# Agile Standards for 5G Slicing

By: Mark Cummings, Ph.D., Christos Kolias, Vinay Devadatta

With 5G Slicing being a critical part of 5G financial success (see also: Orchestration Imperative for 5G Slicing) and the fact that to work, it has to cut across multiple operators and vendors, there will need to be standards. To meet 5G Slicing's full promise, however, these standards must be agile in order to respond to today's rapidly changing world. This means that they must be as minimal as possible. That is the true essence of agile, to be both easy to implement and leave as much as possible up to real time negotiation. Thus, such standards should have the flexibility to enable real time automated functionality that responds to the requirements we can foresee today, as well as those that will inevitably emerge that are not yet anticipated. In this context of agility, initial focus should be on the framework for negotiations and extending the 3GPP/TMF Umbrella model.
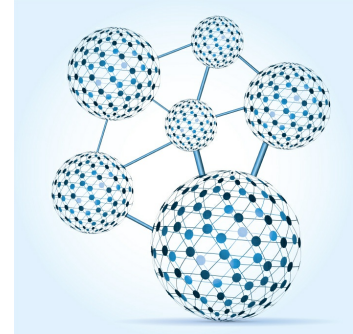
Negotiation, Federation, Orchestration

In "Orchestration Imperative for 5G Slicing," the key capability to work across administrative/ownership boundaries was described. This cross-boundaries capability requires orchestrators on each side of the boundary that can negotiate with each other (see Figure 1.).
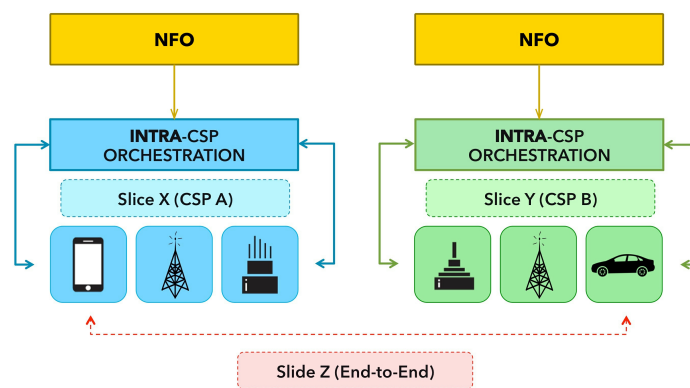

Figure 1: Generic Multi Administrative Unit Multi CSP Slice Example
Click to Enlarge

Based on that description, Caroline Chappell, a leading industry analyst, has suggested that we name these inter-CSP Orchestrators NFOs, or Negotiators, Federators, Orchestrators. This name captures the sense of the functionality of these orchestrators and will be used here.

Two example detailed use cases were described in the aforementioned previous article. They involved two CSPs sharing network resources (for example base stations), and a customer requesting a service that must span two administrative units. In the case of sharing of resources the focus will be on operations—configuration, fault recovery, etc.—that integrate into the two different network structures. In most cases, the parties involved and financial matters will have been predetermined.
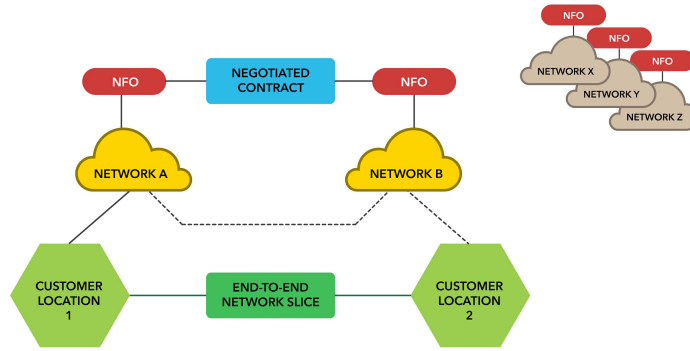
Figure 2. Example Negotiation Process Environment
Click to Enlarge

In the case where a customer requests a service from CSP A that must, by its nature, involve at least one other CSP, the focus can be broader. As shown in Figure 2, there may be several choices of possible administrative units (Networks X, Y, Z, and B). Therefore, the negotiation process must provide for a way to get sufficient information to choose as well as negotiate. Once negotiation starts, it must include the operational factors, but also business factors such as price, settlement, Service Level Agreements (SLA), what happens if SLAs are not met, etc. SLA issues can include who is responsible for creating trouble tickets, troubleshooting, communicating with the customer, financial matters, and so forth.

Negotiation Framework

With this background, the framework for the negotiation process becomes clear. It should follow the process shown in Figure 3, starting at the bottom of the figure and proceeding to the top.

The first step is Discovery. For example, Network A must discover that Networks X, Y, Z, and B are available. Then, Network A must connect to them. Once connected, they need to exchange basic description information; that is, enough information to determine if it makes sense to start negotiation. Given today's Web and Web-related tools, Discovery, Connection, and Description can be done through common Web resources and do not need any special detailed standards. Based on the Description information, Network A has to decide which other networks to start negotiation with. There may be security-related aspects of these and the following steps in the process. These security aspects will be discussed in a subsequent article.
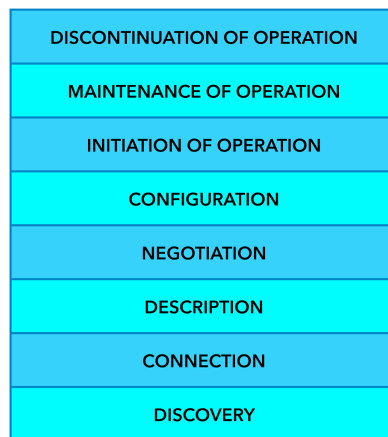
Figure 3. NFO Process
Click to Enlarge

In this automated world, the NFOs have to make decisions. These decisions are based on Objectives, Algorithms, and Constraints. There is experience with this kind of decision-making in SON (Self Organizing Networks). SON promulgated by 3GPP is used for certain tactical activities such as load balancing in neighborhoods of cellular basestations. In SON, the objectives,

algorithms, and constraints are hardwired and the range of choice or decision is very limited. With NFOs, the range of objectives, algorithms, and constraints will be much greater. With networks having multiple NFOs, there will be a need for another component that configures each NFO and gives it its objectives, algorithms, and constraints. Discussion of that component will be the subject of a subsequent article.

Based on its objectives, algorithms, and constraints, and the Description information acquired through Connection, Network A's NFO decides which other networks' NFOs to begin negotiation with. It is possible (and may in fact be desirable) that there will be more than one negotiation going on in parallel (where for example, Network A is negotiating with Networks X, Y, and Z for the same slice, then picking the best one). Typically, the negotiation will start with a bid by NFO A. The corresponding NFO responds with either a bind accepting the bid, or with a bid of its own. This bid/bind process continues until a complete Contract is agreed or it becomes clear that it is not possible or desirable to reach a contract. Creation of a contract completes the Negotiation phase of NFO.

Once created, the contract specifies how the two interacting networks (in Figure 2, Networks A and B) need to configure themselves. Because service initiation may require fine-grained time synchronization, the configuration and initiation stages are separated. As per contract, once the configuration stage has been completed, the service(s) across the network slice are Initiated. Initiation has now created a Federated network slice transiting the two administrative units. Now Maintenance begins.

Maintenance is the ongoing orchestration process that attempts to maintain the slice operating as per the contract. The contract may provide for specific actions if behavior deviates from what has been agreed upon, for example, if an outage occurs such that Network B can no longer provide service. If so, those steps are taken. If the deviation can't be cured by those agreed-upon steps, or there are no such agreed steps in the contract, negotiation attempting to achieve a new contract is started. If that negotiation fails, then Network A goes back to Discovery and proceeds from there. To make this possible, the contract should also specify what monitoring information is to be shared so that deviations from contracted behavior can be measured.

Finally, the contract provides for Discontinuation, which covers how and when the slice should be discontinued. This is important to release committed resources to ensure that old slice components (so-called "zombies") do not proliferate and consume network resources. Unfortunately, this is a problem already in existence today. There are many things running in our networks and infrastructures that were started years ago. Nobody knows what they do, but people are afraid to turn them off. Most of these things just consume resources while doing nothing productive, or worse. Having a requirement for specific terms for discontinuation can go a long way toward avoiding the proliferation of more such zombies.

Umbrella Data Model

Now that we have a negotiation framework, how do the NFOs communicate? The default can be text strings, but it will be much easier if it is possible to communicate by simply referencing objects in a data model. But where does the data model come from? It turns out that there are many forces working on data models in the CSP context. A number of different standards organizations have developed different and sometimes overlapping data models. In an effort to differentiate their products, vendors have developed different data models. And different CSPs have developed their own, too. It is tempting to retreat to an older paradigm and set out to create a single data model for all, but that is both impractical and will create a serious roadblock to 5G slicing. What we really want to do is enable 5G slicing now, not put roadblocks in the way.

The more that is left up to negotiation, the more agile the whole system will be. So, the choice of data model can itself be the subject of negotiation. A common well-thought-out way to start might be to use the Umbrella Information Model (UIM) promulgated by 3GPP SA5 and TMF. The UIM is a Meta Model derived from the previously existing standards of the underlying data models of the two organizations, and those of participating vendors and CSPs.

The Umbrella Model was developed in response to a set of requirements generated by NGMN (Next Generation Mobile Networks—an industry association comprising the largest cellular providers). The development of requirements for the extensions to the Umbrella Model could be similarly supported by an industry group started earlier this year called ZSM (Zero-Touch Service Management).

The UIM was intended to support provisioning and configuration of wireless and wired network elements. To this operational base, we would need to add SLA and Settlement data objects. Those extensions can be done initially with text strings. Thus, while some models are developed in standards organizations, multiple compliant models could evolve in practice. However the umbrella model evolves, there are likely to be a number of versions of the model. So, early in a negotiation, the NFOs could start by bidding model versions until they reach agreement. This way, as the models evolve, the overall system can take advantage of them without significant disruption.

Now, inside the CSPs, there are still likely to be different data models in use. There will have to be a translation capability (sometimes called a Bridge) inside the NFOs to translate to the underlying different model(s). These models reflect the views of each CSP on its day-to-day operations. As an aside, there may be a range of different models in use inside a given CSP and a range of different layers of orchestrators with different bridges. In any case, it will not be necessary to translate all of the vendor-specific, domain-specific, and CSP-specific data models. Only the parameters related to the subjects of negotiation need to be available to the NFOs.

Some people have suggested that we fix certain portions. For example, a common suggestion is that we fix the use of blockchain and a particular crypto-currency for settlement. This is dangerous. A few years ago, we didn't know that blockchain and crypto-currencies were coming. We can't predict what will appear in a few years. But we can say with a high level of confidence that something we don't know about now and which will have the potential to affect settlement will come. So, if, for example, we restrict settlement to a blockchain process using Etherium today, it is likely to be overtaken tomorrow. By leaving it up to the negotiation process, if a CSP wants to use blockchain and Etherium, it can bid it. If the other NFOs have set up this option and it supports their objectives in this slice, they can bind to that.

Over time, a set of widely employed options will emerge. This set of widely used options will also slowly evolve as technology develops. At the same time, CSPs that have close partner relationships may develop specific sets of data models and processes that are tailored specifically for their unique situations.

Thus, by focusing on making standards to support 5G slicing agile and easy to implement, the full economic promise of 5G can be realized. And we have seen how focusing the standards efforts on the negotiation framework and umbrella models is the best way to maintain that agility, enabling near-term deployment of 5G slicing.