

I E T F[®]

In-Network Computing for App-Centric Micro-Services

<https://tools.ietf.org/html/draft-sarathchandra-coin-appcentres-00>

C. Sarathchandra, D. Trossen, M. Boniface

IETF-105, COIN PRG, July 2019

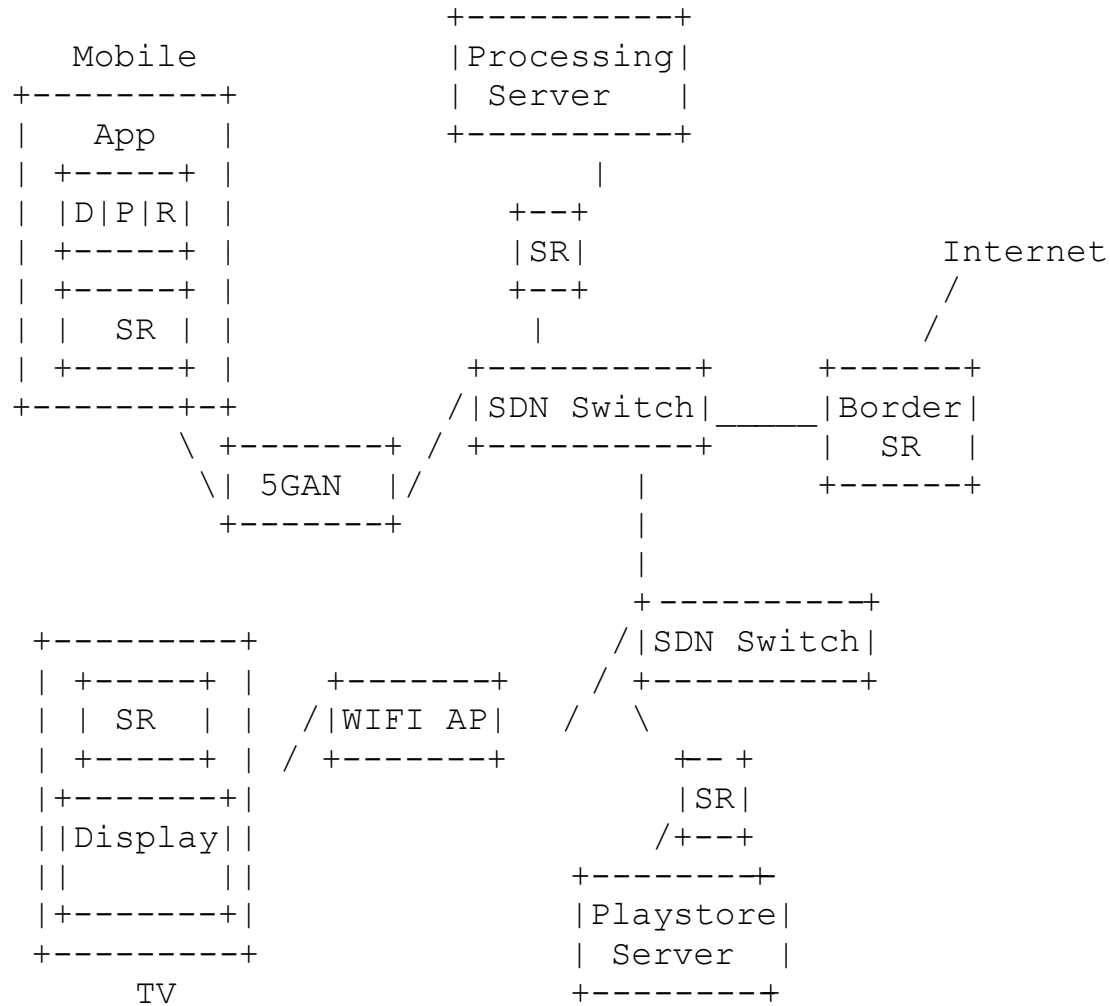
From Apps to AppCentres

- Starting point is the **app economy** of today's smartphone world
 - Drives development of mobile experience as we largely know it
 - Static design of app as SW modules, running as process(es) over existing smartphone OS
 - Extended often with server/cloud interactions
- Move to mental model of micro-service (MS)
 - Decompose app into set of MSs
 - Execute MS on one or more distributed compute resources rather than single smartphone
 - Interpret any such compute resource as a (pico/micro) data centre
 - > turns a data centre into an **app centre**

Use Case: Mobile Function Offloading

- Design app as modularized SW entity
 - Wrap **MS helper classes** around modules, e.g., **d**isplay (an image), **p**rocess (an image), **r**eceive (a video to produce an image)
 - Interpret a given experience (e.g., watching a movie with additional filter effect) as a **chain of MS** (or service functions)
 - When running with all modules/MS active on device, it looks like today's app
 - When one MS is being removed locally, the MS nature of each module will lead to communicate over the network with suitable other MS instances
- > Controlling the overall distributed experience is now about **WHERE** and **WHEN** suitable MS instances are being used

Use Case: Realization



Enabling Technologies

- Application Packaging
 - Usually done at design time through appropriate frameworks
 - Wrapper classes as in use case can be used to turn modules into MSs
 - Profiling can be used to decompose originally monolithically packaged app into set of MSs
- Service Deployment
 - Orchestration framework for in-network deployment (many to choose from)
 - Integration with current app installation model needed though

Enabling Technologies

- Service Routing
 - Usually done as combination of DNS and IP routing
 - Ongoing work on named service function chaining and service routing over L2 environments might be suitable here
- Service Pinning
 - Service routing in environments with several MS instances requires selection of the ‘right’ instance at the ‘right’ execution point -> pinning of the service chain path
 - Pinned relations can change frequently, usually defined through (changing) constraints matched against expressed requirements
- State synchronization
 - State exchange and synchronization is crucial due to distributed nature (of instances)

Conclusion

- Draft positions an evolution of the current ‘app model’ for mobile terminal towards a highly distributed execution of micro-services
- Draft interprets available compute resources as (pico/micro) data centres, where these compute resources can be other terminals, in-network resources, cloud DCs
- Draft presents challenges to be addressed to drive forward this evolution

Authors believe that the establishment of the (proposed) COIN RG will provide the suitable platform for bringing together, evaluating and recommending solutions to those challenges

Next Steps

IF the COIN RG is approved, we plan on

- updating the draft with more information on ongoing work for solutions to the outlined challenges
- providing an overview of realizations of some of the solutions
- demonstrating example realizations in an upcoming COIN RG meeting