, slow access, IP acquisition failure, and slow IP acquisition are already represented by independent dimensions and are therefore not counted again in the Incidents dimension.

Wired Client Health

Wired Client health focuses on whether a wired client can successfully access the network and whether the link remains stable after access—emphasizing both “can it connect” and “is the link healthy after it connects.”

For wired clients, both connection and link error are evaluated using rule-based determination, without unified metric threshold ranges. If either item shows a clear issue, user experience may be affected. Users should troubleshoot by considering both pieces of information together.

Dimension	How users can interpret it	Health thresholds for metrics	Determination basis
Connection	Whether the client successfully obtains network access.	Good: No client abnormality Fair: IP acquisition is normal, but authentication status is abnormal Poor: IP acquisition status is abnormal	Determined by whether IP acquisition and authentication status show abnormalities.
Link Error	Whether the switch port connected to the client has persistent error-packet issues.	Good: No error packets, or no continuous increase in error packets Poor: Error packets keep increasing continuously	Determined by whether error packets continue to increase.

2.2.3 Service-Level Health

Service-level health mainly answers: "Are network services stable and available?" This level includes both WAN quality and overall wireless environment/capacity, because service-level problems often impact user experience across multiple devices and clients.

WAN Health

WAN health describes egress link quality. It focuses on three key aspects: whether latency is too high, packet loss is significant, and jitter is severe. WAN health is evaluated by aggregating these three indicators (i.e., latency, packet loss, and jitter), and does not rely on a single fixed threshold.

Wi-Fi Health

Wi-Fi health reflects the overall wireless experience of the site. The key is not whether a single AP is good but whether the site-wide wireless environment, coverage, capacity, and client experience are all good at the same time.

Dimension	How users can interpret it	Health grading for metrics	Determination basis
Client Onboarding Time	Whether most clients can access successfully and maintain a normal experience.	Good: 80% or more of wireless clients are rated Good on the Onboarding Time dimension Fair: The site-wide average score of Onboarding Time falls within the Fair range Poor: The site-wide average score of Onboarding Time falls within the Poor range	Evaluates whether client access quality and experience metrics meet the corresponding criteria.
Channel Utilization	Whether channels across the site are generally busy.	Good: 80% or more of APs are rated Good on Channel Usage Fair: The site-wide average score of AP Channel Usage falls within the Fair range Poor: The site-wide average score of AP Channel Usage falls within the Poor range	Focuses on whether per-band channel utilization on APs continuously exceeds the corresponding thresholds.

Dimension	How users can interpret it	Health grading for metrics	Determination basis
Channel Interference Rate	Whether there is noticeable co-channel interference and noise impact across the site.	Good: 80% or more of APs are rated Good on Channel Interference Rate Fair: The site-wide average score falls within the Fair range Poor: The site-wide average score falls within the Poor range	Focuses on whether per-band interference rate and noise floor on APs continuously exceed the corresponding thresholds.
RSSI	Whether wireless coverage is uniform and most clients have good signal.	Good: 80% or more of wireless clients are rated Good on RSSI Fair: The site-wide average score of RSSI falls within the Fair range Poor: The site-wide average score of RSSI falls within the Poor range	Focuses on whether wireless client signal strength meets the corresponding coverage thresholds.
Wireless Access Capacity	Whether the number of connected clients exceeds the comfortable capacity	N/A	Evaluated by combining AP client load, recommended capacity, and inter-AP load distribution.

Dimension	How users can interpret it	Health grading for metrics	Determination basis
	range, and whether load is balanced across APs.		

2.2.4 Site-Level Health

Site-level health mainly answers: "Is the site healthy overall?" It does not introduce a new set of underlying metrics. Instead, it aggregates device-level, client-level, and service-level information upward, giving users the most intuitive global assessment.

Site Health

Site health provides an overall site-level judgment. It is not a new set of independent underlying metrics but a unified site score derived by combining four categories: Device, Client, Wi-Fi, and WAN.

Site Health is assessed with reference to Device, Client, Wi-Fi, and WAN results. The system evaluates based on the objects actually present in the site and the availability of valid data:

- If a corresponding object or service does not exist, it is not treated as abnormal data for evaluation.
- If a key dimension degrades significantly, the site status will reflect that risk with higher priority.

2.3 Incident Categories

Omada's incident categorization is redesigned around a troubleshooting perspective, rather than purely by underlying technical modules. The goal is to help users quickly determine roughly where the problem is based on the incident category.

For example:

- For **Access** incidents, users can prioritize checking client association, IP acquisition, DHCP, and related stages.
- For **Authentication** incidents, users can prioritize checking Portal/802.1X configurations and service status.
- For **Wireless Network** incidents, users can prioritize checking signal, SNR, channel utilization, noise floor, interference rate, or negotiation rate.
- For **Link** incidents, users can prioritize checking ports, LACP, loops, and link stability.
- For **Device Status** incidents, users can prioritize checking CPU, memory, PoE, optical modules, and other device operational conditions.

This categorization is designed to guide users from symptoms into a troubleshooting path. Even if a single incident may have multiple technical causes, users can first narrow down the scope by category, and then review root causes and suggestions.

In Omada 6.3, incidents cover five object types: Gateway, Switch, AP, Wired Clients, and Wireless Clients. There are nine incident categories as listed below:

Category	Covered Objects	Full Incident List
Access	Wireless Clients, Switches, Gateways	Client failed to connect; Client failed to obtain IP; Interface IP address conflict; DHCP server conflicted; Device connection tracking session limit reached
Authentication	Wireless Clients, Wired Clients	Switch 802.1X authentication failure; Client Portal authentication failure; Switch 802.1X authentication took too long; Client Portal authentication too slow

Category	Covered Objects	Full Incident List
Roaming	Wireless Clients	Client ping-pong roaming; Client sticky roaming
Wireless Network	APs, Wireless Clients	AP channel utilization too high; AP noise floor too high; AP wireless interference rate too high; Client SNR exception; Client negotiation rate exception
Wired Network	Switches	Host routing entry overflow; ARP table overflow; Multicast table overflow was detected; Illegal multicast packet; Switch port traffic too high; Switch LAG Group member port load too different; Switch port packet loss rate too high; Storm control event
Link	Switches	Loop detected; LACP convergence failure; Repeated address migrations; Switch port detects error packets; Switch Port goes up/down frequently
WAN & Services	Gateways	DDNS service abnormal; Gateway WAN port bandwidth utilization rate too high; Gateway WAN port latency too high; VPN connection failure
Device Status	APs, Switches, Gateways	AP frequently offline; APs on site going offline in batches; Gateway goes offline from the controller frequently; Switch goes offline frequently; AP CPU utilization rate too high;

Category	Covered Objects	Full Incident List
		Gateway CPU utilization rate too high; Excessive PoE power; Switch CPU utilization rate too high; Abnormal optical module DDM
Security	Switches, Gateways	Unauthorized log in from a gateway; Gateway WAN or LAN port detected attack packets; Unauthorized log in from a switch

2.4 Incident Status

Incident status helps users determine whether this issue still needs action now.

Common statuses in Omada include:

Status	Definition	How users can interpret it
Ongoing	The issue is currently active and detected by the Controller. It will be automatically marked as Resolved once no longer detected.	The issue is still ongoing and continues to be detected by the Controller. Once it is no longer detected, it will be automatically marked as Resolved .
Unresolved	The incident has occurred but has not yet been addressed.	The incident has occurred but has not been handled or confirmed yet.
Resolved	The issue is no longer active and has been automatically or manually marked as Resolved .	The issue is no longer ongoing, and has been marked as Resolved either automatically by the system or manually by the user.
Ignored	The incident has been	The user has dismissed the

Status	Definition	How users can interpret it
	dismissed and requires no further action.	incident and considers no further action necessary.

Status does not change the fact that an incident occurred. It is mainly used to distinguish current handling priority. For Ongoing and Unresolved incidents, users are advised to prioritize investigation. For Resolved incidents, users can use them for historical review. For Ignored incidents, users typically do not need to take further action.

2.5 Relationship Between Health and Incidents

Health and incidents are related, but they do not have a one-to-one correspondence.

On one hand, a drop in certain Health dimensions may trigger or correlate with incidents. For example, high channel utilization, high interference, and abnormal client SNR may all manifest as degraded wireless experience. On the other hand, the mapping between incidents and Health is not a simple one-to-one relationship. For issues already expressed directly by independent Health dimensions—such as CPU risk, abnormal SNR, abnormal negotiation rate, access failure, or IP acquisition failure—the system will not count them again in the Incidents dimension. By contrast, incidents such as roaming and Portal authentication—where it is more appropriate to present a clear problem conclusion—participate directly in the Incidents dimension determination.

Users can interpret them as follows:

- **Health** is the overall status view.
- **Incidents** are the issue list.
- **Root causes and suggestions** are the troubleshooting clues.

Therefore:

- When Health is low, users should prioritize reviewing incidents on the corresponding object or at the corresponding time point.
- When there are many incidents, users can return to Health dimensions to see

whether overall experience has been impacted.

Health and incidents are not two parallel lines; they are different layers within a layered diagnostic model. The overall intelligent O&M framework can be understood as three layers:

Layer	Focus	Analogy	Description
Layer 1	Baseline “vital signs” Health dimensions	Physical exam report	Baseline metrics such as CPU, memory, channel, transmission quality, signal strength, and SNR are continuously monitored in the background to form a foundational Health profile.
Layer 2	Incident diagnosis	Doctor's diagnosis	When metrics degrade continuously, states change, or rules are hit, the system generates incidents with clearer problem orientation and suggestions.
Layer 3	Overall Health rating	Comprehensive diagnostic conclusion	Baseline dimensions and incident conclusions are combined to produce the final Health score and Health status.

The core logic of this model can be summarized as: when there is a clear issue, prioritize incidents; when there is no clear issue, rely on baseline dimension metrics. In other words, Health provides continuous sensing, incidents provide explicit diagnosis, and the overall Health rating unifies both into the final assessment.

3 Key Technology Principles

3.1 Data Foundation and Detection Objects

Omada's Health and incident detection capabilities are based on analysis of device, client, and system operational data collected by the controller. This data includes:

- Device operating data: CPU, memory, temperature, PoE, power supply, optical module, etc.
- Wireless environment data: channel utilization, noise floor, interference rate, signal-to-noise (SNR), negotiation rate, retransmission and packet loss, etc.
- Access experience data: association time, authentication result, IP address acquisition result, Portal or 802.1X authentication status, etc.
- Link and service data: port status, port errors, packet loss, WAN quality, VPN, DDNS, DNS, etc.
- Incident data: incident name, level, status, time of occurrence, affected objects, root cause, suggestions, etc.

After being processed, this data will be used for Health scoring, incident detection, root cause analysis, and page display.

3.2 Health Scoring Mechanism

In Omada Controller V6.3, the Health scoring mechanism is responsible for organizing basic vital signs and incident diagnosis conclusions into a unified user view. Users can naturally navigate from score changes to incident location and subsequent repair without having to jump back and forth between two separate modules.

3.2.1 Comprehensive Scoring Model

Health assessment first selects relevant indicators based on the object type, and then makes a comprehensive judgment based on the current status, duration, data validity, and incidents of the indicators. Different indicators use judgment directions that match their meanings: higher resource consumption, latency, packet loss, and interference generally indicate higher risk; lower signal strength, SNR ratio, and negotiation rate generally indicate higher risk.

3.2.2 Metric Normalization and Score Band Definitions

Since different device models, wireless standards, and hardware capabilities have different reasonable operating ranges, the controller does not directly use a single set of fixed raw indicator ranges across different models to evaluate all devices. The system combines indicator type, indicator direction, device capabilities, and valid sampled data to convert the raw indicators into comparable converted values, and then maps them into dimension scores.

This process can be summarized as "removing extreme values and taking the mean + normalizing relative to the benchmark." First, necessary outlier processing and aggregation are performed on the effective samples within the detection window to reduce the impact of instantaneous extreme values. Then, based on the benchmark values corresponding to the device and indicators, the indicators are converted to a unified comparison standard. For indicators such as utilization, latency, packet loss, and interference, higher values generally indicate higher risk; for indicators such as signal strength, SNR ratio, and negotiation rate, lower values generally indicate higher risk. Specific benchmark values, boundary processing, segmented mapping, and device lookup tables may vary depending on the model, standard, and version. This paper does not elaborate on the internal implementation details.

The final Health score adopts a unified classification:

Final Health Score	State	Indication
8-10	Good	Overall, the indicators are in a favorable operating range, and no significant risks have been observed.
4-7	Fair	Some indicators have entered a range requiring attention, which may indicate a decline in user experience or potential risks.
1-3	Poor	One or more key indicators are in the Poor range. We

Final Health Score	State	Indication
		recommend prioritizing investigation.
No Data / Offline	Not Applicable	The lack of sufficient valid data, or the object being offline, makes it impossible to determine whether it is Good, Fair, or Poor.

3.2.3 Incident and Health Score

Incidents are not an alert list independent of Health status. For incidents included in the Health assessment, the system will combine the incident type, severity level, current status, duration, and affected objects to reflect the incident conclusions in the corresponding Health status dimensions; issues already expressed by independent basic indicators will not be penalized again due to repeated incident generation.

Incidents use a four-level severity rating system: **Info**, **Warning**, **Error**, and **Critical**. **Critical** indicates a serious issue that requires priority handling. When an issue is rated **Critical**, it triggers a veto, and the upper limit of the relevant dimensions and the object's Health score is locked in the Poor range to prevent serious issues from being completely offset by high scores on other normal indicators.

The complete incident list is still used for problem localization and historical review. Whether or not it participates in Health scoring depends on the incident scope defined in the current product version.

3.2.4 Dimensional Dynamic Weights

Health weights are used to reflect the relative impact of different dimensions on the operational status and user experience of an object, rather than simply averaging various indicators. Different objects will choose more representative dimensions to participate in the evaluation based on their roles and observable capabilities.

When a dimension is temporarily unavailable due to device capabilities, object status, or data validity, the system will evaluate based on the remaining valid dimensions to avoid

directly interpreting **No Data** as a risk score. Specific weights and handling of missing dimensions may vary depending on object type, version, and product strategy.


3.2.5 Smoothing Strategy

For indicators susceptible to transient fluctuations, the system combines multiple valid samples within the detection window for smoothing, ensuring that the Health score reflects the recent state without frequent jumps due to a single abnormal sampling point. Recent samples generally better reflect the current state, while historical samples are used to provide necessary stability references.

When an object is newly online or currently offline, there are insufficient sampling points, or data is temporarily unavailable, the system will display **No Data**, **Offline**, or other corresponding states based on the data status, and will not directly treat missing data as abnormal scores.

3.3 Incident Detection Process

The overall process of incident detection can be understood as a hierarchical diagnostic model from data collection to user processing. First, the controller collects operational data of the devices and system, then preprocesses and analyzes it to form incidents, root causes, and suggestions, and finally presents them to the user on the page for processing.



3.3.1 Perception Layer

The perception layer collects operational data of the devices and system, such as device resources, port status, wireless environment, access process, authentication results, DHCP results, client connection quality, and related system service status.

3.3.2 Preprocessing layer

The preprocessing layer organizes raw reported data into indicators for detection. For example, it aggregates sampling points over a period of time to identify the key status of a device, client, or site within the detection window.

3.3.3 Analysis Layer

The analysis layer is responsible for determining whether any incidents exist.

Depending on the incident type, detection methods can be divided into two categories:

- **Real-time detection:** Suitable for issues directly triggered by device or user behavior, such as changes in device status, some connection results, or immediate incidents.
- **Periodic detection:** Suitable for problems that require data collection over a period of time before a judgment can be made, such as high channel utilization, interference, access failure rate, port errors, etc.

3.3.4 Storage Layer

When the system detects an abnormality, it generates an incident and attempts to correlate the root cause and suggest solutions. The incident log saves information such as status, time, affected objects, root cause, and suggestions for later querying and display.

3.3.5 Presentation Layer

Users will eventually see Health status and incidents on the Dashboard, device details page, client details page, or Incidents page. These pages will help users filter and locate issues based on information such as incident type, object, time, status, and severity.

3.4 Incident Detection Rules

Incident detection rules are used to transform raw metrics, connection results, and state changes into issues that users can understand. The rules differ for different incidents, but can be broadly categorized into real-time detection and periodic detection.

3.4.1 Real-Time Detection Rules

Real-time detection is suitable for incidents triggered by explicit behaviors or state changes. Its advantage lies in the rapid identification of incidents after they occur, making it suitable for issues such as access failures, authentication failures, device state changes, and partial link state changes.

Incident Name	Incident Definition	Encountered Issue
Client failed to connect	The client connection failed.	A client failed to connect, with related information such as the associated AP, SSID, and frequency band.
Client Portal authentication failure	Client Portal authentication failed.	The client failed to complete Portal authentication.
Client sticky roaming	The client is unable to roam.	The client was unable to switch to a more suitable AP or frequency band.
Excessive PoE power	PoE power exceeds limit.	The powered device malfunctioned or repeatedly disconnected.
Abnormal optical module DDM	The optical module has abnormal parameters.	The device had hardware operational risks.

3.4.2 Periodic Detection Rules

Periodic detection is suitable for problems that require statistical sampling over a period of time before a judgment can be made. It can reduce false alarms caused by short-term jitter and is more suitable for persistent issues such as wireless environment, client quality, port errors, and WAN quality.

Incident Name	Incident Definition	Encountered Issue
Client SNR	Within 15 min, the percentage of	Poor SNR ratio on the

Incident Name	Incident Definition	Encountered Issue
exception	SNR samples below 20 dB is $\geq$ 80%.	client side, which may lead to unstable connection or reduced speed.
Client negotiation rate exception	Within 15 min, the percentage of abnormal negotiation rate samples is $\geq$ 80%.	Low client connection speed, which may cause slow internet browsing.
AP channel utilization too high	Within 15 min, the percentage of channel utilization samples exceeding 70% is $\geq$ 80%.	Busy wireless channel, which may affect access and throughput experience.
AP noise floor too high	Within 15 min, the percentage of noise floor samples exceeding -85 dBm is $\geq$ 80%.	High ambient noise levels in the wireless environment, which may affect coverage and connection quality.
AP wireless interference rate too high	Within 15 min, the percentage of interference rate samples exceeding 50% is $\geq$ 80%.	Significant wireless interference, which may cause retransmissions, reduced speed, or unstable connections.
Client ping-pong roaming	Within 5 min, the client roams between different bands of one AP or different APs 4 or more times.	Frequent switching of connection points on the client, which may cause lag or short-term

Incident Name	Incident Definition	Encountered Issue
		disconnections.
Switch port packet loss rate too high	Within 15 min, 80% of packet loss rate samples of a single port exceed 10%.	Poor port link quality, which may affect the wired connection experience.
Gateway/Switch/AP CPU utilization rate too high	Within 10 min, the percentage of CPU samples at 90% or higher is $\geq 80\%$.	High device resource pressure, which may affect forwarding or management experience.

The focus of these rules is to help users understand why an incident is judged: a problem persists in meeting abnormal conditions within a detection window, or is directly triggered by a clear failure state.

3.5 Incidents and Root Cause Analysis Mechanism

The relationship between incidents and their root causes is not a simple one-to-one relationship. In real-world networks, the following situations may occur:

- An incident may have multiple possible root causes.
- One root cause can lead to multiple incidents.
- Some incidents can also be the root causes of other incidents.
- Some incidents require examining related objects simultaneously to determine the scope of impact and the appropriate course of action.
- Some incidents may initially appear as abnormal phenomena, but users can continue to confirm them by considering subsequent incident status, changes in affected objects, and handling suggestions.

For example, client failed to connect may be related to weak wireless signals, authentication failure, insufficient DHCP addresses, high AP load, or poor uplink quality. The system will attempt to identify possible root causes based on existing data and provide

corresponding suggestions to help users narrow down the troubleshooting scope.

From the user's perspective, the value of root cause analysis lies not in replacing manual judgment, but in reducing unnecessary troubleshooting steps. Users can first review the possible causes the system lists, then confirm them against the on-site network configuration and device status.

3.6 Troubleshooting Workflow and Closed-Loop Verification

Users can access the Health and incident detection functions via the following:

1. Check the Health statistics to determine whether the incident occurs on the device, client, or Wi-Fi.
2. Open the incident list for the corresponding object, and prioritize the Ongoing and Unresolved incidents.
3. Narrow down the investigation by incident category, such as access, authentication, wireless environment, link status, or device status.
4. View the affected objects, first occurrence time, last occurrence time, root cause analysis, and handling suggestions in the incident details.
5. Check the configuration, device status, link status, wireless environment, or client access process as recommended.
6. After processing, observe whether the incident is automatically restored to **Resolved**, or manually mark the processing result according to the actual situation.

The key to this approach is to start with states and incidents users can understand, then gradually drill down to specific metrics and configuration items to avoid getting bogged down in raw data from the outset.

From a closed-loop perspective, Omada organizes the troubleshooting journey into five steps: Identify the problem, locate the target, understand the cause, implement the solution, and confirm recovery. Health status helps identify the problem, incident detection helps locate the target and type of problem, root causes and suggestions help explain the cause and provide a solution, and incident status helps confirm whether the problem is still ongoing or has been resolved.

4 Pro Controller Compatibility

After upgrading from Pro to Standard V6.3, the Omada Controller underwent a deep restructuring of the incident interaction logic and configuration system. To ensure a smooth controller upgrade and configuration compatibility, this section compares the differences between Pro and Standard V6.3 in incident handling, notification mechanisms, and configuration windows, providing a reference for network architects, system administrators, and other users before and after the controller upgrade.

4.1 Changes to the Incident Analysis Mechanism

4.1.1 Changes to Incident Counting Rules

Pro: The tabs for the incident list are based on statistics of incident occurrences.

CATEGORY	INCIDENT NAME	DEVICE	OCCURRENCES	LAST OCCURRED	ACTION
Throughput	PRO L3 port 20/14.2/0/1 traffic was too high	PRO L3	2	Jul 15, 2026 04:15 pm	

Standard V6.3: The system has been adjusted to display statistics based on the number of aggregated incident types.

Status	Last Occurred	First Occurred	Root Cause Analysis	Improvement Suggestions	Affected Devices	Affected Clients	Occurrences	Action
Unresolved	Jul 15, 2026 06:04:49 pm	Jul 15, 2026 04:22:46 am	Interface IP occupied	Please check whether the conflict device or local device SW is using a static IP address.	2 Devices	-	3	


4.1.2 Handling Historical Ongoing Incidents

Pro: Ongoing incidents will automatically resolve once they stop occurring, requiring no manual handling.

CATEGORY	INCIDENT NAME	DEVICE	OCCURRENCES	START TIME	ACTION
Hardware	PRO AP CPU was at high risk	PRO AP	1	Jul 15, 2026 07:00 pm	
Hardware	PRO L3 CPU utilization rate was too high	PRO L3	1	Jul 15, 2026 07:00 pm	

Standard V6.3: After the controller upgrade, historical ongoing incidents inherited from the


original Pro Controller need to be handled manually. Users can batch process these incidents by selecting the corresponding incidents under the **Ongoing** tab in the Standard V6.3 interface. Newly generated ongoing incidents in Standard V6.3 will automatically resolve once they cease to occur.




4.2 Changes to Notifications

4.2.1 Refactored Notification Categories

Pro: Allows modification of Incident Object, Incident Category, and Incident Level for email notifications.




Standard V6.3: The Incident Category has been completely refactored in Standard. After upgrading from Pro to Standard V6.3, Incident Category will be reset to the default configuration. Users need to reconfigure it according to the adjusted categories. Everything else is the same as Pro.




4.2.2 Merging Account-Level Notification Configurations

Pro: In Account Settings, Alert Emails (for Logs) and Incident Notification are isolated configurations.



Standard V6.3: The two configurations mentioned above are integrated into Alert/Event/Incident Notification, and users can adjust the configuration as needed.



4.3 Streamlining and Optimization of Incident Settings

Architecture Evolution: Based on user troubleshooting, the incident categories have been optimized from the original 10 technical descriptions in Pro to 9 more intuitive categories in Standard V6.3.


Incident Attribution Adjustment: For example, a Portal authentication failure that originally belonged to the Access category in Pro was reclassified to the Authentication category

after upgrading to Standard V6.3.


Incident Library Trimming: Standard V6.3 has adjusted the types of incidents supported and removed some redundant incidents from the original Pro version that were not helpful for troubleshooting. Subsequent versions will continue to add and optimize the supported incidents.

4.4 Function Support Status Display

Pro: The device's support for incident detection will be indicated in the Device Name column of the Device List page.




Standard V6.3: The mixed display within the Device List has been removed, and all items are now displayed uniformly on the Configuration Result page.




4.5 Device Detail Health Correlation

Pro: The device's Health data and incidents are presented separately. Stacking incidents together makes it impossible to see their correlation with the Health dimension.



Standard V6.3: Breaking away from isolated displays, it achieves vertical correlation where

incidents follow dimensions. Incidents will be displayed in the corresponding Health dimension according to the incident type, thereby establishing an intuitive mapping between incident data and service metrics.



5 Application Scenarios and Solutions

Omada's intelligent incident detection and Health scoring function, based on a hierarchical diagnostic model of perception-analysis-presentation, provides a practical intelligent operation and maintenance solution for typical scenarios such as chain supermarkets, office networks, and hotels.

5.1 Multi-Branch Chain Store Networks

Scenario Description:

A chain retail brand has multiple stores nationwide, each equipped with a standard network environment consisting of access points (APs), switches, and gateways. Each store lacks dedicated IT personnel; all network operations and maintenance are centrally managed by the headquarters' IT team. The headquarters needs to monitor the real-time network status of each store and respond quickly to any issues that arise.

Operation and Maintenance Pain Points:

1. **Lack of Transparency:** The headquarters cannot know the true Health status of each store's network in real time, and can only passively wait for store employees to report problems such as network lag and inability to process payments.
2. **Low Troubleshooting Efficiency:** When stores report problems, HQ maintenance personnel need to remotely access the controller to check a large number of scattered indicators such as the CPU, memory, ports, and wireless channels of the devices one by one, which is a time-consuming troubleshooting process.

3. **Lack of Guidance:** For inexperienced operations and maintenance personnel, even if they discover an abnormal indicator (e.g., AP channel utilization too high), it is difficult to determine whether the problem affects the business, what the root cause is, and how to solve it.
4. **Difficulty in Problem Closure:** After a problem is resolved, there is a lack of an effective confirmation mechanism, making it impossible to confirm whether the store network has returned to normal, which can easily lead to lingering issues.

Solutions:

1. **Global Health Monitoring:** When operations and maintenance personnel access the controller, they first check the Site Health or Device Health in the Dashboard. This allows for quick identification of which stores' overall network experience or key device (e.g., the gateway) is in a **Poor** or **Fair** state, enabling a shift from passive waiting to proactive discovery.
2. **Automatic Reporting of Incidents:** The system automatically identifies and generates incidents. For example, an AP in a store may experience AP channel utilization too high. The incidents will clearly list the problem type, severity level, and affected devices, helping maintenance personnel quickly pinpoint the problem.
3. **Root Cause Analysis and Suggestions:** Maintenance personnel can click on the incident to view the **Root Cause** and **Suggestions** on the details page. For example, for the "AP channel utilization too high" incident, the suggestion might be "Check if there are high interference sources near the AP, or try adjusting the AP channel". This clear guidance allows even non-experienced maintenance personnel to quickly get started.
4. **Problem Closure Confirmation:** After addressing the issue according to the suggestions, operations and maintenance personnel can see the incident status automatically change to **Resolved** in the incident list, confirming that the issue has been resolved. For issues requiring manual handling, the status can also be manually marked, achieving complete troubleshooting closure.

5.2 High-Density Office/Conference Networks

Scenario Description:

In the office area of a large enterprise during working hours, a large number of employees simultaneously use wired networks (e.g., desktops and IP phones) and wireless networks (e.g., laptops and mobile phones) for daily office work, video conferencing, file transfers, and other tasks. During large meetings or company-wide training sessions, the network access density increases dramatically.

Operation and Maintenance Pain Points:

1. **Unclear Problem of User Experience:** Users report slow network speed, video buffering, or intermittent IP calls, but it is difficult for administrators to determine whether the problem is in the wireless environment (e.g., interference or channel congestion) or the wired network (e.g., packet loss at switch ports, access layer bandwidth bottlenecks, or uplink congestion).
2. **Lengthy and Fragmented Diagnostic Paths:** Administrators need to check both the wireless metrics of the AP (e.g., channel utilization, interference rate, etc.) and the wired metrics of the switch or gateway (e.g., port error packets, packet loss rate, traffic load, CPU utilization, etc.). The information is scattered across different pages, making it difficult to quickly correlate and locate the problem.
3. **Difficult Bottleneck Identification:** When both wireless and wired network metrics fluctuate, it is difficult to determine whether the bottleneck is located on the wireless access side or the wired aggregation/uplink side, resulting in an unclear direction for optimization.
4. **Lack of Targeted Optimization Basis for High-Density Scenarios:** For high-density scenarios, there is often a lack of decision-making based on actual operating data of how to reasonably plan AP deployment, adjust switch port policies, and optimize uplink bandwidth.

Solutions:

1. **Unified Health Monitoring:** Administrators can view the Health status of all devices simultaneously through the Site Health or Device Health in the Dashboard on the

controller. When network performance deteriorates in a certain area, it can quickly identify whether the problem is due to poor AP Health, poor switch/gateway Health, or both simultaneously.

2. **Correlation Analysis of Wired and Wireless Incidents:** When a problem occurs, the administrator first reviews the incident list. For example, an AP generates a "AP channel utilization too high" incident, while the switch port connected to that AP simultaneously generates an "Switch port traffic too high" or "Switch port packet loss rate too high" incident. The system will display both types of incidents together, helping the administrator realize that this may be a problem caused by multiple factors and should not be analyzed in isolation.

3. **Root Cause Pinpointing:**

On the wireless side: Check the Channel Usage and Channel Interference Rate in the AP Health dimension, as well as the wireless client Health, to determine if there is wireless congestion or client signal problems.

On the wired side: Check the Packet Loss Rate, Error Frame Rate, and CPU Usage in the switch Health dimension to determine if there are bottlenecks in the wired access layer or uplink layer.

4. **Precise Optimization and Adjustments:** Based on the guidance of root cause analysis, administrators can take targeted measures. For example:

If the root cause of the problem is **wireless channel congestion**, it is recommended to increase the number of access points, enable higher frequency bands (e.g., 5 GHz/6 GHz), or adjust the radio frequency power.

If the root cause of the problem is **a bottleneck in the wired port**, it is recommended to aggregate critical users or wired clients to ports with higher bandwidth, or enable link aggregation (LACP).

If the root cause of the problem is **insufficient uplink bandwidth**, it is recommended to upgrade the uplink port speed or increase the number of uplinks.

5. **Post-Event Review and Planning:** By reviewing historical events and Health trends, administrators can understand the patterns of network bottlenecks in high-density office scenarios, providing data support for subsequent network expansion, device

updates, and configuration optimization.

5.3 Campus/Hotel Networks

Scenario Description:

Campus networks or hotel networks are characterized by high user mobility, diverse client types, and concentrated access times. Students or guests frequently report problems to the IT department such as "unable to connect to Wi-Fi," "connected to Wi-Fi but unable to access the internet," or "extremely slow internet speed."

Operation and Maintenance Pain Points:

1. **Difficulty in Locating Access Problems:** Many access problems (e.g., connection failures and slow IP acquisition) are difficult to determine whether they are caused by client problems, AP signal problems, DHCP service problems, or authentication server problems. The troubleshooting process is lengthy.
2. **Scattered Information:** Logs and metrics from different components such as DHCP server, authentication server, AP, and gateway are scattered across different pages, requiring administrators to manually link them based on experience, which is inefficient.
3. **Lack of Batch Processing Approach:** When a large number of users report problems, administrators often need to check them one by one, lacking an entry point to quickly locate the problem from a global perspective.
4. **Insufficient Early Warning Capabilities:** For problems that can be prevented in advance, such as "excessive utilization of wireless frequency band channels leading to local paralysis," traditional operation and maintenance methods lack early warning mechanisms and often only deal with them passively after user complaints erupt.

Solutions:

1. **Client Access End-to-End Monitoring:** The Omada system monitors key access and authentication incidents such as Client failed to connect, Client failed to obtain IP, and Client Portal authentication failure. When a large number of users report

problems, administrators can immediately check the Incidents page to see if corresponding events are generated in real time, quickly determining whether the problem is widespread.

2. **Quick Problem Pinpointing:** By viewing detailed information about specific incidents, such as "AP channel utilization too high" or "Client Portal authentication failure", administrators can quickly pinpoint the problem to the wireless air interface radio frequency interference or the authentication server, instead of checking each AP, switch, and gateway one by one, greatly shortening troubleshooting.
3. **Root Cause Analysis and Suggestions:** For the "AP channel utilization too high" incident, root cause analysis will point to "too many surrounding devices" or "APs placed too densely and on the same channel," with suggestions such as "switch channels/frequency bands" or "check for densely placed APs." For "Client Portal authentication failure", root cause analysis will point to "abnormal interaction between the controller and Twilio" or "abnormal interaction between the controller and the RADIUS server," providing administrators with clear troubleshooting directions.
4. **Batch Processing and Prevention:** Administrators can quickly adjust DHCP Server configuration or check authentication server status based on suggestions, resolving all user issues at once. In addition, early warnings can be issued for "AP channel utilization too high" incidents, so problems can be addressed before they escalate, moving network management from reactive troubleshooting to proactive prevention.

6 Future Outlook

As enterprises deepen their digital transformation, network scale and business complexity are growing exponentially. In the future, intelligent network operation and maintenance systems will continue to evolve around three core directions: deterministic positioning, predictive operation and maintenance, and self-healing closed loops.

6.1 Deeper Determination of Root Causes

- **Multidimensional Correlation Algorithm Upgrade:** Industry technology will further

deepen the correlation analysis algorithm that integrates topology, time, and space.

- **Precise Location Within Seconds:** When facing complex end-to-end packet loss problems spanning wireless, wired, and wide area networks, graph neural network (GNN) inference through the entire network event flow graph has become a core research direction, aiming to accurately pinpoint the true root cause with the highest confidence from multiple intersecting factors.

6.2 From Fault Diagnosis to Fault Prediction

- **AI Trend Prediction Models:** Introducing time series prediction and machine learning models is an inevitable path for upgrading network operation and maintenance technology.
- **Network Risk Prediction:** The network Health scoring mechanism is breaking through the limitations of existing historical status summaries and evolving towards in-depth prediction based on device resource consumption and wireless interference patterns. This enables early warning of potential risks such as "AP channel utilization exceeding limits" or "continuous degradation of optical module data indicators," allowing operations and maintenance to stay ahead of network failures.

6.3 Fully Automated Troubleshooting and Self-Healing

- **Deep Integration of Intelligent Recovery:** The ultimate goal of technological development is to completely establish a closed loop across the entire chain of perception-decision-execution-verification.
- **Unattended Operation and Intelligent Self-Healing:** Future technology architectures are trending towards automated decision-making by system. After identifying the root cause, the management platform not only provides handling suggestions but also supports a policy-authorized autopilot mode—where the system automatically performs RF channel optimization, isolates faulty ports, applies corrective configuration, and allocates backup links based on the site environment. After execution, the system can automatically compare service Health, thereby driving the enterprise network to achieve true unattended operation.

7 Appendix

7.1 Glossary

Term	Description
Incidents	Structured problem logs identified and generated by the system, used to describe anomalies in the network.
Ongoing	Incidents that are still occurring and being continuously detected by the controller.
Unresolved	Incidents that have not yet been confirmed and handled.
Resolved	Incidents that have been resolved. They can be manually marked or automatically marked as resolved by the system.
Ignored	Incidents that the user confirms require no further action and actively ignores.
Health	A comprehensive evaluation score of the operational status of devices, clients, Wi-Fi, or sites, calculated through a weighted average of multiple dimensions.
Health Status	Corresponds to a status level, including Good , Fair , Poor , and No Data , which is used to intuitively reflect the current operating status of the object.
Health Dimension	Basic dimensions that make up the Health score, such as CPU utilization, memory utilization, signal strength, SNR ratio, and channel utilization.