DELL Technologies

Whitepaper

iDRAC9: Enabling Telco Transformation

Accelerating automation from the core to the far edge

Accelerate Transformation

Communications service providers (CSPs) are amid rapid transformation. This change is driven by an increasing demand for low-latency, high-bandwidth services from both enterprises and consumers, and enabled by a new generation of advanced technologies.

5G and edge computing have unique network requirements, including the rapid scaling of resources and modular, agile infrastructure. Also, edge deployments have their own set of considerations, such as limited real estate and highly distributed sites with little or no on-site technical support.

To keep pace with this network transformation, CSPs are turning to two key enablers: automation and DevOps mode of operations. While network virtualization and cloud-native architectures have become increasingly important to CSPs, particularly in mobile networks, physical servers remain the bedrock of modern telecommunications networks. Servers have a key role to play in the core, the radio access network (RAN) and OSS/BSS infrastructure.

To streamline and optimize the management of hardware resources in the network, Dell Technologies offers the integrated Dell Remote Access Controller (iDRAC): a powerful tool for remote deployments, monitoring, updating, troubleshooting and re-mediating Dell EMC PowerEdge servers. iDRAC is the first touch point in Dell PowerEdge server provisioning and plays a critical role in the creation of an automated, self-healing telecommunications network.

This paper provides an explanation on iDRAC and how CSPs can leverage its capabilities to build out 5G and edge networks to deploy new services to customers.

Table of contents

The evolving telecom landscape

5G and edge computing are two of the critical focal points for CSPs as they transition into digital service providers, so it's worth looking at each point and how iDRAC9 can efficiently enable and automate the mass deployment of these initiatives. As 5G and edge take center stage, telcos are becoming a convergence point for the needs of all other vertical industries, bringing a new wave of innovation.

iDRAC9 and 5G networks

The mmWave frequency bands for 5G will require antennas to be placed closer to users. This placement will result in the dense distribution of compute infrastructure required to support radio signal processing and device-generated data at these locations. These compute resources will primarily serve the following types of workloads: telco workloads such as virtual RAN (vRAN) and 5G core deployments, and real-time data processing and analysis for artificial intelligence (AI) and machine learning (ML) tools, IoT devices, and augmented/virtual reality (AR/VR) applications.

In the 5G RAN architecture as specified by 3GPP, three main functional modules can be placed in the vRAN 5G radio node (gNB): the radio unit (RU), distributed unit (DU), and centralized unit (CU). These modules can be deployed in various combinations and at different locations (such as cell sites, edge sites). The DU and CU workloads will be deployed on bare-metal servers with each requiring its own server configuration.

iDRAC9 enables efficient, automated, mass deployment of infrastructure. A key feature is network operators can easily generate server profiles based on a wide variety of different workloads, and dynamically apply those profiles in coordination with their network management and orchestration layers. This capability is critical in the deployment of dynamic network slices, an important feature of 5G.

Another consideration for 5G deployments is its requirements for ultra-low latency. To reduce latency, CSPs will need to place compute resources as close to the user as possible, while also employing graphics processing units (GPUs) or field-programmable gate arrays (FPGAs) and SmartNICs to offload data plane traffic and accelerate network performance. iDRAC9 supports this configuration by automatic discovery and inventory of GPU resources and small form-factor pluggable (SFP) transceivers. As Figure 1 illustrates, CSPs can dynamically place and monitor 5G network elements with iDRAC's server profiling and inventory capabilities.

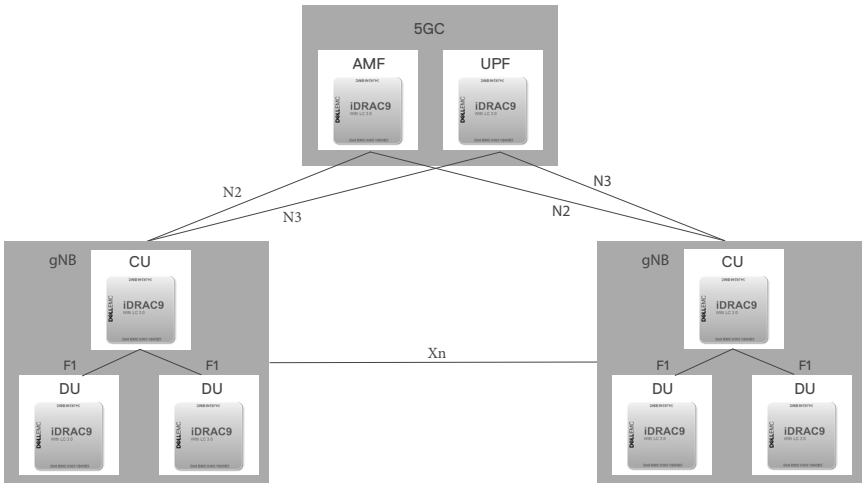

Figure 1 5G architecture and iDRAC9

Network element	Sample iDRAC9 configs support
DU	GPU, SmartNIC based server profiling
CU	CPU, RAID, NIC, BIOS
5G Core	CPU, RAID, NIC, SFPs, BIOS

Table 1 Deployment and configuration sample

iDRAC9 and edge deployments

Edge computing is exactly what its name suggests: the placement of compute resources at the network edge rather than in a centralized data/network center. Edge computing is a standards-based approach defined by ETSI's Mobile Edge Computing (MEC) architecture and 3GPP's definition for edge computing as detailed in TS 23.501 (see Figure 2).

For edge computing to be effective, software must be portable across different server platforms and locations to optimize network performance. Also, these servers must support remote management. iDRAC9 is designed specifically to enable this portability, allowing server administrators to be more productive while optimizing the performance of Dell PowerEdge servers in the network. In the event of a server issue, iDRAC automatically alerts server administrators and allows them to remediate the problem remotely without physically visiting the server.

While CSPs may prefer to use a single, homogeneous hardware platform for their edge deployments, different use cases will often result in a heterogeneous hardware platform. iDRAC9 addresses this potential problem by interoperating with any server hardware platform, using common APIs and interfaces to deliver remote management capabilities across heterogeneous hardware environments. iDRAC9 uses Redfish-compliant RESTful APIs to manage edge computing resources.

Unlike core networks, edge compute resources often face unique real estate and power/cooling restraints. Therefore, they must be optimized for low operational overhead, a small power/cooling footprint, and high reliability. iDRAC9 addresses these requirements with remote management of customizable thermal and airflow settings for PowerEdge servers. With iDRAC9, server administrators can easily monitor, customize, and optimize PCIe airflow and temperature, exhaust control, Delta T control, and overall airflow consumption remotely. In addition, iDRAC9 allows server administrators to easily pre-define power and cooling settings as part of the server configuration profile.

iDRAC9 provides RESTful API in compliance with Redfish 2018 R1/R2 to achieve programmatic management of edge compute resources.

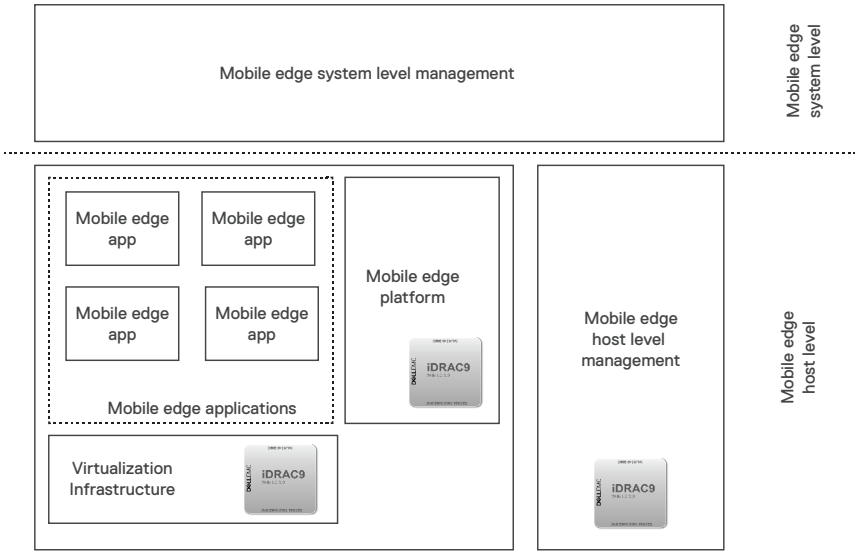

Figure 2 ETSI MEC and iDRAC9

Automation

Automation is an essential part of the next-generation network transformation. iDRAC9 supports it in two very important ways: by automating the server lifecycle and enabling AI-driven network operations. CSPs can choose to create their own automation framework using iDRAC9 or seamlessly integrate it into an existing automation framework.

Server lifecycle automation

iDRAC9 features the auto-discovery of servers in the network via Link Layer Discovery Protocol (LLDP) messages that are sent to the various network switches. Once a server is installed and instantiated in the network, no more manual intervention is required as iDRAC automates the entire server lifecycle from discovery through to configuration (see Figure 3).

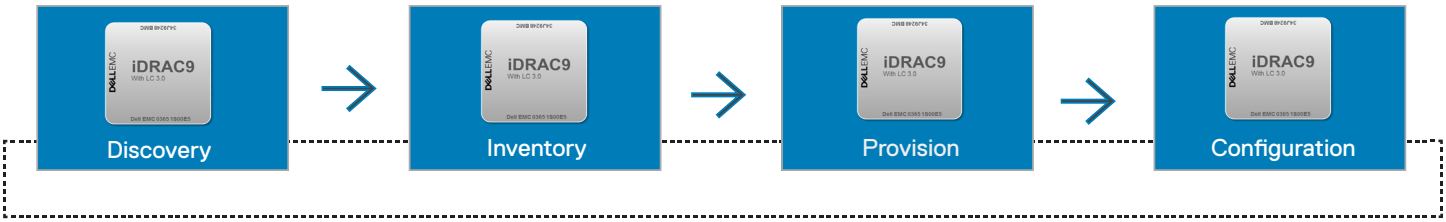

Figure 3 iDRAC9 automation

When a server is powered-on, iDRAC requests and obtains the server's IP address from the Dynamic Host Configuration Protocol (DHCP) server. The DHCP server provides the link to the Server Configuration Profile (SCP), which contains all the parameters that can be provisioned by iDRAC including the BIOS, PERC, RAID and NIC systems of the PowerEdge server. The SCP captures all the server settings in a profile that can be applied across a single server or group of servers, based on their network function. The iDRAC9 auto-configuration feature eliminates manual configuration steps and ensures a version-controlled deployment environment.

In addition to auto-discovery, iDRAC9 prepares a detailed inventory of all server components (including firmware versions) to ensure the correct assignment of workloads. The SCP can be configured with a reference pointer to a network-based firmware repository, so iDRAC9 can update firmware based on the initial profile setup (an important feature for security compliance). Firmware version updates can be automated based on the SCP's configuration, with iDRAC9 comparing deployed and stored versions and, in the event of a mismatch, automatically triggering a firmware update to the new version.

Other server lifecycle automation features that are new in iDRAC9 include:

- An operating system install option in the SCP that supports remote OS deployment, eliminating the need for scripting or external tooling;
- Initial server setup for any network workloads that require the configuration of NIC and/or RAID devices;
- Hardware failure recovery through human-readable snapshots of the entire server system configuration, including easy restoration to previous system configuration states;
- An embedded diagnostic tool for advanced remote troubleshooting;
- Risk mitigation features including real-time BIOS scanning, system erase and alert on USB insertion.

Enabling AI-driven operations

As CSPs look to deliver better network experiences to their subscribers and deliver new revenue-generating services, they are shifting from event-based network monitoring to analytics-based monitoring with the aid of new telemetry-based tools. The new iDRAC9 telemetry streaming feature is one such tool. Designed to provide deeper insights into server performance, iDRAC9 streams real-time telemetry data into any external analytics tool via Rsyslog or Redfish SSE. Metrics captured by iDRAC9 include data on GPU, CPU, memory, serial data logs, and optical network interfaces. Over 180 unique server and peripheral metrics can be streamed or pulled from iDRAC9, with no agent required.
CSPs can leverage these additional metrics by using the Dell EMC Streaming Data Platform or by injecting these real-time data into any external analytics tool. Streaming telemetry data can be used to train machine learning models to support various use cases such as identifying idle servers, power management through intelligent workload placements, and auto-healing. iDRAC9's rich set of metrics can also be used to implement closed-loop automation, where deviation from any standard range will trigger a corrective action on the server (see Figure 4).

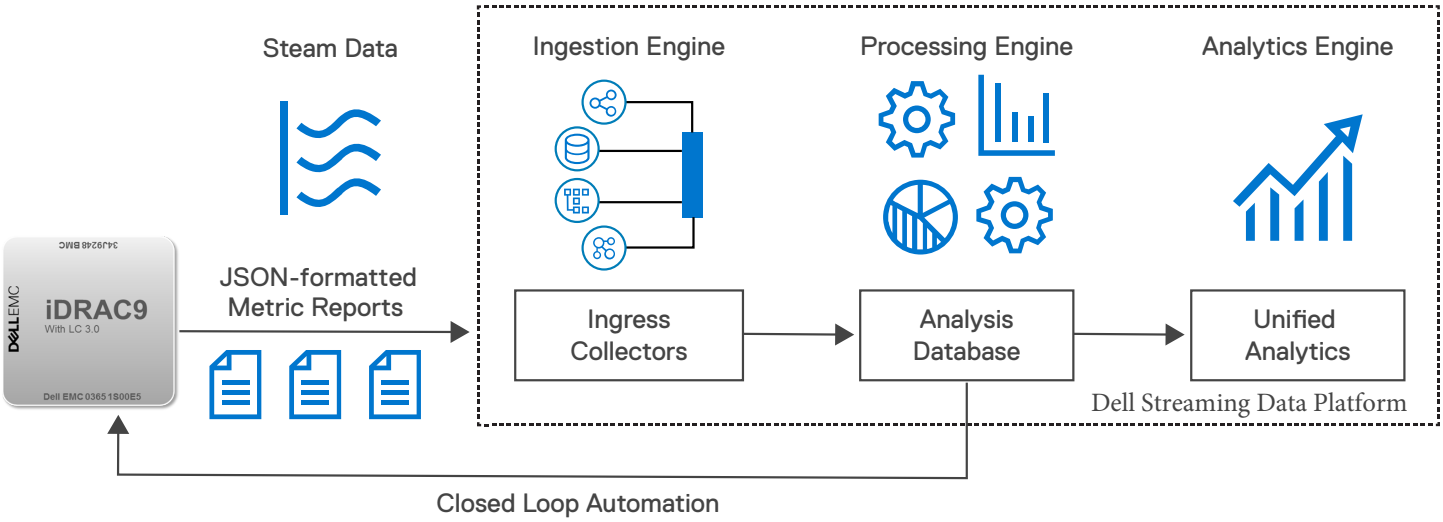

Figure 4 iDRAC9 Streaming Telemetry

Closed-loop automation enhances network operations by enabling dynamic resource allocation, fault decision and prediction, security threat alerting, performance degradation correction, hardware-aware workload placement, idle server detection, and demand prediction. iDRAC9 provides standards-based APIs that can be used to integrate with existing Management and Orchestration (MANO) layers, thus enhancing intelligent operations by sharing server state with the orchestration layers.

Sustaining automation through DevOps

Network Functions Virtualization (NFV)—an important concept in 4G that plays an even more pivotal role in 5G—embraces a microservices approach that includes unique microservices orchestration components. Microservices bring many advantages to 5G network operations including: flexibility/agility, continuous delivery, integration, time-to-market reduction for new services, faster resolution of problems, and a more stable operating environment. Transformation in the application stack will demand equal agility from the NFV infrastructure (NFVi) layer, as CSPs drive toward one-click deployment that will cover both the application and infrastructure stacks.

Infrastructure as code

Orchestration and automation platforms allow for physical infrastructure to be constructed in a version-controlled, immutable manner. In this case, the desired server state is defined declaratively in a template that is then used to build out the network. iDRAC9 supports such templating in the form of an SCP that covers details on the OS deployment, firmware, hardware configuration, and so on—making it a single source of truth for version-controlled deployments across edge locations.

The best practice will be to create a configuration profile for each specific use case (for example RAN-DU/CU, 5G core). The configuration profile can be created from blank templates or exported from a "golden" server configuration by exporting them via iDRAC9's RACADM API or Redfish APIs. Configuration profiles are supported in JSON and XML formats; CSPs can choose whichever format is most convenient. iDRAC9 can also be configured using Redfish API calls. iDRAC's Redfish API supports many GET, POST and PATCH calls that allow for configuration changes.

Agile adoption

CSPs want fast, reliable, and repeatable outcomes. They also want an on-demand runtime environment that is consistent across deployments and locations.

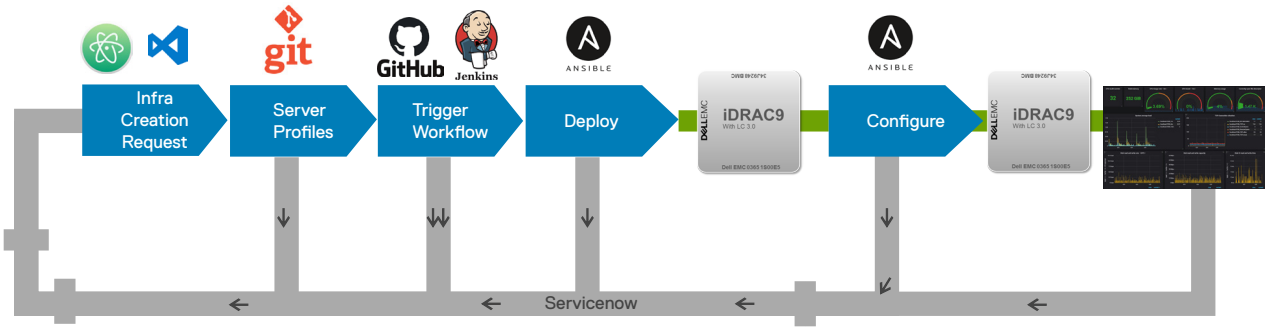

Figure 5 Integrating iDRAC9 with CI/CD tooling

iDRAC9's open, standards-based APIs can be easily integrated with any continuous improvement/development (CI/CD) process (see Figure 5). SCPs can be stored in a centralized, version-controlled management system (such as Git), and various artifacts including OS and firmware can be part of the repository (such as Bitbucket, Nexus). Any update or change in the SCP will trigger the server deployment workflow application (such as Jenkins Workflow, CircleCI). iDRAC9, through its APIs, can easily be managed by any configuration management tool, such as Ansible. The server's end state can then be integrated into monitoring tooling (such as Grafana, Prometheus) to certify that all server components are in a healthy state and ready for workload deployments.

Summary

Automation has become the new byword for business agility. As CSPs look to deploy and accelerate 5G and edge computing rollouts, network infrastructure and service automation will be an integral part of their success.

Beyond automation, real-time monitoring of infrastructure resources will be critical to ensure 5G services meet the high expectations of subscribers.

The latest version of iDRAC—iDRAC9 v4 with data center license— improves upon Dell Technologies' legacy of server automation with new features designed to better automate remote server configuration, reduce errors, maximize server uptime and optimize server performance. With iDRAC9, CSPs can count on fast, reliable, repeatable rollouts of 5G and edge computing services.

FY21Q1_1499_idrac9_white_paper_041320

DELL Technologies