# Research on Implementation Path and Risk Control of Large Enterprise Campus Network Upgrade and Renovation Projects

**Jianying Gao***

*Guangzhou Kinte Industrial Co. ,Ltd., Guangzhou, 510860, China*
**Corresponding author: jygao08@139.com*

***Abstract:*** *With the routine deployment of services such as ultra-high-definition video, real-time industrial signaling, and distributed storage, the backplane bandwidth, forwarding capability, and concurrent connection count of traditional campus networks are approaching their design limits, and multi-vendor protocol differences and irregular topologies further exacerbate upgrade complexity. This study focuses on the implementation path and risk control of upgrade and renovation projects, analyzes the driving factors from three aspects—performance degradation, compatibility constraints, and change vulnerability—and proposes a progressive decoupling deployment of control and forwarding planes, a hybrid migration sequence of virtualization and physical devices, and a service-unaware cutover scheduling model; at the risk control level, it defines atomic rollback boundaries, designs a real-time suppression mechanism for flapping convergence, and constructs a fault injection pre-validation method, thereby forming a complete technical framework for upgrade and renovation.*

***Keywords:*** *Campus Network Upgrade; Layered Decoupling; Traffic Cutover; Risk Control; Fault Injection Testing*

## Introduction

After long-term operation and multiple local expansions of large enterprise campus networks, their topologies often exhibit irregular hierarchical characteristics. The redundant link configurations between the core layer, distribution layer, and access layer may contain unreasonable loop dependencies. In addition, the implementation differences among multi-vendor devices in terms of Spanning Tree Protocol extension features, OSPF LSA flooding mechanisms, and Link Aggregation Control Protocol negotiation state machines make network upgrade and renovation projects face significant technical complexity. Studying the implementation path and risk control methods for network upgrade and renovation has important engineering significance and theoretical value. On the one hand, it can provide operable technical solutions for the smooth evolution of campus networks from existing architectures to software-defined and virtualized architectures. On the other hand, by introducing control measures such as atomic operation boundaries, real-time suppression mechanisms, and fault injection pre-validation, it can effectively reduce the risk of service interruption and cascading failures during change operations. This study starts with an analysis of driving constraints, then proposes a layered decoupling implementation path, and finally constructs a closed-loop risk control mechanism, in order to provide technical references for similar network upgrade and renovation projects.

## 1. Analysis of Drivers and Constraints for Large Enterprise Campus Network Upgrade and Renovation

### *1.1 Campus Network Performance Degradation and Service Carrying Capacity Gap*

During the long-term operation of large enterprise campus networks, the backplane bandwidth of the switching architecture and the processing capability of the distributed forwarding engines gradually approach their design upper limits. With the routine deployment of services such as 4K/8K ultra-high-definition video conferencing, real-time industrial control signaling, and large-scale distributed storage synchronization, the end-to-end delay jitter and packet loss rate of the network exhibit nonlinear growth. The concurrent connection count and NAT table entries of the egress gateway

approach saturation, causing some traffic flows to trigger TCP retransmission window oscillation. In wireless access areas, co-channel interference and channel overlap among adjacent access points intensify, and the terminal roaming handover delay exceeds the tolerable threshold for VoIP sessions, manifesting as session establishment failures or audio interruptions[1].

The network quality of service mechanisms cannot maintain deterministic forwarding guarantees when traffic models change abruptly. The original class-based weighted fair queuing shows a significantly increased probability of queue overflow under burst traffic, and the critical traffic flows in the low-latency queue are affected by best-effort traffic bursts, causing priority inversion and instantaneous spikes in buffer occupancy. The aging mechanism for NAT and firewall session table entries fails to match the keepalive period of long-lived connection services, resulting in abnormal reset of connection states. The aforementioned performance degradation phenomena indicate that the service carrying capacity margin of the existing campus network architecture has been approached, and the upgrade and renovation need to introduce higher-dimensional forwarding and control resource scheduling models.

### *1.2 Compatibility Constraints in the Heterogeneous Environment of Multi-Vendor Devices*

Large enterprise campus networks typically include switching, routing, wireless, and security devices from multiple vendors, and each vendor has partial differences in its implementation of protocol standards. In the Layer 2 network, different vendors' devices exhibit inconsistent support for Spanning Tree Protocol extension features, which may lead to discrepancies in port role calculation and abnormal root bridge election. At the Layer 3 routing level, the fine-grained implementation differences in the OSPF LSA flooding mechanism and the link state database aging timer cause routing convergence time to deviate from design expectations. The negotiation state machines of VLAN and Link Aggregation Control Protocol have subtle deviations in the ASICs of different chip vendors, resulting in uneven traffic distribution among member ports within an aggregation group or link flapping.

Compatibility constraints are even more prominent in the control plane and management plane. The structure of the Management Information Base and the return value encoding rules of external interfaces vary across vendors, causing field parsing errors or data missing when the unified network management platform collects performance metrics. The SNMP trap message format and event definition lack a standardized cross-vendor mapping, necessitating additional development of a middleware adaptation layer for alarm normalization. The SDN controller faces the complexity of southbound interface protocol adaptation in a mixed-vendor environment, as vendors have varying degrees of support for the match fields and action sets of OpenFlow tables, and the distribution of globally consistent forwarding tables requires the introduction of shim logic at the abstraction layer. During the upgrade and renovation process, such compatibility constraints require that the implementation path must include cross-vendor interoperability verification[2].

### *1.3 Vulnerability Assessment of the Existing Topology Structure to Changes*

The existing campus network, after multiple local expansions and adjustments, often exhibits irregular hierarchical characteristics in its topology. The redundant link configurations between the core layer, distribution layer, and access layer may contain unreasonable loop dependencies, and there may be inconsistent mapping relationships between the physical topology and the logical topology. The bandwidth utilization of some trunk links shows extremely skewed distribution over the time dimension, with a few backbone links carrying aggregated traffic that exceeds the designed redundancy ratio. During the upgrade and renovation process, operations performed on an intermediate node or link may trigger abnormal recursive lookup of routing next hops, or cause temporary disconnection of the route propagation path of BGP. The probability that single points of failure, which are implicitly present in the topology, are amplified during change operations increases significantly.

Change vulnerability is also reflected in the risk of conflicts between existing configurations and new configurations. The original device configuration files have accumulated a large number of residual Access Control List entries, obsolete static routes, and invalid Network Address Translation mapping rules, and these configuration items may have unintended interactions with the post-upgrade network policies. The route leakage control policies between Virtual Routing and Forwarding instances may be bypassed during topology changes, causing service isolation failure. For Layer 2 networks that employ Multiple Spanning Tree Protocol, the original mapping relationships between instances and

VLANs are difficult to keep completely consistent after device replacement, causing the positions of blocked ports to drift. To assess these vulnerabilities, it is necessary to establish a change impact analysis model based on a configuration dependency graph, so as to identify the cascading failure paths that may be triggered by non-atomic operations[3].

## 2. Implementation Path for Network Upgrade and Renovation Based on Layered Decoupling

### *2.1 Progressive Decoupling Deployment of the Control Plane and Forwarding Plane*

The decoupling of the control plane and forwarding plane is a fundamental operation for achieving elastic resource scheduling in network upgrade and renovation. In hybrid networking scenarios, while retaining the forwarding capability of existing physical devices, control functions such as route computation, label distribution, and topology discovery are migrated to centralized or distributed controllers. This process adopts an incremental replacement strategy, first deploying independent control domains at aggregation nodes with lower geographic priority or lower service priority, and establishing secure channels with forwarding devices through OpenFlow or NETCONF. A state consistency verification mechanism is maintained between the controller's Routing Information Base and the device local routing tables to ensure that no black-hole routing or forwarding loops are generated during decoupling. For legacy devices that do not support standard southbound interfaces, a protocol adaptation gateway is used to convert their original management protocols into unified control primitives, thereby enabling the coexistence and operation of heterogeneous control planes.

Progressive decoupling needs to address the response priority conflicts of forwarding devices when facing multiple control sources. When the decoupling ratio exceeds a certain threshold, the original distributed routing protocols should be gradually degraded or suppressed to avoid timing contention with the flow table entries issued by the controller. A flow table merging logic based on version numbers and timestamps is designed, specifying that entries issued by the controller take precedence over those generated by local routing protocols. In the fully decoupled target state, forwarding devices retain only the capabilities of local link state detection and directly connected route generation, while all other routes are computed and issued by the controller. A fallback mechanism is preserved during the transition period: when the number of heartbeat timeouts between the controller and forwarding devices exceeds a preset threshold, the devices automatically enable the original routing protocol stack to maintain basic connectivity.

Example:

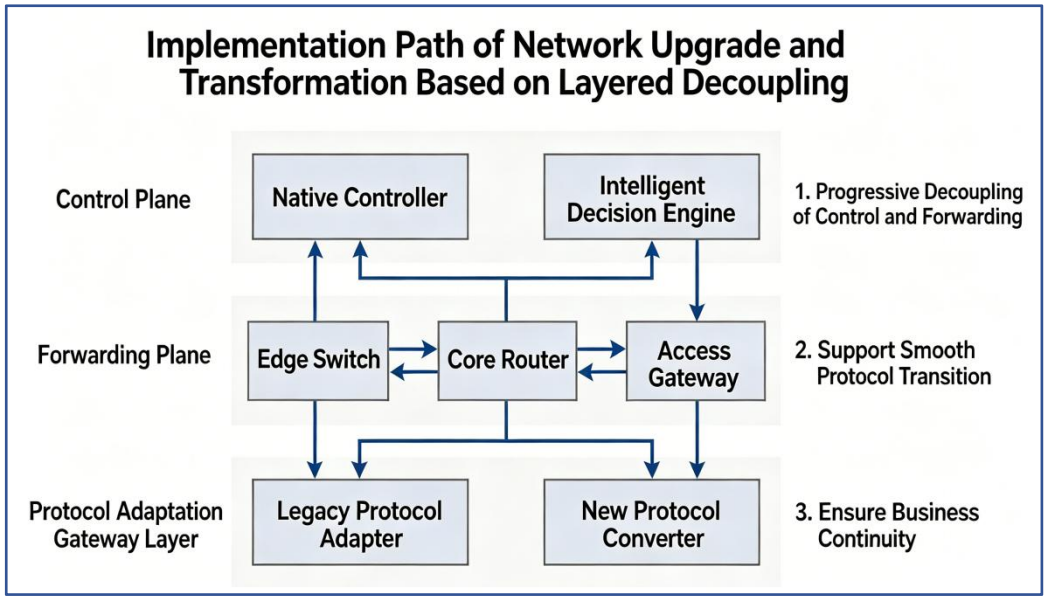


### *2.2 Migration Sequence for the Hybrid Configuration of Virtualized Network Elements and Physical Devices*

Virtualized network elements, as the carriers of software-based network functions, assume roles such as border gateways, firewalls, and load balancers in campus network upgrades. The orchestration

of the migration sequence follows the principle of dependency inversion, that is, service-based network elements with lower dependence on underlying physical links are migrated first, followed by dedicated devices that are tightly bound to specific hardware. The specific sequence is divided into three phases: in the first phase, virtualized route reflectors and virtualized NAT instances are deployed to run in parallel with physical routers, sharing part of the session traffic; in the second phase, access control policies and Quality of Service policies are replicated from physical devices to virtualized network elements, and a policy synchronization mechanism is used to maintain rule consistency between the two systems; in the third phase, service instances on physical devices are gradually reduced, and the traffic ingress is migrated to the virtualized aggregation gateway[4].

The migration sequence in a hybrid configuration needs to address the link identifier mapping problem between virtualized network elements and physical devices. Virtualized network elements run on general-purpose servers or hyper-converged platforms, and their virtual network interfaces need to establish correspondence with the ports and VLAN tags of physical switches. Tunnel encapsulation technologies such as VXLAN are adopted to encapsulate the internal communication traffic of virtualized network elements within the transmission tunnels of the physical network, thereby avoiding broadcast storms caused by the flooding of Layer 2 broadcast domains. During the migration process, bidirectional health probes are maintained between virtualized network elements and physical devices, and the probe packets carry session affinity labels to ensure that the same service flow traverses a path with a consistent five-tuple before and after migration. The devices that have completed migration enter a read-only maintenance mode, with their original configurations retained as the rollback baseline.

### *2.3 Service-Unaware Traffic Cutover Window Scheduling Model*

The scheduling of traffic cutover windows needs to balance the requirements of service continuity and the time constraints of network change operations. A dual-parameter scheduling model based on service priority and delay sensitivity is constructed, in which cutover operations are decomposed into atomic task units, each atomic unit corresponding to one or a set of related flow table entry changes, routing policy adjustments, or link state switches. The scheduler calculates the service perturbation coefficient of each atomic task based on its impact radius, where the impact radius is defined as the number of service flows that may experience packet loss or increased delay after the execution of the task. Atomic tasks with coefficients below a preset threshold can be placed into the near-real-time cutover queue and executed online; tasks with coefficients above the threshold must be assigned to the planned cutover window and executed in batches during off-peak service hours[5].

Example:

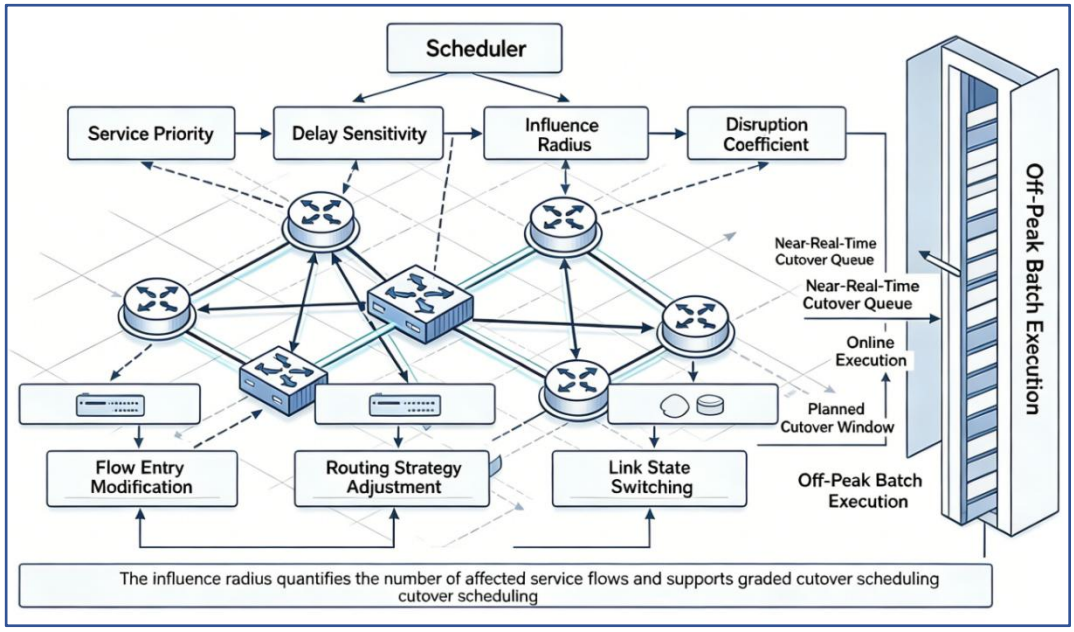


The realization of service-unaware cutover relies on the prefix-independent switchover and Bidirectional Forwarding Detection (BFD) linkage mechanism. For a given link to be cut over, forwarding entries are pre-installed on the backup path first, and after the Forwarding Information Base (FIB) installation is completed, the metric or priority of the original link is reduced hop by hop, so that

traffic flows naturally migrate to the new path. During the entire process, no global Spanning Tree recalculation or full convergence of routing protocols is performed, and a smooth switchover is achieved solely through local metric adjustments. The scheduling model incorporates a rollback timer, which continuously monitors the packet loss rate and round-trip delay of critical service flows within the cutover window. If abnormal indicators exceed the baseline range for three consecutive sampling periods, the timer triggers an automatic rollback, revoking all current and subsequent atomic tasks, and restoring the original forwarding state.

## 3. Risk Identification and Closed-Loop Control in the Process of Network Upgrade and Renovation

### *3.1 Definition of Atomic Operation Boundaries Triggered by Configuration Rollback*

The atomic operation boundaries of configuration rollback determine the reversible granularity and rollback scope of change operations. For configuration changes across devices or across protocol layers, a two-phase commit protocol is adopted to coordinate the operation status of each participating node, where the pre-commit phase completes syntax validation and resource reservation, and the formal commit phase is entered after confirming that all nodes are ready.

The definition of atomic operation boundaries needs to address the asymmetry of configuration elements, that is, the inverse operation mapping between creation-type commands and deletion-type commands. The system maintains a configuration state version difference matrix, which records the incremental fields between the current running configuration and the target configuration, and each incremental field corresponds to a forward operation script and a reverse rollback script. After receiving an abnormality signal, the rollback trigger executes the reverse scripts in the reverse order of the commit timestamps of the atomic units. For atomic units involving routing policies or tunnel establishment, the status of neighbor relationships or the tunnel peer must be verified before reverse execution, so as to avoid secondary failures caused by neighbor restarts or session timeouts during rollback. The atomic boundary metadata is recorded in persistent storage, forming a unique mapping relationship with the change ticket, and supports post-mortem audit and operation review.

### *3.2 Real-Time Suppression Mechanism for Link Flapping and Abnormal Routing Convergence*

Link flapping is characterized by frequent switching of interface physical states or protocol states within a short period of time, which triggers routing protocols to repeatedly compute the shortest path tree and consumes control plane computing resources. For Layer 2 Ethernet links, the timeout parameters of the slow protocol negotiation in Link Aggregation Control Protocol are tuned to merge multiple state jitters within microseconds, thereby reducing the frequency of alarm reporting to upper-layer protocols[6].

The real-time suppression of abnormal routing convergence needs to distinguish between two modes: convergence delay and convergence oscillation. The method to suppress such anomalies is to introduce a route damping algorithm, which increments a penalty value for each route withdrawal event; when the penalty value exceeds the suppress threshold, the route entry enters a hold-down state and temporarily does not participate in forwarding decisions, and is reactivated only after the penalty value decays to the reuse threshold according to the half-life period. The control plane monitors the oscillation states and convergence anomaly logs in real time, and feeds the abnormal events back to the management plane of change operations.

### *3.3 Risk Pre-Validation Method Based on Fault Injection Testing*

Fault injection testing reproduces typical failure modes that may occur during the upgrade and renovation process in a simulated environment, and it is used to identify potential vulnerability points and trigger conditions. The test cases are designed based on the historical fault database and the topology dependency graph, and the injection types include link failures, node failures, configuration errors, control plane overload, and the like. Each injection case defines the pre-failure state baseline, the injection timing and duration, and the expected recovery behavior after fault elimination. Before performing serialized migration operations, a digital twin model is constructed for the devices or links that are about to be changed, and the twin model synchronizes the running configuration and state information in real time to serve as the target carrier for fault injection. During the test execution, the

variation curves of global network performance metrics are recorded, including end-to-end delay, packet loss rate, jitter, and routing convergence time, among others.

The output of the risk pre-validation method is a risk level assessment matrix for change operations. The abnormal behaviors observed after fault injection are classified as risk items, and each risk item is assigned quantitative scores in two dimensions: occurrence probability and impact severity. Before the formal upgrade, the current network state is compared with the most recent successful pre-validation snapshot, and if the difference exceeds the set tolerance threshold, the fault injection test is re-executed.

## Conclusion

This study focuses on the implementation path and risk control for large enterprise campus network upgrade and renovation, and completes the analysis of drivers and constraints from three aspects: performance degradation and service carrying capacity gap, heterogeneous compatibility constraints of multi-vendor devices, and vulnerability of the existing topology to changes. On this basis, it constructs a layered decoupling-based implementation path for upgrade and renovation, which includes progressive decoupling deployment of the control plane and forwarding plane, a hybrid configuration migration sequence of virtualized network elements and physical devices, and a service-unaware traffic cutover window scheduling model. Furthermore, it designs atomic operation boundaries for configuration rollback, real-time suppression mechanisms for link flapping and abnormal routing convergence, and a risk pre-validation method based on fault injection testing, thereby forming a complete technical system covering the phases before, during, and after the upgrade.

Future research directions can focus on the following aspects: machine learning-based mining of historical change data to optimize the parameter adaptive adjustment of atomic operation boundary division and risk scoring models; construction of a digital twin platform for large-scale heterogeneous campus networks to improve the fidelity and coverage of fault injection testing; and research on state consistency and conflict resolution mechanisms among multiple controllers in cross-domain collaborative scenarios.

## References

*[1] Shao Yan, and Qi Jing. "Synergy of Wi-Fi 7 and 10-Gigabit Optical Network Facilitates Campus Network Evolution." Communications World, 14 (2025): 25-27.*
*[2] Xiao Ze. Research on Full-Lifecycle Risk Management of Jinan YK Industrial Park Project. 2025. Inner Mongolia University of Science and Technology, MA thesis.*
*[3] Wang Jiesheng. Research on Schedule Management of YG Company's Campus Network Optimization Project. 2023. Nanjing University of Aeronautics and Astronautics, MA thesis.*
*[4] Ma Pianpian. "Application of Network Simulation Software in Large Campus Three-Tier Network Teaching." Journal of Nanyang Normal University, 19.06 (2020): 59-66.*
*[5] Wang Lei. Design and Implementation of Network System and Network Security for a Digital Campus. 2019. Xi'an University of Architecture and Technology, MA thesis.*
*[6] Xie Wenna. Design and Simulation of Enterprise Campus Network Based on IPv4/IPv6 Dual-Protocol Stack. 2019. Xiamen University, MA thesis.*