White Paper

The Cisco 5G Strategy Series: Packet Core, Transport, and Identity Management

What You Will Learn

Among the current standards recommendations for the fifth-generation (5G) mobile network are exciting technologies and architectures that promise more flexibility and agility, less complexity, and higher performance while lowering an operator's total cost of ownership. This paper in the Cisco® 5G Strategy Series includes an overview of general architectural design principles for 5G; a look at the evolution of flexible mobile services; and descriptions of service function chaining, the 5G core architecture, network transport protocols, and fixed-mobile convergence (FMC) in a 5G network.

5G Core Objectives Overview

Organizations such as 4G Americas (in the white paper "4G Americas' Recommendations on 5G Requirements and Solutions") and the Next Generation Mobile Network (NGMN) Alliance have included a series of requirements for 5G standards. Cisco believes that these following recommendations will be of great significance as 5G standards are defined.

Support for any access in a manner that is simple and broad. Simplicity means that as new access types emerge, they can easily be integrated into the 5G framework. Target access technologies include not only 3rd Generation Partnership Project (3GPP) standard access (5G as well as legacy 3G and 4G) but also non-3GPP technologies such as Wi-Fi, satellite, community antenna television (CATV), and the range of Internet of Things (IoT) technologies that address short range as well as low-power, wide-area (LPWA) needs. In addition, 5G should be relevant to providing personalized services in a fixed and mobile convergence environment. So 5G should provide an evolutionary path for 3GPP technology specification (TS) 23.402, which include architectural enhancements for non-3GPP access.

Support for network slicing. Network slicing enables the management of multiple logical networks as virtually independent business operations on a common physical infrastructure. In practice, this corresponds to the idea that the mobile network could be partitioned into a set of resources that might be virtual. Each one is called a "slice" that can be allocated for different purposes. For example, a slice can be allocated to a mobile virtual network operator (MVNO), an enterprise customer, an IoT domain, or some other convenient set of services (for example, mobility as a service). A network slice extends the access point name (APN) concept used in the mobile network today.

Acknowledgement of the central roles of network functions virtualization (NFV) and software-defined networking (SDN) in networks today. While standards should not be prescriptive on implementation, Cisco believes that the 5G solution should acknowledge these approaches. We believe that NFV and SDN will be the way to implement control and user plane separation, which will have tremendous value. Additionally, NFV and SDN can enable the high degree of scalability and elasticity of network configurations demanded by 5G requirements. The flexibility created by an NFV and SDN-centric network provides a basis for the network automation tools and mechanisms that will be needed in the future.

The 5G core should be able to easily accommodate new protocols, particularly because it is well understood that the IP suite needs to evolve to be more radio friendly. Mobility can be achieved by using transport protocol identify flows by 5-tuple (a source IP address and port number, destination IP address and port number, and the protocol in use). So no IP address preservation is required on the handover. Quick User Datagram Protocol (UDP) Internet Connection (QUIC), named data networking (NDN), and Multipath Transport Control Protocol (MP-TCP) all fall into this category. We believe that the introduction of new protocols could be a primary enabler to address the mobility concerns expressed in 5G requirements.

The 5G core should also be an open platform for services. Normally, a set of capabilities known as "value-added services" are deployed in the (S)Gi-LAN domain. (S)Gi-LAN describes the network domain that straddles the boundary between the mobile network and the Internet packet data network (PDN). In mobile networks today, (S)Gi-LAN service functions are connected in sequential order between the PDN gateway (PDN-GW) and the Internet PDN edge router. (S)Gi-LAN service functions optimize traffic, sometimes by compressing it (through media optimization), charging differentially or shaping it (as with deep packet inspection [DPI]), or by preparing it for the radio access network (RAN) (through TCP optimization). (S)Gi-LAN service functions also provide valuable analytics information. They can support partner services through HTTP header enrichment. Service functions deployed in the (S)Gi-LAN can support firewall and network address translation (NAT) functions as well. In the future, (S)Gi-LAN must address the growing use of encrypted traffic and will include different forms of proxies as well as interworking functions for the IoT. Cisco plans to build a 5G core to be an open platform for (S)Gi-LAN services, capable of onboarding any new services from any vendors.

General Architectural Design Principles

For the 5G core architecture, shown in Figure 1, Cisco sees the following functional segments.

Figure 1. Architectural Structure for the 5G Core

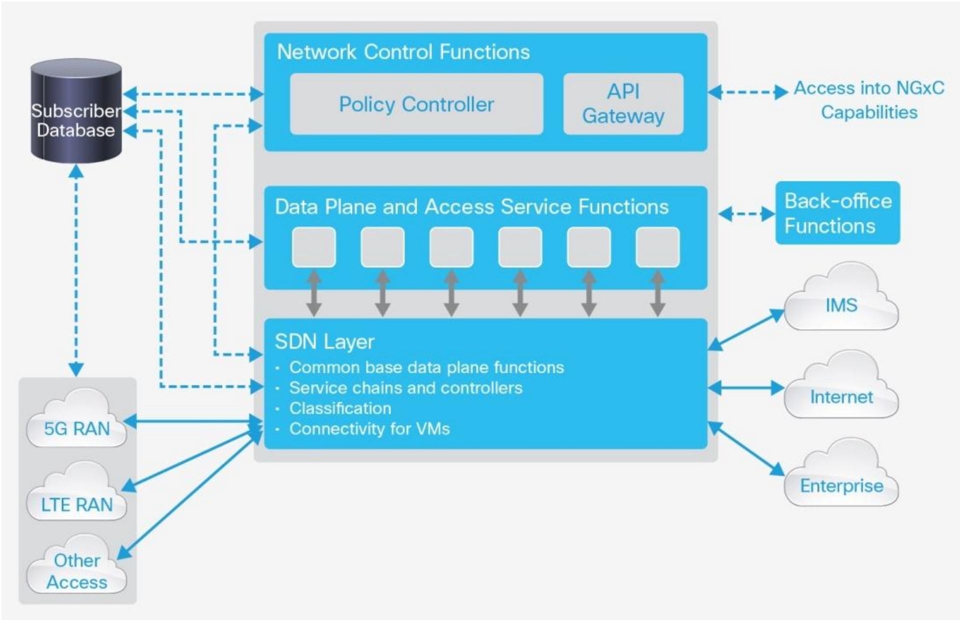

- **The SDN layer** provides bearer capabilities. Its primary purpose is to create service chains in a virtualized environment. Traffic will be classified into service chains dynamically, based on information from the network control functions. Cisco uses industry-standard SDN controllers such as OpenDaylight and open-source SDN network operating system (ONOS) for SDN configuration functions and OpenFlow for rapid packet classification. We also use the incipient IETF standards track network service header (NSH) for implementation of the service chains created by the SDN layer. NSH is highly beneficial based on its ability to flexibly implement forwarding of any protocol and for the ability to pass policy metadata among the service functions composing the service chain.
- **Data plane** and access service functions manage the bearer plane traffic as well as access-specific functions as needed. Access tunnel encapsulation/decapsulation, charging and policy functions, and additional value-added service (VAS) functions are implemented as virtual machines (VMs) in an NFV environment. The SDN layer provides connectivity to allow service functions to plug in to the architecture in a more dynamic fashion. Basic common data plane functions may also be supported in the SDN layer and supported by the SDN classification protocol to allow some data plane services to be passed on to the SDN layer for express operation, offloading the more expensive compute resources. Additionally, some service functions will be specifically designed and allocated to support security and other access-specific functions for an access technology.
- **Network control functions include policy-based control plane and the emerging traffic management and steering function** being proposed in 3GPP. Cisco sees policy as a foundational component for the 5G core. Policy provides the framework for interactions with the bearer plane, determining how SDN is going to treat different flows. Policy determines the charging treatment. It also handles bearer control and quality of service (QoS) assignment for specific flows. In addition, the policy layer exposes capabilities that enable third parties to access functionality in the 5G core. At a high level, Figure 1 represents the Cisco view of how these functions should be implemented.
- **Subscriber database functions** are used for authentication, security management, and holding subscriber-specific data. We propose using the 3GPP user data convergence model in TS 23.335 as the foundation.
- **Back-office functions** including automation and orchestration (based on the European Telecommunications Standards Institute [ETSI] NFV Management and Orchestration [MANO] stack); charging; and operations, administration, and maintenance (OAM).

This 5G core solution is consistent with ETSI NFV architectural principles. This means it operates on commercial compute and uses commercial switching. It is also in alignment with control plane and user plane separation principles. The user plane is based on SDN plus the service functions. One of the advantages of the control plane and user plane separation is that it allows the operator to distribute the data plane from the control plane. For example, in an information-centric networking (ICN) environment or something similar, where no IP anchor is required, a centralized control plane manages bearer control and QoS, charging, and common regulatory functions. The data plane can meanwhile be distributed on an NDN edge router, for example, or to the access edge into the 5G base station element. This split provides the flexibility to support mobile edge computing and other edge functions.

The SDN layer enables the deployment of an access-independent core. This means that access-specific logic can remain in the access network, but when you want to centralize access-specific functions, you can treat them as "plug-ins" in the service chain, as needed. This architecture is well suited to address both network slicing and the more advanced mobility needs described in 5G requirements.

Flexible Mobile Services Steering and 3GPP

3GPP work on the topic of service function chaining has resulted in two technical reports (TR 22-808 and TR 23-718). They cover the topic of flexible mobile services steering, or FMSS. Stage 2 FMSS work is finished in Release 13. Stage 3 work has just started and is expected to affect the policy and charging control (PCC) procedures in TS 29.212 and the PCC architecture. There are several elements to the approach pursued by 3GPP, including:

- A functional architecture for services steering. In the 3GPP architectural view, the policy and charging rules function (PCRF) is used to create, modify, or remove traffic steering policies defined with the granularity of subscriber, application, and service data flow in a packet gateway (PGW) or a traffic detection function (TDF).
- Introduction of a traffic steering support function (TSSF). The TSSF is the control plane part of the TDF without the charging-related aspects of the TDF, which are unnecessary for services steering. The TSSF can be deployed as a modular component for further control-data plane separation. The role of the TSSF is to provide classification control.
- Identify interfaces in the PCRF that can be used to create, modify, or remove traffic steering policies defined granularity of subscriber, application, and service data flow. In this respect 3GPP has (1) enhanced the “Gx” if steering rules are implemented in the PGW, (2) enhanced the “Sd” if they are implemented in the TDF, and (3) introduced a new “St” interface in case the TSSF is identified as a separate control function. The latter has two options. One is based on diameter and is basically a subset of the “Sd,” and the other is based on representational state transfer (REST).

The R13 PCC architecture is shown in Figure 2. The gray shading indicates the user plane elements. The RAN congestion awareness function (RCAF) and the TSSF are unique elements. The TSSF is a function that receives traffic steering control information from the PCRF and makes sure that the related traffic steering policy is enforced in the (S)Gi-LAN.

The RCAF emerged as part of the user plane congestion (UPCON) work and is a functional entity that reports RAN UPCON information to the PCRF to enable the PCRF to take the RAN user plane congestion status into account for policy decisions. The functional description of the RCAF can be found in TS 23.203 and TS 23.401 (and TS 23.060 for 3G).

Figure 2. Emerging Release 13 PCC Architecture

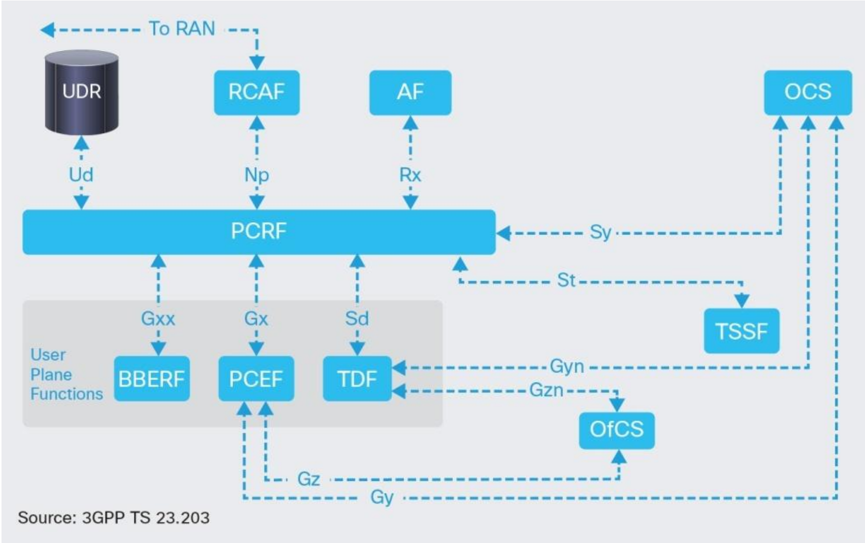

A traffic steering policy rule is preconfigured in PGW, TDF (or TSSF as a PCC rule) and can be referenced for service data flows (SDFs) in uplink, downlink, or both directions. Activation of a steering rule is done through interaction with the PCRF, which can be either "pull" or "push" modes. Should there be a change in policy rules, the reauthorization request procedures can be used to update policy (such as in a diameter reauth answer [RAA] message).

Service Function Chaining

A foundational concept for the 5G core implementation is service function chaining. Service chaining has emerged to describe composite services that are constructed from one or more Layer 4 to Layer 7 services using SDN and NFV. These technologies are helping to move service deployment into modern networks.

A service chain steers traffic through a set of functions in a set order. With NFV, a network service may be decomposed into a set of VNFs, VNF components (VNFCs), network functions (NFs), or a combination of these for flexibility and performance. (VNFC is sometimes used to emphasize being part of a particular service chain, but we prefer to use the term service function.)

The main challenge in the creation of service chains is service insertion. Existing service insertion models suffer from a number of limitations, detailed in RFC 7498. Today, service functions that must be applied to traffic for a given service are physically inserted on the data-forwarding path between communicating peers. Traffic is directed through them using VLANs and policy-based routing techniques. Consequently, services are tightly coupled to the physical network topology, creating constraints on service delivery and potentially inhibiting the network operator from optimally utilizing service resources. New application deployment or the addition of new services into the network is constrained, and this topological coupling limits scale, capacity, and redundancy.

If the necessary service functions are not available to support a new application, then the network operator has no other option but to deploy additional hardware resources and to reconfigure the network to accommodate the new service requirements.

Service functions deployed in this manner are not easily moved, created, or removed, even when virtualized service functions are deployed. This rigidity is the antithesis of highly elastic environments that demand rapid creation, destruction, or movement of the service functions required for application delivery. Additionally, the transition to virtual platforms requires an agile service insertion model that supports elastic and very granular service insertion, modification after services are deployed, and the movement of service functions and application workloads in the existing network. All this while retaining the network and service policies and the ability to easily bind service policy to granular network-centric identifiers such as subscriber information.

A basic form of service chaining may be realized using existing transport encapsulations. This method of chaining relies on the tunneling of selected data between service functions. Using this model, packets are selected based on some classification criteria and then encapsulated in a network transport header (such as GRE or VXLAN, or VXLAN-GPE). The packets are delivered to the initial service function in a service chain. After service function execution, packets are reclassified, and, if further service functions need to be applied, the packets are encapsulated again in a tunnel for transport to the next service function. This process is repeated until all required service functions have been applied.

Although this form of service chaining achieves some level of abstraction from the underlying topology, it does not truly create a service plane. This distinct identifiable plane can be used across all transports to create a service chain and to exchange metadata along the chain. From an architectural standpoint, a number of significant limitations remain for complex deployments. These include:

- Configuration complexity
- Limited service chain construction capabilities; rigid service ordering that prevents the formation of true service graphs
- The need to support many overlays and mapping between overlays
- No way to pass additional metadata
- Limited visibility and audit capabilities
- Limited application of service policies
- Per service classification and reclassification
- Limited integration of multivendor service functions

The solution to the service chaining difficulties is the NSH currently being standardized in the IETF with broad industry support along with an identified set of mobile network use cases. NSH is a data plane protocol that represents a service path in the network. It provides a common service plane fully orchestrated top to bottom. The path information is akin to a subway map: It tells the packets where to go without requiring per flow configuration. Metadata is information about the packets that can be consumed or injected by service functions, so it can enable policy. In Figure 3, NSH is added to packets using a classifier. NSH is carried along the chain to services. Intermediate nodes do not need to be NSH aware. A proxy can take over the NSH function so that non-NSH enabled services are supported. The addition of metadata through the NSH header in the IETF service function chaining (SFC) protocol will enable the evolution of service creation from a dumb set of nonrelated building blocks to composable services that can pass and react to imputed data and midservice computational results.

So beyond simple linear graphs, any kind of directed graph can be supported with a dynamic component whereby a graph node can, using policy, make an autonomous determination to resteer a flow. Additionally, the presence of programmable metadata will enable new opportunities for compliance and assurance. Note that the functions in gray are defined in 3GPP Release 13.

Figure 3. Service Function Chaining

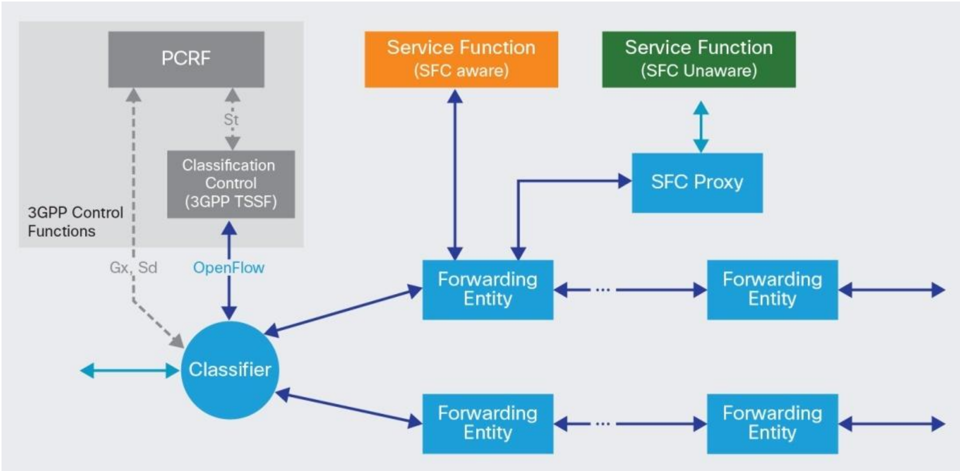

[Figure 4](#) shows a simple linear service chain being steered using NSH. In NSH, packets are steered solely by the value of a special ID, the service path identifier (SPI), and so metadata can be conveyed to participating nodes.

Figure 4. Service Chaining with NSH

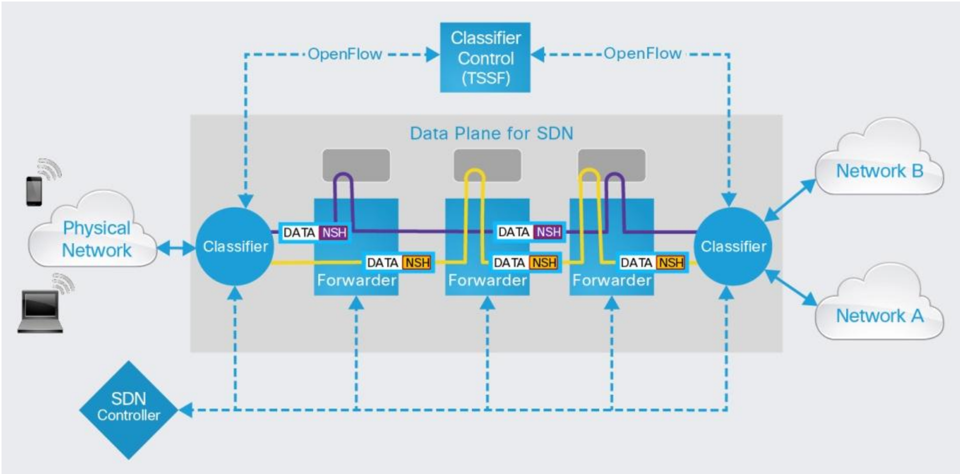

5G Core Architecture

The 5G Core Reach

The design objective for the 5G core is to reach any endpoint over any kind of access network and with any kind of service. Figure 5 shows 5G access, evolved LTE, legacy RATs, fixed access (including CATV, satellite), and IoT applications over various access technologies.

5G will support evolved access interfaces based on the S1 and the S2 family (S2a, S2b, and S2c). We also add new access independence shown in purple for common services such as entertainment video, bulk Internet access, and any managed content that will be common across access technologies in a converged network. The red and yellow interfaces are meant to indicate legacy Internet access and IP multimedia subsystem (IMS).

Figure 5. 5G Scope and Reach

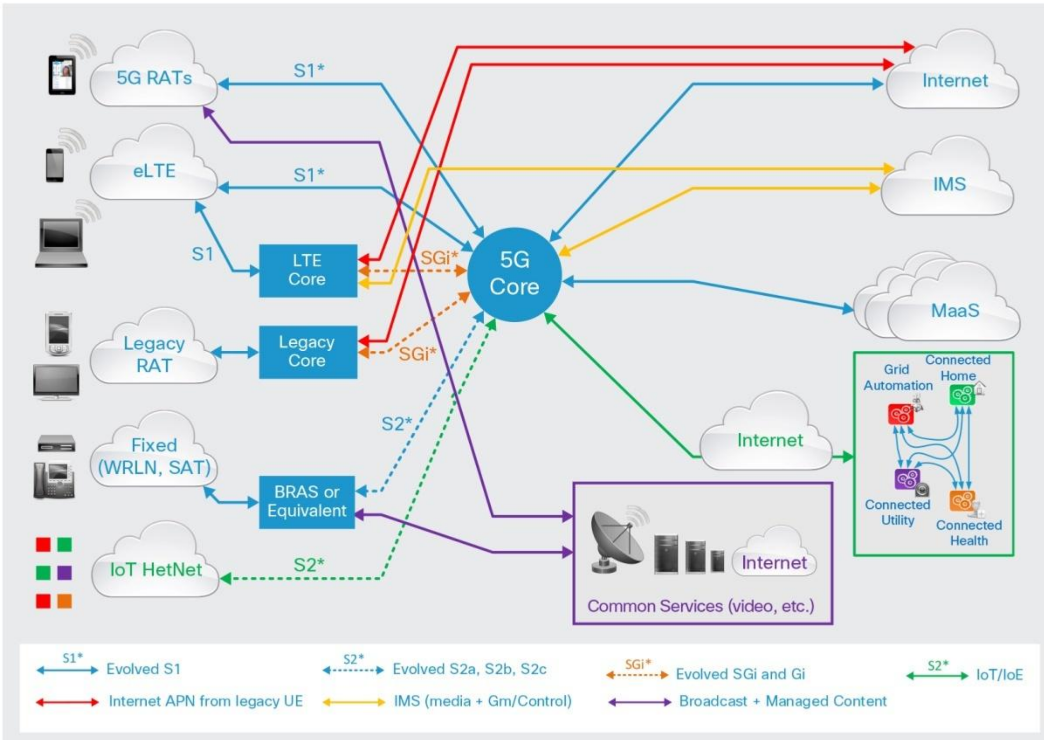

S1* is a tunnel-based interface plus control plane interface. This interface will most likely be an evolution of the S1 control and user plane as applied to LTE. SGi* is a simple IP interface. The 5G core will then only apply VAS to the traffic and its role collapses to that of an (S)Gi-LAN. It will classify traffic to a service chain comprised of virtualized VAS deployed as service functions based on resolving the IP address to a unique international mobile subscriber identity (IMSI) or mobile subscriber international subscriber identity number (MSISDN), which can then be used by the policy layer to determine the classification into a service chain. This is a standard Gx+ call flow (or Sd) as in TS 23.203. When there is IPv4 address overlap, the policy layer can use other information to resolve an IP address to a unique identifier for a user.

The concept of the SGi* is worth emphasizing. The way in which the IP address is resolved to a unique device identity (a satellite, cable, FTTX set top, or LoRA EUI) will be unique to the access technology and will need to be better defined and documented in call flows. For a Wi-Fi device, the registration procedures allow the device identity to be exposed. Therefore, following TS 23.402, it is possible to resolve an IP address to a unique identifier associated to a Wi-Fi user.

The 5G reach capability includes fixed-mobile convergence (FMC). The idea of FMC is sometimes confused with convergence of the infrastructure supporting fixed and mobile networks. The idea of infrastructure convergence has its merits (for example, when an IMS should equally well support fixed and mobile but does not). One example where infrastructure convergence does not make much sense is the idea that the broadband remote access server (BRAS) and mobile gateway subscriber termination points should be converged into a common platform to handle both wireless and wireline endpoints. The problem here is that in a fixed network, a subscription is associated with a location such as a home or residence.

The individual users who are in the scope of the fixed wireline subscription are not visible to the network as individual users. By contrast, in the mobile network a user typically maps to a single person independently of where that person is located.

That person is identified to the network through possession of a unique mobile identity presented through a mobile device or UE equipped with unique identifiers such as those embedded in a SIM card.

So, when it comes to fixed and mobile networks and the implications of convergence:

The focus of FMC should be delivery of personalized converged services. Personalization means that the service capability is personalized to a subscriber. Examples of such services include single-number reach (as implemented with IMS), personalized video, and personalized Internet access (as implemented with parental control filters, for example).

When technically feasible, infrastructure delivering services should be converged. When not feasible or not economically viable, separate infrastructures should be used. In other words, Cisco acknowledges that using a common delivery platform is not an objective of FMC.

In the mobile network, a specific services delivery capability is oriented toward the individual subscriber, a person who consumes networking services. There are important attributes of mobile network services delivered through the packet core. They:

- **Require an identity management framework focused on the individual subscriber**. Such a framework relies on 3GPP databases such as the usage data records (UDRs) that is defined as part of subscriber data convergence. The fixed network has no such databases.
- **Use subscriber policy as a fundamental attribute of 3GPP mobile architectures.** The fixed network has no personalized policy management. The broadband forum has recommended how 3GPP’s policy and charging architecture solution can be applicable in fixed networks.
- **Correct for the lack of inherent mobility support in the TCP/IP protocol suite.** Without support for anchoring sessions in the mobile network through the PDN-GW and because IP addresses are location dependent, handovers will break TCP sessions and hopelessly ruin UDP. The mobile network infrastructure corrects for the absence of mobility support in TCP/IP by introducing a mobility anchor and a comprehensive mobility management signaling framework (the mobile packet core) to make sure of IP address preservation when a mobile device moves about.

For the packet core specifically, the bulk of the bandwidth does not benefit at all from traversal of a packet core subscriber termination point, in this case the PGW. The specific bandwidth services referred to here include linear broadcast video, video on demand (VoD), and general purpose Internet access. These capabilities are delivered as massive raw bandwidth services, which are delivered at the location of the fixed subscription (for example, the home). Some of these services may be personalized on the back end. An example is a VoD service that provides a user interface with entries populated from a recommendation engine tied to the user identity and previous views. However, the delivery itself does not require any special processing in the packet core.

In practice, combining personalized services and massive raw bandwidth services in one converged fixed and mobile platform might result in an architecture that is not optimal for both capabilities. Therefore Cisco does not recommend pursuing the concept. The architectural solution embraces the concept of offload. Bandwidth that does not benefit from traversal through the mobile core is not sent through it.

The concept of offload from the packet core is going to be short-lived. New protocols are being introduced such as multipath TCP, QUIC, ICN, and NDN for which it is not necessary to bind a protocol session to an IP address. This is because, in the case of QUIC and multipath TCP, a connection identifier is used that persists across handovers between locations and in the case of ICN and NDN because data chunks themselves have a name.

With such protocols, a mobility anchor is not required, and much of the reason for having a packet core vanishes. The prospect of new protocols that have built-in mobility support is a huge potential boon for FMC.

The present reality, however, is the TCP/IP stack and a huge base of applications that rely on it and are not inherently mobile at the protocol level.

New Transport Protocols for 5G

The industry focus on developing new 5G networking technology offers new opportunities to address many of the fundamental and persistent issues that exist in today's data networks. Attempts to address these issues in the past have been characterized by the proposed introduction of new incremental changes or overlay solutions. These are often expensive to fully deploy or cumbersome to manage. Some of these fundamental issues are articulated in the following questions:

- Given emerging and new access network technologies for IoT along with unlicensed spectrum and Wi-Fi evolution, how do I implement mobility across this diverse access environment?
- IP address preservation has been a primary tenet of existing mobility solutions in order to support session and service continuity during mobility handover. Yet this results in a complex network infrastructure. Are there more efficient ways to implement session and service continuity to enable a more efficient mobility network architecture?
- How do I manage the huge expected traffic growth in the network, all driven by increased content (particularly video) consumption? In the past, circuit-based networks have been built to optimize for circuit switched traffic. More recently the focus has been on packet-based networks to efficiently transport packets. As network traffic becomes increasingly dominated by video content, do I need a network optimized for the type of information and content that composes the bulk of transport needs? Can I cache content in the network in a distributed fashion to create an intrinsic, scalable information-centric network?
- How do I deal with security and associated encryption issues while providing reasonable network management (HTTP 2.0)? Can the industry more efficiently address increasing security concerns by securing the content rather than the container (for example, router, host) or the communication channel?
- How do I more effectively deploy broadcast and multicast services ubiquitously given today's cumbersome overlay methods?
- What can be done to address the existing smartphone issue of access network selection: for example, how to choose between Wi-Fi and cellular to get the best user experience? The more general question might be: Are there more effective ways to utilize simultaneous multiple access network technologies for improved throughput and better QoS?

As technology options and solutions for 5G are proposed and agreed to, it is important that these questions are carefully considered. The goal is to implement solutions that are cost effective and result in an overall simpler network infrastructure.

In considering 5G network solutions, three main components are perhaps the most prominent: the access network (the 5G RAN), the core network, and the underlying protocols used over the network. Much work has been done so far on the access and core networks, but unfortunately the underlying protocols have been largely ignored, even though they might be the keys to greater efficiency and cost effectiveness in 5G.

Mentioned among the preceding questions are two fundamental attributes that can provide solutions in 5G network standards. First, note that capabilities such as mobility, service continuity, and access network selection are all constrained and complicated today by the fact that existing Internet protocols (such as TCP) effectively use the IP address not only for routing but also to identify the session related to the communication exchange.

By separating the session identifier or content name from the routing, things such as mobility, session continuity, and support for multipath transport can be integrated into the underlying protocol, offloading the network from these complex tasks. Additionally, if the new underlying protocol operates based on the name of the data, issues related to caching in network storage and security are further addressed.

Embedding the content name into network primitives allows for a more agile connectionless transport model driven by the end user that is not bound to a network addressable interface. This simplifies the management of multihomed mobile and fixed communications and provides fine-grained control in data delivery. Because of content awareness carried by content names, the network is able to route requests toward nearest content replicas, exploiting temporary in-network caches and adaptive request routing for a more efficient and cost-effective data delivery. Finally, the symmetric routing attribute endows the networking layer with natural mobility support. The inclusion of support for mobility, security, and storage in the networking layer leads to a simpler, more cost-efficient architecture that intrinsically supports modern communication usage patterns.

ICN and its implementations in NDN and content-centric networking (CCN) identify a new networking paradigm that utilizes content awareness at the network layer to simplify the network architecture. Network operations are driven by content names rather than location identifiers such as IP addresses to gracefully enable user-to-content communication.

The basic idea of ICN is to enrich network layer functions with content awareness so that routing, forwarding, caching, and data transfer operations are performed on topology-independent content names rather than on IP addresses. Data is divided into a sequence of chunks uniquely identified by a name and permanently stored in one or more servers. Naming data chunks allows an ICN network to directly interpret and treat content according to its semantics with no need for DPI or delegation to the application layer.

The naming convention does not need to be specified and can be application-specific. Only a hierarchical structure, similar to that already adopted by HTTP, is required for entries aggregation in name-based routing tables. The hierarchical naming scheme uses hierarchically ordered labels such as URIs. More precisely, a name is composed by a variable number of components, which is not necessary human-readable as URIs, organized in a hierarchical structure.

One of the most important differences between name-centric and traditional host-centric networking is that data is retrieved by name rather than location. Hence, in ICN architectures data authenticity verification (that is, the verification of the publisher of a named data object) is an important challenge. Data authenticity is achieved by applying a digital signature such as a hash of a name plus a data object through a publisher's key to a named data object with a hierarchical naming scheme.

Because forwarding in the data plane is based on NSH and not specific packet fields, ICN and associated network caching mechanisms can easily be integrated into the 5G architecture shown in Figure 1. For example, ICN routers can exist in the Gi-LAN services to support a migration strategy to ICN-based protocols using the subscriber-aware SDN layer.

Later, an ICN element can be integrated with elements close to the access edge such as an aggregation router and even the base station or cloud RAN itself. The control/user plane split functions of the architecture are used to allow an integrated and more centralized control plane to manage QoS, charging, and other functions even when no data plane anchor is necessary. Additionally, a Layer 2 PDN type could be used to route ICN traffic directly to an ICN PDN, avoiding IP routing that to the edge of the network.

Fixed-Mobile Convergence

The essential solution component for FMC is provided by access through tunnels into the mobile core from a non-3GPP network. Such access is enabled through mechanisms detailed in TS 23.402.

The (S)Gi-LAN can support FMC services. In Figure 6, devices that attach to the fixed wireline network can use services that are also delivered as part of the (S)Gi-LAN by following the path indicated in the tunnel. The tunnel interface is an S2a or an S2b from TS 23.402 that extends from the mobile core to a multiaccess device supporting the FMC capability. Such a device could use split tunneling to access local resources such as a printer but would utilize services delivered from the (S)Gi-LAN.

Figure 6. Fixed-Mobile Convergence Use Case

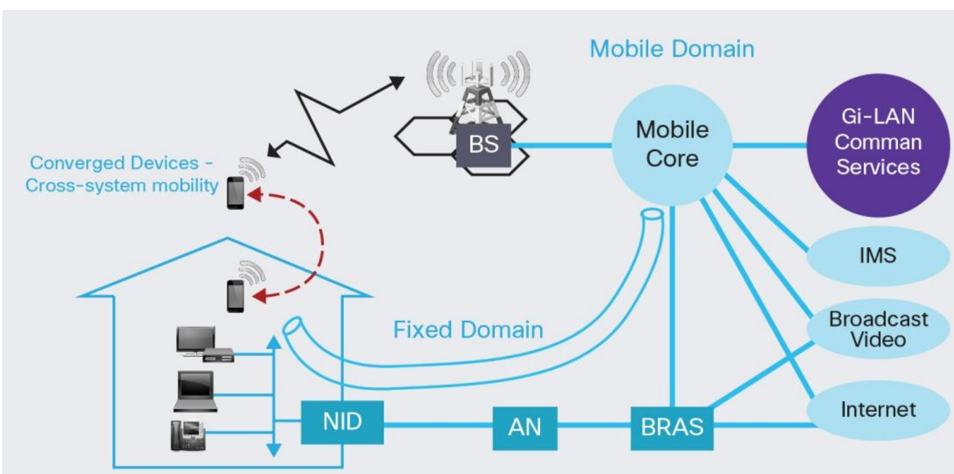

Subscriber Management

Cisco recommends evolving to the user data convergence (UDC) framework in TS 23.335 as in Figure 7 as the foundation for subscriber management in the 5G core. Note that the names of current interfaces in the diagram are presented to depict existing functionality. Existing operator databases will play a role in this evolution for subscription management.

Figure 7. Database Solution Based on User Data Convergence

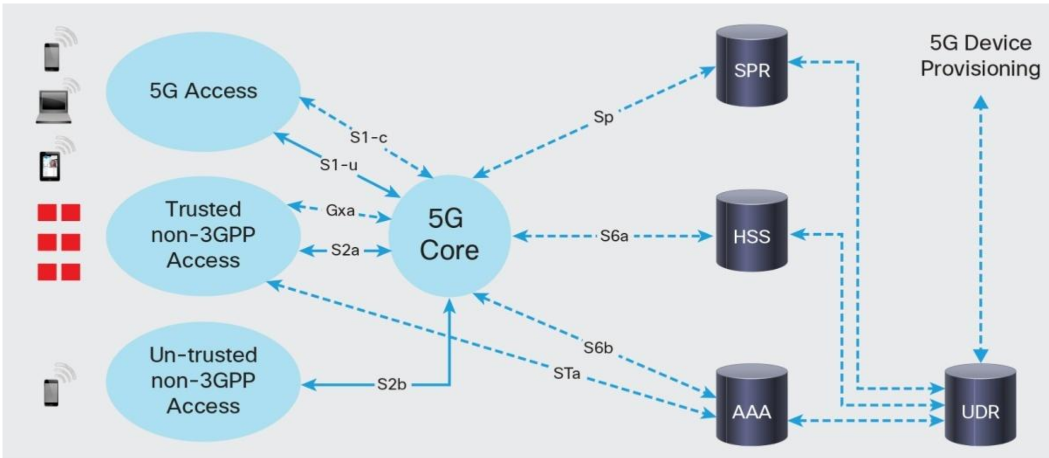

Summary

5G standards recommendations include many exciting new approaches to some of the most daunting problems affecting mobile operators. They call for support for an array of access methods with the emphasis on simplicity. They promote network slicing and a prominent role for NFV and SDN. Flexibility in the 5G core is a top priority, including new protocols to allow networks to be more radio friendly and to be an open platform for services of all kinds.

Alongside the standards bodies and others that are helping to define and design 5G, Cisco is working on a variety of approaches and technologies similar to those discussed in this paper to ease the evolution from 4G LTE to 5G and to make sure of the success of 5G for years to come.

For More Information

Go to http://www.cisco.com/go/mobileinternet

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-736765-00 06/16