



# **PURSUING HYPERSCALE ECONOMICS: THE ROLE OF DISAGGREGATION AND CLOUD NATIVE**

Webinar: Questions and answers

# Pursuing hyperscale economics: The role of disaggregation and cloud native

*This document outlines the questions and answers received from the STL Partners webinar, **Pursuing hyperscale economics: The role of disaggregation and cloud native**, which was hosted on Thursday 9<sup>th</sup> September 2021.*

*In this document, we seek to address the questions raised in the webinar that we were unable to address in the time available.*

**You can also watch the recording of the session, and also access the slides, using the link [here](#).** We have included the following timestamps for the webinar recording:

- **02:20** for the introduction to our presenters and panellists
- **03:27** for STL's presentation on "Pursuing hyperscale economics: The role of disaggregation and cloud native"
- **24:11** for DriveNets' presentation on "The right way to pursue hyperscale economics: It's the data plane, stupid"
- **38:54** for the live panel session with:
  - **Bryce Mitchell**, Vice President, Core Network Engineering, Rogers Communications
  - **Joe Baeumel**, Vice President, Business Development & Partner Strategy, KGPCo
  - **Inbar Lasser-Raab**, Chief Marketing Officer, DriveNets

---

*If you have any questions not addressed in the webinar or this Q&A document, or want to hear more about our research findings or from our speakers, please contact:*

- Yesmean Luk, STL Partners – [yesmean.luk@stlpartners.com](mailto:yesmean.luk@stlpartners.com)
- Dudy Cohen, DriveNets – [dcohen@drivenets.com](mailto:dcohen@drivenets.com)

*For more details on the research findings, you can access the 'Pursuing **hyperscale economics: What, why and how telcos can do it**' research report [here](#).*

# Webinar questions and answers

The below questions were received from the webinar audience during the live session, this does not include the questions asked to our panellists.

For the Q&A session with our webinar panellists, please refer to the recording of the session [here](#).

---

- 1. From a telco perspective, at what point, or what are the key triggers to make the leap from a traditional locked-in solution to move to a disaggregated cloud-based solution? For instance, is it part of a technology cycle, or a CapEx-OpEx decision, or another reason?**

**STL Partners:** Technology and (re)investment cycles can often be the trigger for operators to evaluate disaggregated cloud-native approaches. 5G standalone is a good example of this. However, the experience of some operators is that they face time pressures (market deadlines) that add further risks to making the 'leap'. These are big changes and those operators that have made a strategic commitment and prepared the ground well in advance are more likely to act.

**DriveNets:** The main trigger is the need for hyperscale economics. The ability to scale efficiently, to innovated fast and to break the capabilities-expenses linkage.

- 2. As operators move to hyperscaler economics, this new model is typically very OpEx driven and different than the operator business model where they prefer to spend CapEx (due to how they are valued by the market). Do you see this view of operators' financials changing in this area and a "reset" needed with how they are valued by the market?**

**STL Partners:** The move to hyperscale economics does not necessarily entail a move to OpEx although this becomes increasingly possible and likely. Operators need to bring investors with them on the journey.

**DriveNets:** When looking at this move from a TCO perspective, moving to hyperscale economics does not necessarily mean skewing expenses towards OpEx. There are multiple examples in which both CapEx and OpEx were reduced when a cloud-native, disaggregated solution was adopted.

**3. How much do operators anticipate they will build in house, versus run in public clouds? We see this already occurring particularly at the edge with many recent partnership announcements between operators and public cloud providers?**

**STL Partners:** We expect that for the next 3-5 years most operators will run most (network) workloads in-house and gradually blend in more public cloud infrastructure (e.g. IMS workloads). This could change quite quickly if public cloud (including edge) proves itself with some of the early adopters.

**DriveNets:** There are multiple considerations leading to this selection, including cost-performance, scalability, security and time to market. In general, when it comes to data-plane network functions such as core-routing, peering, UPF, BNG etc., compute-based platform such as the public cloud infrastructure do not suffice, in terms of performance and scalability. The fact is that you can do everything with CPU, but you cannot do it all efficiently. This is why you have Network Interface Cards (NIC) and smart NICs, and GPU and now also Network Processing Units (NPU) to take care of specialized tasks. Networking as in our case, is such an example.

**4. Do you have any examples of telco operators in mind that have already made progress in achieving hyperscale economics?**

**STL Partners:** Some of the enterprise & wholesale players such as PCCW Global and CityFibre come to mind, given the competitive market they are in and the types of customers they serve (including the hyperscale cloud providers themselves). Rakuten and Jio are also examples of operators that have embraced hyperscale economics.

**DriveNets:** AT&T is the most significant one. Many others are there, though we cannot disclose their names, yet.

**5. What do you see as the main challenges for a telco operator to move from an “on-premise” environment controlled by the telco, to utilizing a cloud (aaS) environment??**

**STL Partners:** This assumes that operators want to make such a move, which is not always the case. At least, not yet for network functions. For those who do, the challenges will be around performance (particularly for data plane functions) and total cost of ownership (TCO). Challenges around regulation and the need for data sovereignty may also constrain them.

**DriveNets:** See answer to #3 above

**6. CORD by the Open Networking Foundation (ONF) has been a type of first blueprint for cloudifying the fixed network edge. What has changed/evolved since then? What may have gone wrong or were they too early?**

**STL Partners:** The challenge is that rearchitecting central offices as data centres requires significant investment and the business case is often unclear. In many cases, the premises are not especially suitable in terms of space, power consumption, racking, cooling etc. so often it becomes easier to build 'new' instead. Interesting to note too that M-CORD later emerged, which is built on the CORD infrastructure platform, which may have addressed some integration problems with the mobile network.

**DriveNets:** CORD gave a blueprint of how to build a cloud in closer proximity to the network's access areas. It moved the overhead of building a cloud to the network experts which in turn takes away the value of their key asset which is the network. CORD wasn't wrong nor early, but they missed out on the key assets of the potential users, network know-how and real estate attributes. In a way, OGA (Open Grid Alliance) is offering an evolved approach where the cloud knowhow is built in with the solution and the IaaS offering already integrates the real estate aspect and the intelligence of the network plays an effective role.

**7. Other vendors such as Cisco have claimed that big iron core routers will stay as a domain of specific hardware. Where do you see the limits of scaling a cloudified approach in comparison with the box solution?**

**STL Partners:** Without getting into specifics of different vendors' offerings, we think that this distinction is blurring. Merchant silicon is being adopted by many vendors who are evolving to support greater disaggregation.

**DriveNets:** As proven in the AT&T core network, this is not the case. The DriveNets Network Cloud platform scales up to 691Tbps capacity. This is more than enough for any core network on earth. A distributed model will **always** out scale a monolithic chassis simply because it removes the metal enclosure that acts as a limit for a chassis to scale (# of line cards, cooling, power, connectivity, etc).

**8. How do the commercial aspects of vendor agreements impacting decisions to move to cloud native moves?**

**STL Partners:** In the minds of many stakeholders within operators, cloud-native disaggregation is closely linked with changing their relationship with vendors (e.g. more opportunity to consider best-of-breed). The procurement teams may see this as a means of strengthening their negotiating position, but they will find they no longer have a single the 'neck to choke'.

**DriveNets:** Just as the architecture changes, so does the engagement model(s). Operators can make the leap and engage separately with different providers of various solution pillars, or mitigate this change by engaging with a system integrator.

**9. [To the DriveNets team] Once you disaggregate the router, do you also see or plan to have intense mobile core applications – e.g. EPC or 5G Core – to run on this disaggregated model?**

**DriveNets:** Any user-plane function (such as the 5GC UPF) can leverage the networking optimized capabilities of the network cloud.

**10. On the economics of disaggregated architecture, how do we mitigate the interconnect costs between leaf spine since price points are skewed in favour of optics price than port price as we observe on the journey towards 100G to 400G ports and beyond?**

**DriveNets:** Fabric interconnect exists in chassis same as it exists in a distributed chassis. The difference is that within a chassis the cost of it is integrated within the chassis cost and you don't see it as a line item but it's still there. In a distributed model, this "backplane" is sold separately which makes it dynamic, repairable, evolving, and multi technology.

To the point asked, you can use optics or copper (passive or active) on these interfaces depends on the distance you require in order to control your cost. We have demonstrated a router of 350Tbps (world's largest router) with using Active Electric Cable (AEC) only and we deployed over clusters which used optics only and had a relative price tag of 3X that of an AEC. The point is that the customer has his own say here and can choose his preferred option, our business is of enabling this choice and not by burying hidden costs and charging for them elsewhere.

# PARTNERS



Research



Consulting



Events