Perspective

Deutsche Telekom's Horizontal TelCo Cloud (HTC) establishes a new industry blueprint for telco clouds

February 2026

Gorkem Yigit and Joseph Attwood

Contents

List of figures

This perspective was commissioned by Deutsche Telekom, Amdocs, HPE Juniper Networking, Mavenir and Nokia. Usage is subject to our disclaimer and copyright notice. Analysys Mason does not endorse any of the products or services mentioned in this paper.

1. Executive summary

Deutsche Telekom's (DT's) Horizontal TelCo Cloud (HTC) marks a shift to a disaggregated, horizontal cloud platform for fixed and mobile core networks.[1] HTC reflects DT's investment in cloudification and automation across its core networks over several years and is guided by earlier initiatives such as the operator's next-generation IP multimedia subsystem (NIMS) programme in 2018. HTC represents a large-scale, industry-leading implementation of a common cloud-native platform and unified network automation.

This perspective details DT's implementation of HTC, and explores the motivations and components of the HTC architecture, the outcomes the project has delivered and the key learnings and best practices for other operators. It positions the HTC project as a blueprint that DT's peers can use for their own network transformations, while also demonstrating what is possible to achieve with the right investment in technology, people and processes.

1.1 HTC represents a foundational transformation in how DT deploys and operates its fixed and mobile core networks

DT initiated the HTC project to address limitations in the agility, quality and resilience of its networks, in line with its ambition to become the 'Leading Digital Telco' by 2030. To support this ambition, DT wanted to adopt cloud-native network functions (CNFs) and cloud-native ways of automating its network, that is, using declarative, intent-based approaches based on the use of GitOps and Kubernetes (K8s). As such, the HTC project focused on breaking down operational silos and moving towards unified platforms, tooling and processes for cloud infrastructure and network automation. The aim of this was to improve efficiency by optimising the time and resources allocated to the systems, platforms and frameworks underpinning DT's networks and operations.

HTC is built around a layered architecture, as illustrated in Figure 1. Network functions (NFs) and services sit between two shared horizontal layers: a common cloud infrastructure layer and a unified automation layer. This replaces monolithic network stacks with reusable platforms; the two shared horizontal layers support multiple NFs, thus allowing NFs to be added, changed or scaled independently without needing to rebuild the entire stack.

1 This perspective describes what DT has done with its operations in Germany. It will apply a similar approach to its operations in the rest of Europe in the future.

Figure 1: Overview of the layered architecture of HTC

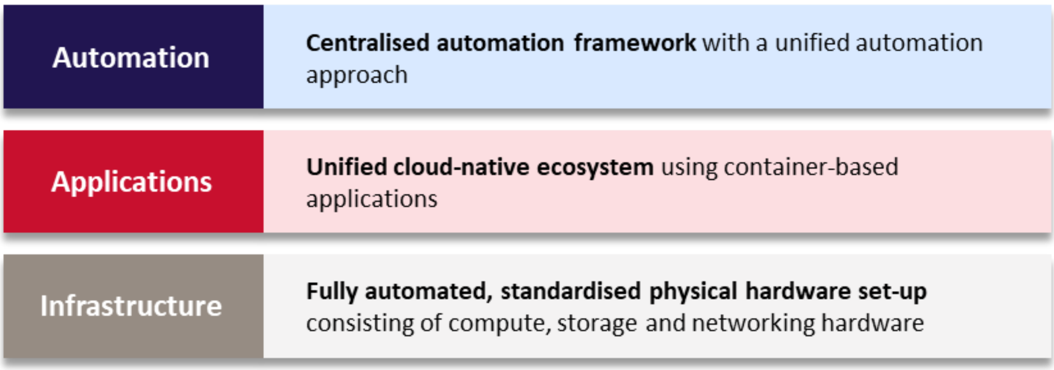

Source: Analysys Mason

HTC is part of DT's wider pursuit of greater architecture harmonisation across both its European footprint and its technology areas (access, transport and core networks, as well as fixed and wireless networks) to improve scale and efficiency. The cloudification and high level of automation of core networks is a key component of this pursuit.

1.2 HTC's success required a focus on standardisation, radical end-to-end automation, organisational transformation and the development of in-house capabilities

The HTC project required DT to overcome significant architectural, organisational and ecosystem challenges, while also making changes to its operating model, organisational structure and mindset. DT chose to take on the role of prime integrator and invest in in-house development because off-the-shelf solutions often struggled to meet its specific operational and architectural requirements and did not provide enough control and flexibility.

The main aspects of HTC's implementation were as follows.

- **Radical end-to-end automation** to support Day 0, 1 and 2 operations with a unified automation platform.[2]
- **Organisational restructuring** to break down the siloed teams that were responsible for the full network stack into horizontal teams that focus on just a single layer of DT's network infrastructure. These teams provide common components that are consumed by other layers of DT's infrastructure.
- **System integration ownership** to enable DT to take responsibility for the successful integration of network components, rather than relying on vendors to deliver fully pre-integrated and/or pre-certified solutions. This allowed DT to ensure the smooth interworking of components and develop expertise that would later support its network operations.
- **The adoption of DevOps and governance** to ensure that the pieces of HTC fitted neatly together.
- **In-house development** to create the cloud infrastructure layer of HTC and the automation framework used by HTC using the skills and expertise that DT had accumulated in telco cloud and automation over the

2 Day 0 refers to the planning and design process before the deployment of systems. Day 1 refers to the provisioning and deployment of systems. Day 2 refers to the scaling and maintenance of systems.

years. DT took overall responsibility for designing and implementing its automation layer, though this layer consists of a combination of in-house-developed, open-source and vendor proprietary components.

- **Vendor ecosystem alignment** to ensure that the 40 vendor partners that supported the HTC project were in alignment with DT's vision for this project and delivered solutions that were interoperable with the HTC architecture.

1.3 HTC has delivered significant benefits to DT, especially in regards to the ability to quickly deploy new software onto networks

HTC has reduced the complexity of network operations, increased the speed of deploying software upgrades and security patches, and improved the resilience of DT's core networks by reducing human error in network operations and allowing fixes to be introduced more quickly.

Some of the metrics for success that DT has been able to measure are as follows.

- The time to validate and roll out new software for voice applications has been reduced from **6–9 months to 7 days**.
- New software for voice applications can be deployed onto networks **at least 12 times per year**, as opposed to a **single software upgrade every 1 or 2 years** previously.
- HTC has enabled in-service software upgrades (ISSUs) in DT's packet core network, without needing to take off traffic from the cluster. This has accelerated network roll-out from **several weeks to several days**.
- Multiple packet core feature sets can be introduced **in parallel**. This is enabled by robust Git branching and version control as opposed to sequential work in traditional environments.
- Network fabric software upgrades are now fully automated; they take only **a few days** to implement, down from **4 weeks**.
- DT is able to keep up with the release cadence of K8s (three releases per year).
- DT can have cloud infrastructure ready for consumption within **40 days** of ordering.

1.4 DT's HTC project offers a range of learnings for the telecoms industry

The telecoms industry is increasingly recognising the importance of cloud-native technologies and horizontal network cloud architectures. The open-source project Sylva is illustrative of a growing interest in the latter.[3] The idea behind Sylva, which shares many similarities with HTC, is to develop a common cloud architecture, harmonised across many operators, to ease the integration of NFs in a cloud-native way.

Given this backdrop, Analysys Mason and DT have produced this paper with the aim of encouraging other operators to start and/or accelerate their own journeys towards cloud-native, horizontal and highly automated networks.

---

3 Sylva is a Linux Foundation Europe project that was founded in 2022 as a collaboration between operators in Europe, including DT, and vendors to develop a common, open network cloud software framework. Linux Foundation Europe, *Sylva*.

The main takeaways for operators include:

- beginning the horizontal cloud journey as soon as possible to ensure that that they drive the direction rather than waiting on their vendors to set it
- capturing learnings and best practices from smaller-scale initial projects and leading operators' transformation efforts to validate architectural choices, operating models and automation approaches before scaling them across the network
- clearly defining the target architecture and operating model, and executing against this vision with strong cross-organisational governance and discipline to ensure compliance
- addressing organisational mindset and structure changes early, recognising that these are often among the most challenging aspects of the transformation
- setting and enforcing strict expectations for vendors to ensure alignment with the operator's long-term network vision.

2. Background to HTC

2.1 The telecoms industry is moving towards cloud-native networks, and horizontal architectures are set to play a key role

Telecoms operators are increasingly progressing with the cloudification of their networks. Traditionally, physical networks, as well as early cloudified networks, existed as operational silos. Vendors provided hardware, software and management solutions as closed, vertically integrated stacks that were managed end-to-end by a single team within operators with specialised knowledge. Bridging the gaps between these silos was challenging, and solutions built in house or provided by third-party vendors could not be easily applied. This resulted in inefficiencies in operations, limited automation capabilities and slower innovation.

Moving towards more disaggregated and open networks gives operators the opportunity to tear down these operational silos. This has provided operators with a greater ability to mix and match components from multiple vendors and components built in house. Additionally, cloudified networks enable operators to adopt horizontal architectures whereby a set of common components and processes are used to support various network workloads/applications. A consequence of this is that the teams responsible for NFs can focus on their primary task because they are supported by easy-to-use platforms that provide capabilities for resource management and NF lifecycle management.

There is a growing industry consensus that using a horizontal network cloud architecture is the 'right' way to operate a network because it enables operators to take greater control and reduce their reliance on any single vendor, thus resulting in more rapid innovation, improved operational agility and customer-focused network excellence. We forecast that the adoption of a horizontal cloud approach will outpace that of vertically integrated stacks in the 5G mobile core; spending on the former will grow at a CAGR of 16% between 2026 and 2030, while that for the latter will grow at a CAGR of 8%.[4]

Cloudified applications are also evolving as operators shift from virtualised NFs (VNFs) to CNFs that run as containers on K8s-based cloud platforms. Operators expect CNFs to conform to the core principles of cloud-

[4] See Analysys Mason's *Network cloud infrastructure: worldwide forecast 2024–2030*.

native architectures.[5] The decomposition of NFs into microservices and the ability to use cloud-native approaches for the architecture and lifecycle management of applications enable operators to achieve operational agility and efficiency benefits. Analysys Mason estimates that spending on CNFs as a percentage of total spending on mobile core NFs will grow from 36% in 2025 to 76% in 2030. As operators transition to CNFs, they also have the opportunity to implement more cloud-native ways of automating their networks (for example, by using GitOps and K8s's orchestration capabilities to manage NFs).[6] Such an approach is fundamental to unlocking the operational efficiency benefits promised by CNFs. Additionally, GitOps allows each change to be tracked, reviewed and audited via Git and subjected to automated tests. Subsequently, the occurrence of incidents caused by manual changes is significantly less likely.

Other advantages of a move to CNFs include the ability to scale systems dynamically with load and being able to deploy updates more frequently as operators move from making bulk updates to making small, more frequent improvements. Additionally, configuration environments become more consistent and deterministic, which simplifies the recovery of systems after a failure, for example.

2.2 DT sought to evolve its network architecture to improve agility and efficiency, while progressing towards greater network harmonisation

Telecoms networks face strict requirements in terms of performance, uptime and security. DT identified that it would be challenging to better fulfil these requirements if it continued to rely on its legacy approaches. For example, the use of physical NFs (PNFs) and VNFs limits the pace at which changes can be made to networks, thereby restricting the speed at which performance and security issues can be addressed. Furthermore, the use of operational silos increases the complexity of managing and optimising networks.

As a result, DT invested in evolving its networks in line with the trends explained in the previous section. As part of this transformation, DT aimed to transition towards a more flexible and adaptable network infrastructure that would allow it to:

- quickly validate and deploy the latest features onto its network following a new software release from a vendor or internal team in a matter of days rather than months
- deploy security patches as soon as they become available in order to withstand the growing frequency and intensity of attacks that are targeting telecoms networks
- make changes to the network without creating network downtime or otherwise lowering the quality of service (QoS).

To achieve these feats, DT aimed to implement more automated update procedures. The desire for higher levels of automation extended past Day 0 operations to include Day 1 and Day 2+ operations.

DT was one of the first operators to rethink the network and telco cloud and commit to breaking down silos by moving to a horizontal network cloud operating model. DT recognised the efficiencies that this could create by reducing the duplication of effort, tools and processes. Additionally, this change was in line with DT's strategic ambition to increase the harmonisation of its architecture across its footprint. This ambition means that DT

---

5 According to the Cloud Native Computing Foundation's (CNCF's) definition, this means that CNFs should also use approaches such as service meshes, microservices, immutable infrastructure and declarative APIs. They should be loosely coupled, resilient, manageable and observable, and be supported by robust automation and orchestration. Linux Foundation, *Cloud Native Computing Foundation*.

6 GitOps approaches make use of tooling and concepts from modern software development. A Git repository is used as a single source of truth for the deployment and configuration of applications, as well as for additional K8s manifests.

encourages the use of common platforms and components across its various European national companies (natcos).

2.3 DT first set out what a cloudified network should look like with its NIMS project and has applied learnings from NIMS to its wider network transformation

DT began its next-generation IP multimedia subsystem (NIMS) project in Germany in 2018. In this project, DT applied a highly automated architecture consisting of virtualised applications running on a horizontal network cloud platform to transform its IP multimedia subsystem (IMS) fixed-line services.[7] The NIMS project was noteworthy for:

- using an open, vendor-agnostic, disaggregated and horizontal telco cloud environment
- adopting a radical end-to-end automation approach based on DevOps automation principles that extended automation far beyond existing levels
- involving extensive collaboration and co-creation with vendors.

NIMS allowed DT to achieve game-changing levels of automation and prove the viability of a horizontal, virtualised network architecture. It also meant that DT gained crucial expertise relating to horizontal network cloud architectures and the implementation of the organisational changes needed for these architectures. DT was satisfied with the high-level architecture adopted for NIMS and the associated operational benefits, but it recognised that its future network transformation efforts would have to diverge from those involved in the NIMS project in certain ways.

- DT wanted to move to CNFs rather than continuing to work with VNFs (as were used in NIMS), and subsequently adopt new automation approaches that are more suited for cloud-native networks than ETSI NFV MANO is.[8]

- DT identified the need to replace OpenStack with K8s at the infrastructure layer; this would allow DT to support CNFs while also improving the upgrade processes for live networks and operations (because roll-outs of new OpenStack versions, as well as the maintenance of these versions, was a cumbersome process).

- DT wanted to be able to orchestrate networks with less reliance on a vendor's professional services.

- DT believed that it would struggle to scale its automation framework if it continued to rely heavily on vendors for development.

- DT felt that serving as its own prime integrator would allow it to develop knowledge and expertise that would help it to operate its networks.

2.4 With HTC, DT moved to deploying cloud-native networks based on unified infrastructure and automation platforms

DT aim to use its HTC architecture to further the speed and efficiency benefits that it achieved with NIMS to the rest of its fixed and mobile core networks. HTC used common platforms, tools and processes at the infrastructure and automation layers for all of its core network workloads (see section 3) to reduce the effort

---

7 See Analysys Mason's *A move to cloud changes the game for Deutsche Telekom's next-generation IMS*.

8 ETSI NFV MANO stands for European Telecommunications Standards Institute network functions virtualisation management and orchestration.

required to achieve this. Changes were also made to DT's organisational structure in order to break down silos. Previously, a single team was responsible for both a given workload and providing the infrastructure and automation for the workload. With HTC, each of these responsibilities is now assigned to different, horizontal teams (see section 4).

With HTC, DT stopped onboarding new VNFs and moved to CNFs. Additionally, it set a mandate to use cloud-native methodologies to automate its networks, through the use of GitOps, for example. However, DT believed that none of the commercially available cloud platforms could support its need to run containerised applications in its data centres while meeting its requirements for scale, performance, privacy and resiliency. As such, it decided to build the cloud infrastructure layer of HTC in house from scratch, drawing on open-source technologies where appropriate. Similarly, DT developed a custom automation and observability solution using a combination of in-house-built, vendor proprietary and open-source solutions.

DT played the role of the platform engineer and prime integrator for the project. It had the capabilities to do this work in house thanks to the experience it had gained working on NIMS.

HTC was designed to provide DT with a futureproof architecture that is capable of supporting the best possible network experience for its customers while reducing the complexity of network operations (see section 5). Additionally, DT wanted to adopt a radical end-to-end automation approach for its HTC-based networks. Figure 2 shows DT's timeline for the HTC project.

Figure 2: HTC project timeline

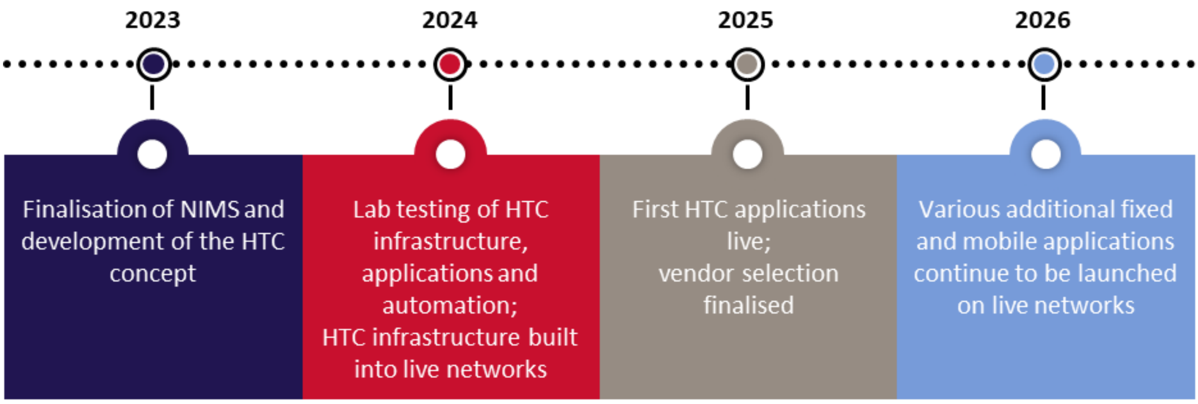

Source: Analysys Mason

3. The HTC architecture

DT internally refers to the HTC architecture as its 'burger' architecture. A common cloud infrastructure layer sits at the bottom of the stack; this includes a platform for providing bare-metal infrastructure ('connected bare metal for everything as a service (BM4X)) as well as a platform that provides K8s clusters on top of bare-metal infrastructure ('Telekom Containers as a Service (T-CaaS)'). BM4X and T-CaaS act as a hardware factory and a software factory, respectively, within the cloud infrastructure, and together they form the foundation of the HTC architecture. Cloud-native applications provided by an array of NF vendors sit on top of this. There is then a common automation and observability platform ('Automation Core') for application lifecycle management at the top of the stack.

The architecture also includes connecting elements that stitch together the infrastructure, applications and automation layers. These elements facilitate the onboarding of applications onto the cloud infrastructure and integration with the automation platform. The high-level architecture of HTC is shown in Figure 3 and the following sections describe each layer in more detail.

Figure 3: Overview of the HTC architecture

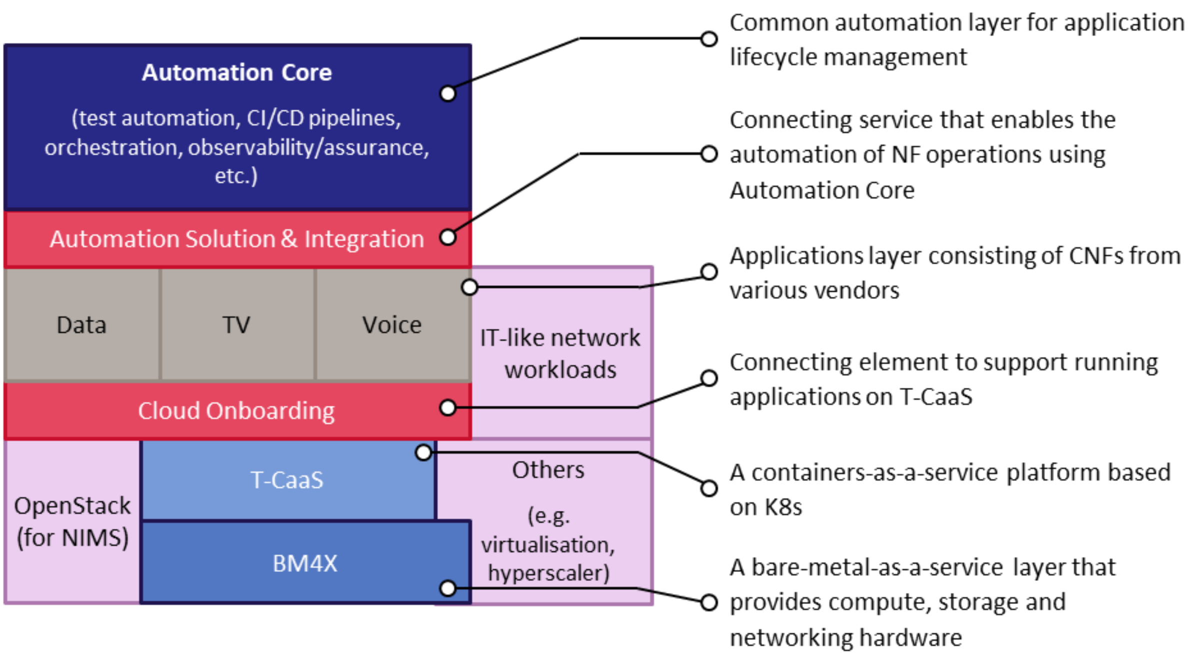

Source: Analysys Mason and Deutsche Telekom

DT has adopted three pillars as the architectural foundations of its HTC project: **one unified infrastructure**, **one unified cloud-native ecosystem** and **a centralised automation platform**. As well as applying these three architectural pillars, DT adopted the following approaches for the HTC.

- **One consolidated network environment.** DT uses a single environment (HTC) to run its various core networks (mobile, fixed residential and enterprise).
- **Cloud-native technologies.** DT expects its NF vendors to deliver their solutions as CNFs and continually pushes these vendors to ensure that cloud-native principles (such as the twelve-factor app methodology[9]) are applied to their solutions. Additionally, DT is taking a cloud-native approach to automation, that is, it is using declarative, intent-based automation based K8s's inherent orchestration capabilities, K8s operators and GitOps.
- **Radical end-to-end automation.** This uncompromising, automation-first philosophy was established by DT and is embedded in the HTC architecture. Applications are provided with a comprehensive set of automation capabilities from the outset.

9 For more information, see *The twelve-factor app*.

- **Modularity of tooling.** The HTC architecture adopts a common, standardised toolset model and allows for tooling to be replaced if a better alternative is found because the architecture describes desired functionality rather than specific tools. DT aims to gradually consolidate its tooling over time.

- **Open-source components.** DT aims to build its network using open-source and industry-standard software as much as possible.

- **Common enablers.** DT aims to offer software components to internal teams as common, 'as-a-service' solutions as much as possible, minimising the need for individual teams to build and manage these components themselves. Instead the teams may consume (preferably through a self-service model) tooling that is pre-integrated into DT's platform.

- **Zero-trust network.** DT is using a zero-trust architecture for its network, which means removing implicit trust from the network and shifting responsibility for security to individual applications teams. This approach is well aligned with cloud-native principles and helps to achieve resilience across the network.

3.1 BM4X delivers standardised compute, storage and networking hardware with automated ordering, inventory management and provisioning

DT introduced BM4X (pronounced 'b-max') to solve inefficiencies with setting up data centre infrastructure. It provides API-consumable compute, storage and networking hardware in more than 10 data centres in Germany (BM4X is currently only available in Germany). BM4X increases infrastructure standardisation to enable DT to simplify and automate the management of infrastructure for thousands (and moving to tens of thousands) of servers, and significantly reduce the time taken to deploy new servers. DT's work on BM4X introduced standardised rack profiles, floorspace designs/rack layouts, cable management, and power and cooling solutions. Additionally, BM4X automates many of the processes involved in the delivery of this infrastructure.

BM4X supports a comprehensive set of infrastructure lifecycle management processes, including documentation, procurement, centralised inventory and operations. BM4X allows infrastructure to be ordered at a press of a button as soon as the specification is received. Rack and cabling plans are generated automatically to simplify the physical installation of hardware. BM4X then enables this hardware to be provisioned with zero touches from engineers; this includes installing firmware, configuring the hardware and managing IP addresses. This automation produces pre-configured data centre infrastructure that is ready for use by the T-CaaS platform (T-CaaS is the first and main customer of BM4X, though BM4X can support other use cases as well). Infrastructure inventory management systems are updated automatically after data centre infrastructure has been delivered.

3.2 T-CaaS enables DT to deploy cloud-native applications on bare-metal infrastructure

T-CaaS is a K8s-based platform developed by DT that sits at the core of the operator's horizontal cloud architecture and hosts its multi-vendor core network workloads using BM4X infrastructure. T-CaaS provides managed K8s clusters for running CNFs at scale. It brings strong multi-cluster fleet management via a common framework and standard YAML manifests, thus enabling the consistent creation and orchestration of K8s clusters using GitOps tools such as Flux CD and Cluster API. It provides DT's application teams with standardised, integrated and security-hardened clusters. DT adopted zero-trust networking (ZTN) as a design principle for T-CaaS workloads.

DT developed host-based routing for T-CaaS to reduce network complexity

Every server in a telco cloud deployment (abstracted as a node in a T-CaaS cluster) needs connectivity; this is provided through hundreds of virtual routing and forwarding (VRF) instances. However, the role and network profile of each server/node changes frequently. Servers/nodes are reprovisioned and repurposed between clusters during rolling upgrades or as the resources assigned to particular workloads are scaled in/out or in response to failures.

Traditionally, a network fabric would use static VLAN configurations to specify how data centre switches should communicate with K8s nodes/servers. However, because servers are frequently reassigned and workloads may move between them, the management of static configurations quickly becomes impractical. Network management done in this way requires an external orchestrator to continually reconfigure the network, resulting in a highly complex and fragile network fabric.

To overcome this issue, DT developed a solution called host-based routing (HBR). HBR enables DT to solve the connectivity challenge in an efficient and automated way in a large-scale environment. HBR deploys a containerised host router on each node; this router instance is managed by a K8s operator in a declarative manner. The node can then directly peer with Ethernet virtual private network (EVPN)/virtual extensible LAN (VXLAN)-compatible networks thanks to the border gateway protocol (BGP) control plane capabilities of the containerised router. This approach means that K8s's orchestration capabilities are used to deliver Level 3 routing to workloads/applications; all connections can be configured solely within the K8s cluster, without touching the network. The BM4X fabric provides seamless integration support for HBR connectivity.

DT has donated HBR to become part of the Sylva architecture to share it with the telecoms industry. DT now contributes to HBR upstream and consumes it back from the open-source project. DT has also successfully partnered with a third-party company, a virtualised and secure networking software developer, to integrate its vSR enterprise-grade router software into DT's HBR deployment and thus add a service-level agreement (SLA) component to it.

3.3 DT has made strong progress in transitioning its core networks from legacy applications to cloud-native applications running on T-CaaS

The applications layer of the HTC architecture consists of CNFs from vendors such as Amdocs, Mavenir and Nokia that run on top of T-CaaS.

This horizontal, vendor- and domain-agnostic HTC architecture provides DT with reusable, standardised building blocks. These can be composed and instantiated quickly to deliver existing and future data and voice services (mobile, fixed residential and enterprise) on a single platform with logical separation. DT has already migrated its 5G core, as well as several other data, voice and messaging applications, to the HTC architecture and has concrete migration plans for nearly all other workloads. DT expects to complete the full migration of applications to the HTC architecture by 2027/2028. HTC currently supports nearly 40 workloads (using approximately 200 000 vCPUs) and is expected to support nearly 80 workloads by the end of 2028.

DT's approach has created challenges related to sourcing and onboarding applications. DT was one of the first operators to adopt CNFs at scale, so it often found that its suppliers lacked readiness. DT reports that it sometimes still faces challenges getting NFs that strictly embody the properties expected of cloud-native software. However, DT also says that these NFs still deliver significant benefits over VNFs if they are fully integrated into the HTC architecture, even if they are not fully cloud-native (for example, due to the need to comply with standards and the use of point-to-point protocols). Some fine-tuning of solutions from vendors has been required to get NFs working on DT's cloud and a balance must be struck between carrier-grade requirements such as latency, performance and resiliency and cloud-native features.

Integrating CNFs into DT's Automation Core has also presented challenges. Even though DT's automation environment is largely based on common tooling, suppliers have historically preferred to use their proprietary and specific management tools for all operational tasks. DT required vendors to provide NFs that were decoupled from their element management systems (EMSs) to support onboarding CNFs into the generic architecture of HTC. Vendors also needed to support changes to NF lifecycle management such as the use of Git and declarative approaches.

Nonetheless, DT has progressively migrated more of its core network to the HTC architecture and established HTC as the target architecture for all future core network investment.

3.4 HTC's Automation Core provides a common set of pre-integrated tools that unify NF lifecycle management

HTC's Automation Core provides centralised and standardised automation and observability capabilities for NFs. Traditionally, applications used various automation tooling and processes that were often unique to each operational silo. Automation Core replaced these siloed, 'snowflake' automations with a common framework that can be used for all applications and platform teams. This ensures an automation-first approach, where automation based on a common framework is available from the outset; this is opposed to automation being developed separately and in parallel for each function and infrastructure component, which previously resulted in fragmentation.

DT designed its overarching automation architecture by identifying the set of functionalities that it needed and then implementing these functionalities using a combination of in-house-developed, open-source and commercial solutions. This was carried out via collaboration between the automation, architecture and DevOps solutions teams. The approach of describing the functionality rather than the tool and defining standardised points of integration allows tooling for a given capability to be exchanged if a better alternative is found. A centralised governance team ensures that the constituent tools of Automation Core can be integrated together and that they work as an end-to-end service for application teams. DT is continually evolving Automation Core in response to technology advances and common operational demands of NFs.

As discussed in section 3.3, DT moved to rely on the common, standardised and pre-integrated toolsets provided by HTC's Automation Core instead of vendor-proprietary automation and observability solutions to eliminate function- and vendor-specific silos. This required DT to push its NF vendors to provide EMS-less solutions with open interfaces that align with DT's automation roadmap. DT initially faced resistance both internally within its application teams and externally from its vendors, and several workarounds and adjustments were needed to address constraints. Nonetheless, the transition proved to be successful for DT's automation ambitions and also enabled reductions in hardware footprint, resource usage and maintenance requirements.

HTC's automation is now fully GitOps-based using DT's Operator-based Cloud Automation Solution (OCAS). OCAS was developed in house but uses open interfaces, and the integration with external NFs is very successful. OCAS uses K8s operators to achieve declarative and intent-based network automation and is described in an upcoming DT whitepaper *Deutsche Telekom's Cloud Automation and Orchestration based on the Kubernetes Operator Pattern (2026).*[10]

10 Deutsche Telekom's Cloud Automation and Orchestration based on the Kubernetes Operator Pattern. Boyan Banev, Dr. Patrick Derckx and Dr. Roland Werding (2026).

Automation Core uses AI extensively for use cases such as root-cause analysis (RCA) and network assurance, for example, for predictive operations and controlling workload scaling. DT is also building a baseline for more advanced AI use cases that will follow in the future.

The Automation Solution & Integration connecting service enables the automation of NF operations using the Automation Core platform

The Automation Solution & Integration layer in Figure 3 refers to the work of DT's internal professional services and systems integration teams to support the automation of individual applications. Automation Core is a common enabler, but automation solutions still need to be tailored so that NFs can be onboarded and their operational needs are met. The Automation Solution & Integration layer takes the generic functionality of Automation Core and tailors it to meet domain-specific requirements. It also supports observability integration, from data extraction up to integration with solutions that use this data. DT's teams tailor automations for applications as they are prepared for onboarding into T-CaaS. As a result, applications are onboarded into HTC with a comprehensive, context-specific set of automation capabilities ready for immediate use.

HTC Automation Solution & Integration layer introduces a 't-shirt size' approach to automation, whereby pre-defined sets of use cases are aligned to varying levels of complexity of NF architectural and operational requirements.[11] DT has built a large repository of standardised automation use cases (over 80) based on both existing and future needs, ranging from basic (small minimum viable products (MVPs)) to advanced (closed-loop, scale-in/scale-out) packages. This repository is shared across DT and continues to grow as more applications are onboarded into HTC.

3.5 DT has worked with vendors such as Amdocs, HPE Juniper Networking, Mavenir and Nokia to implement networks based on its HTC architecture

Approximately 40 vendors provide products and services across the HTC domains of infrastructure, applications and automation. It was critical for DT to engage with vendor partners that support the HTC project's cloud-native, open and disaggregated network principles and roadmap, as well as its vision for radical end-to-end automation. This was particularly important given the significant variation in the compliance and maturity of these capabilities across the market. This section details the involvement of four major partners that supported DT with transforming its networks to the HTC architecture.

Amdocs

DT has sourced its 4G Policy and Charging Rules Function (PCRF), 5G Policy Control Function (PCF) and 5G Binding Support Function (BSF) from Amdocs. Amdocs's capabilities in this space come from its acquisition of Openet in August 2020. Openet had been working on designing new 5G NFs as CNFs for some time (since around 2015), meaning that it was ready to provide DT with a policy management system that conformed to cloud-native principles. For example, the PCRF, PCF and BSF are based on containers/microservices architecture, are able to be orchestrated by K8s, use a service mesh, conform to the principle of immutable infrastructure and are manageable via APIs. One of the key benefits of Amdocs/Openet designing the PCRF, PCF and BSF as cloud-native solutions is that they support efficient elastic scaling. Their cloud-native nature means that individual microservices can be dynamically scaled as needed (as opposed to monolithic applications, which must be uniformly scaled in their entirety).

The Amdocs/Openet CNFs do not rely on an EMS; the NFs' microservices deliver alarms, logs and metrics directly to the T-CaaS observability stack. The cloud-native principles on which the Amdocs/Openet 5G core

---

11 'T-shirt sizing' is an approach to project work that involves assigning so-called 't-shirt sizes' (such as small, medium and large) to particular tasks depending on their expected relative difficulty and the effort needed to complete.

NFs are built drive capabilities that support and further DT's objectives for HTC. These capabilities include 'hitless' (that is, resulting in zero errors) ISSUs that are enabled by the service mesh managing message routing during NF upgrades, updates and roll-backs, and multi-site geographical redundancy. The Amdocs/Openet CNFs also support rapid-cadence release cycles, which enables the necessary agility to conform to DT's 3-2-1-0 vision.[12] The Amdocs team has collaborated with DT and integrated with DT's GitOps, SecOps and continuous delivery (CD) automation platforms, and ensured that Amdocs's CNFs comply with HTC's policies.

HPE Juniper Networking

DT implemented its data centre underlay network fabric using HPE Juniper Networking's technologies. Juniper, which was acquired by HPE in July 2025, has been a long-standing player in DT's cloudification journey, including having major involvement in the NIMS project. DT and HPE Juniper Networking learnt from NIMS (relating to both the technology that underpins cloudified networks and how these networks should be operated) that have been invaluable for the further evolution of cloudified network architectures.

HPE Juniper Networking has advocated for horizontally scalable network clouds since the early days of telecoms cloudification. As such, it designed its portfolio to support operators that are seeking flexibility and the ability to adopt selected components from its stack. For HTC, this means that DT can use HPE Juniper Networking's data centre fabric while retaining the freedom to develop its own network management and automation solution. HPE Juniper Networking's openness to modular deployments, enabled by open interfaces, has been a key factor in its suitability for the HTC project.

DT maintained ownership of the data centre network design, but it worked closely with HPE Juniper Networking to jointly drive the architecture and innovation. HPE Juniper Networking's technical teams, including its professional services team, continue to provide ongoing support to DT. Furthermore, HPE also serves as one of DT's compute infrastructure providers.

HPE Juniper Networking and DT's long-standing relationship enables both parties to collaborate on developing strategic solutions and driving innovation. This close co-operation allows HPE Juniper Networking to align its developments with DT's requirements. These requirements often spark innovations that ultimately benefit the entire telecoms ecosystem.

Mavenir

Mavenir has also been a long-term strategic player in DT's cloud transformation journey; it supplied IMS NFs for the NIMS project and then applied learnings from this project to the next generation of DT's cloud. Mavenir is delivering complete, cloud-native solutions that are designed to meet the evolving needs of HTC. For example, its fixed-line IMS core solution is used for both consumer and business services and supports over tens of millions of endpoints. Mavenir also offers a standalone Breakout Gateway Control Function (BGCF), which serves as the central intelligence for interconnect and wholesale transit traffic routing decisions. Mavenir's IMS platform is used to provide critical emergency infrastructure; it enables nationwide connectivity to emergency call centres in Germany via Mavenir IMS core components. Furthermore, Mavenir delivers messaging solutions for DT's SMS Gateway and for spam protection, and has also successfully achieved the first call for IMS CNFs on T-CaaS.

Mavenir has provided DT with a cloud-native, converged (4G/5G) packet core with proven network slicing applications. This fully containerised packet core has already been deployed on T-CaaS and is serving live

---

12 See section 5.1 for an explanation of DT's 3-2-1-0 vision.

customers; it enables advanced 5G standalone (SA) services including live video production, mobile gaming and RedCap.

Mavenir provides its solutions as CNFs that are ready to use GitOps-based automation and take advantage of the inherent automation capabilities of K8s. To achieve this, it needs to deliver NFs as 'automation artifacts' that fit into the cloud-native automation framework used by DT. The successful implementation of GitOps-based automation has been a collaborative effort between Mavenir and DT, with both parties exchanging learnings about how this should be done. The result of this collaboration is that DT has transitioned to declarative and intent-based operations and can use increased levels of automation to make changes to networks more rapidly and with less human error. The GitOps-based approach has enabled DT to deploy in-service upgrades of Mavenir's combined packet core at a rate of several per month, as opposed to making just a few major releases per year.

Mavenir recognised the importance of DT's vision to apply a unified operational framework across domains. As such, it is providing NFs without an EMS; observability, automation and configuration management are carried out by components supplied by DT. Mavenir has conversed extensively with DT to ensure that all the functions traditionally performed by the EMS are supported in the HTC architecture by other means.

Going forward, Mavenir will continue to evolve DT's HTC in the areas of energy savings, K8s operators and custom resource definitions (CRDs), and AI to drive more efficient network operations.

Nokia

Nokia previously provided Subscriber Data Management (SDM) functionality to DT as appliances/VNFs and is now providing SDM as CNFs for DT's HTC-based networks.[13] Nokia is also providing DT's Session Border Controller (SBC) and Call Session Control Function (CSCF), again as CNFs.

Nokia was quick to recognise the need to make its NFs cloud-native, meaning that it is able to offer its entire portfolio of NFs as CNFs and has considerable experience of deploying CNFs with Tier 1 operators. Nokia's multi-cloud strategy means that its CNFs can operate on a diverse range of K8s cloud environments including on platforms offered by AI and cloud providers and software-only cloud platform providers. This helped DT to deploy Nokia's NFs on its in-house-built T-CaaS platform with reasonable effort. Nokia is working towards aligning with the T-CaaS lifecycle and release cadence.

DT's automation architecture for HTC (that is, Automation Core) is based on the use of generic, rather than vendor-specific, components. Nokia has been supporting DT for years as a strategic partner in shaping the future of cloud-native automation for operators. Around 3 years ago, Nokia set out a vision for cloud-native automation that was based on the use of GitOps, K8s operators and 'everything as code'. Nokia's early work in defining what the automation of cloud-native networks would look like means that it has a plethora of learnings that it has been able to share with DT. Nokia and DT report having close alignment on the principles of automating cloud-native networks, which makes it fairly straightforward for DT to build automation for Nokia's CNFs. Nokia's portfolio is designed with the flexibility to allow operators to build certain automation components in house or use open-source solutions, rather than taking a full suite of automation components from Nokia. Nokia is currently working on feature alignment to realise its complete vision for DT's HTC environment.

13 Nokia's SDM functionality includes the Shared Data Layer (SDL), Unified Data Management (UDM), Authentication Server Function (AUSF), Home Location Register (HLR) and Home Subscriber Server (HSS).

Nokia has served as a strategic partner for DT for the HTC project. It has developed features to support DT's requirements, shared learnings and aligned its roadmaps with DT's to ensure that the pair continually move towards achieving their shared goals relating to the operation of cloud-native networks.

4. DT's new operating model for HTC

4.1 Target organisational structure

The successful adoption of the HTC architecture required significant organisational changes within DT. Teams that originally managed full stacks (including infrastructure, applications and automation) had to be transformed into horizontal organisational structures that reflected the HTC architecture. Achieving widespread acceptance for this required a concerted effort from DT. Application teams needed to be convinced that the infrastructure and automation tooling provided by horizontal teams would work well and would be carrier-grade.

4.2 Cross-team collaboration and expertise sharing

The horizontalisation of teams and the sharing of platforms and tools also increased the potential for learnings to be shared internally. For example, DT formed a team with responsibility for tailoring the automation capabilities provided by Automation Core to individual applications, essentially following a centre of excellence (CoE) model. This team sets up tailored automations for the applications teams and provides education on how to manage them to avoid each individual team needing to set up this tailored automation themselves. DT is adopting a similar approach for other functions whereby individuals with expertise in a certain area are encouraged to support all teams rather than each individual team trying to develop this expertise themselves.

4.3 Upskilling and organisational learning

The HTC project required DT and its employees to upskill and develop their expertise. DT took the experience that it gained from the NIMS project and used this to guide its development work for HTC. DT and its vendors also learnt from each other over the course of the project, and DT ran a comprehensive upskilling programme for its employees.

The time taken to onboard applications onto DT's cloud infrastructure and develop automation pipelines for these applications has decreased as DT has gained more experience, which demonstrates that learning has indeed been ongoing. In addition, DT has now standardised parts of the list of requirements that it uses in RFQs and has stored them centrally after learning from previous RFQ processes.

DT has been a long-time adopter of DevOps and agile ways of working, which has helped it to drive the development of the HTC architecture and to implement networks based on this architecture. However, a further mindset change has been needed to encourage teams to ensure that everything is automated from the start and to avoid the use of command line interface (CLI) tools.

5. Benefits of HTC and next steps

5.1 Key benefits of HTC include the increased speed at which software updates can be deployed and the reduced effort required for DT to build and operate its networks

The automation associated with HTC allows DT to make software updates (including updating software that underpins existing services, deploying new services or applying security patches) much more quickly than was previously possible. In addition, multiple packet core feature sets can be introduced in parallel, enabled by robust Git branching and version control, rather than sequential work in traditional environments. Furthermore, HTC has enabled ISSUs, which means that clusters can be upgraded without needing to migrate traffic away, thereby reducing the preparation and effort needed for network upgrades. HTC is also a fundamental enabler for real-time and near-real-time data use cases, because every layer of the stack (including network fabric, compute infrastructure and CNFs) generates operational telemetry that can be streamed into common Kafka-based data pipelines. Without this full-stack data foundation, higher-level automation capabilities, including autopilot functions and root-cause analysis (RCA), would not be possible.

As a result, HTC has helped DT to get close to achieving its 3-2-1-0 vision: new product releases are ready for market in 3 months, new product releases can be tested and deployed within 2 days, patches can be tested and deployed within 1 day, 0 night shifts are required for software roll-out and 0 service outages occur during the roll-out process. HTC also supports DT with achieving its 0-15-1 vision: 0 outages that affect customers occur, a maximum of 15 minutes is needed to identify an incident that could affect customers and a maximum of 1 hour is needed to resolve the incident. Additionally, automation of Day 2+ processes reduces the effort required for DT to operate its network. Data points that show the benefits that have been achieved with HTC are given in Figure 4.

Figure 4: Examples of quantifiable benefits that DT has achieved by implementing HTC

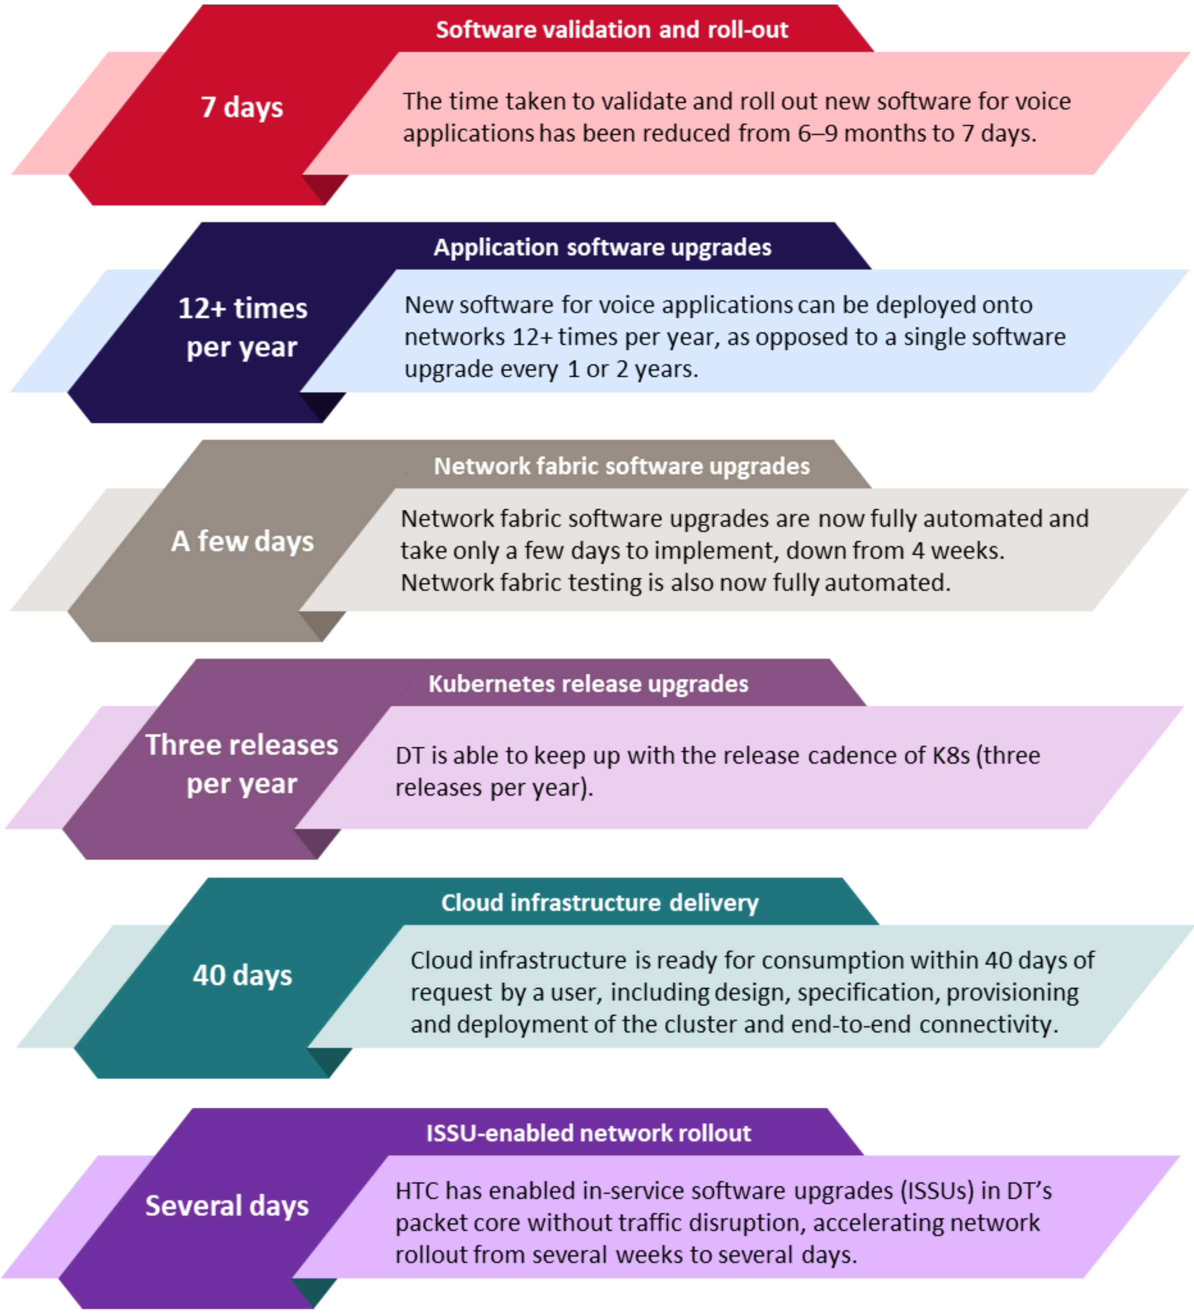

Source: Analysys Mason

DT's transition to using BM4X data centres containing standardised infrastructure for its live networks means that the operator is easily able to replicate core workloads across more than 10 sites. This helps to bring workloads closer to end users, allows for greater georedundancy and allows redundancy to be achieved in a more hardware-efficient manner compared to if fewer data centres were used.

DT has also introduced a new virtual network, based on zero-trust architecture, to simplify connections between its core sites as well as between workloads and hyperscalers' data centres. This has helped to simplify a network that has grown complex over time; it also shifts the primary security layer from the network to the applications and teams behind them. It improves network security by making industry-standard zero-trust mechanisms, transport encryption and strong authentication between network services mandatory.

The harmonisation that results from adopting a horizontal architecture has helped DT to lower costs and efforts by reducing the need to duplicate solutions and development work. However, DT reports that achieving cost reduction in this way was not the primary focus of this horizontalisation initiative. Far more important was the delivery of a better network for its customers and the ability to unlock efficiencies and (cost) savings by using automation to make networks easier to operate.

Analysys Mason views DT as one of the most advanced operators when it comes to implementing a horizontal network cloud. We believe that other telecoms operators are increasingly recognising the importance of adopting such an architecture. However, DT is likely to have developed a lasting advantage over other operators by being quick off the mark and doing much of the development work in house. DT's approach positions it strongly to unlock the agility and efficiency benefits that are becoming recognised as the primarily benefits of adopting cloud-native technologies for networks. DT expects that 6G will primarily be a software upgrade, and its HTC project has already established both the technical and operational foundations for this transition. Additionally, the HTC project sets DT up well to facilitate the greater integration of AI and AI agents with its networks.

5.2 Next steps for HTC

DT will continue to migrate more workloads to HTC, scale the physical infrastructure that underpins HTC and optimise the automation (and observability) capabilities of the HTC architecture over the coming months and years. Power efficiency is one area that DT is targeting for optimisation: DT is working on an AI-based solution that causes the power usage of its core networks to follow data traffic trends, rather than being largely decoupled from it. Additionally, DT will increase harmonisation across its various natcos by eventually extending the implementation of HTC to its operations in all countries.

DT is actively involved in various industry and open-source communities to promote and establish the HTC concept and contribute back to the components used for the realisation of HTC. For example, DT participates in the maintenance of, and/or regularly contributes to, projects such as Kube-vip, Coil CNI and Network Plumbing. DT has also donated innovations that it developed for T-CaaS, such as host-based routing, to project Sylva.[14] DT believes that T-CaaS is still ahead of Sylva but it says that, with further contributions to and collaboration with Sylva, the two can converge at some point in the future because the concepts are very similar.

6. Conclusion and key takeaways for the telecoms industry

The telecoms industry has a long-standing vision to build networks based on horizontal architectures and cloud-native technologies, but few operators have been willing to fully commit to this transition. DT's HTC project has demonstrated that this vision can be made a reality and that operators that do this successfully will be rewarded in the form of being able to achieve highly automated, agile and efficient networks. This journey has not been without challenges for DT: it had to spend a lot of effort building internal capabilities and says that it has often had to act as a 'pipe cleaner' for the industry given that many of its vendors were not ready for this

14 Linux Foundation Europe (July 2025), *Sylva Accelerates Telco Cloud Innovation with Host Based Routing, v1.4 Release and New Industry Backing*.

change. However, DT aims to demonstrate that this journey has been a success for its networks and encourage other operators to follow in its footsteps.

The key takeaways for other operators looking to transition to a network architecture similar to HTC are as follows.

- **Operators should begin the horizontal cloud journey early and take ownership of it.** It takes considerable time to design the appropriate architecture, build internal capabilities and promote the organisational/mindset change that is needed for a project like DT's HTC project. Operators should start on this journey as soon as possible to avoid being left behind. They should not wait until the 6G era before making this transition. Next-generation networks and services will be an evolution of today's cloud-native, horizontal network architectures rather than a revolution, so operators need to lay the foundation now with a futureproof platform and organisational set-up. It is also important that operators take ownership of, and drive forward, this network transformation rather than waiting for their vendors to prescribe the direction for them.

- **Use precursor projects and industry best practices to de-risk transformation.** DT's HTC represents a large-scale, step-change transformation. However, it was underpinned by years of earlier experimentation via smaller, focused initiatives that helped to shape the target architecture, tooling choices and operating model. Transitioning to a horizontal cloud architecture is therefore not limited to Tier 1 operators. Other operators can begin with contained projects that build cloud-native skills, automation discipline and organisational confidence, while maintaining a clear long-term vision for a horizontal, network-wide platform. Drawing on lessons from leading operators such as DT, as well as industry initiatives and open-source communities, can further reduce execution risk without diluting strategic ambition.

- **A change in mindsets and organisational structures is a must.** Operators need to change how they build and operate networks from an organisational point of view in addition to changing the technical architecture. They need to break down operational and organisational silos and structure their teams to mirror the technical architecture that they are trying to achieve. Getting employees onboard for such an organisational change may be challenging but is critical. Operators need to invest time and effort in getting employees to see the value of this change.

- **Devise and enforce a new operating model with clear non-negotiable factors.** Operators should define a new operating model for the standardised, cloud-native, horizontal network, underpinned by clear non-negotiable factors. At a minimum, this should include Day 0 automation for all workloads, GitOps as the default model for configuration and lifecycle management, open and standardised interfaces and tooling, EMS-less operation and the reuse of standard pipeline patterns across onboarding. Operators should then make this operating model enforceable by embedding these non-negotiable factors into RFQ requirements and applying them as onboarding gates and day-to-day guardrails.

- **Set firm expectations with vendors and be explicit about the target architecture.** DT found that some vendors were initially not ready to provide solutions that met their needs; for example, some vendors were not initially willing/able to provide their NFs without an attached EMS. Convincing partners to make changes to their solutions was a beneficial step towards broader acceptance. The compatibility of vendors' solutions with HTC was also key to the successful introduction of an open ecosystem. For projects like HTC, strict requirements should be established for suppliers; potential suppliers should be rejected if needed to ensure consistency in the network.

7. References

- The Sylva project: https://sylvaproject.org/
- Analysys Mason’s *Network cloud infrastructure: worldwide forecast 2024–2030*: https://www.analysysmason.com/research/content/reports/network-cloud-forecast-2024-rma16/
- The CNCF: https://www.cncf.io/about/who-we-are/
- Analysys Mason’s *A move to cloud changes the game for Deutsche Telekom’s next-generation IMS:* https://www.analysysmason.com/research/content/perspectives/dt-cloud-nims-rma16/
- The 12-factor application methodology: https://12factor.net/
- Deutsche Telekom's Cloud Automation and Orchestration based on the Kubernetes Operator Pattern. Boyan Banev, Dr. Patrick Derckx and Dr. Roland Werding (2026).
- Sylva’s press release on host-based routing (July 2025): https://sylvaproject.org/sylva-accelerates-telco-cloud-innovation-with-host-based-routing-v1-4-release-and-new-industry-backing/

8. About the authors

Gorkem Yigit (Research Director) is a Research Director within the *Networks and Cloud* research practice. He leads the *Cloud and AI Infrastructure* and the *NaaS Platforms and* Infrastructure research programmes. He guides clients through the transition to software- and cloud-based network technologies and architectures that support automation, AI capabilities and API-enabled services.

Joseph Attwood (Senior Analyst) is based in our London office. He is part of the *Networks and Cloud* research practice and contributes to the *Cloud and AI Infrastructure*, *AI and Data Platforms* and *NaaS Platforms and Infrastructure* programmes. Key focus areas for his research include the cloud-native transformation of telecoms operators’ networks, data and AI platforms offered by vendors targeting operators and the application of generative AI in the telecoms industry. He studied computer science at the University of Surrey.