

The logo for SANDOG, consisting of the word "SANDOG" in a bold, white, sans-serif font, centered within a solid black rectangular background.

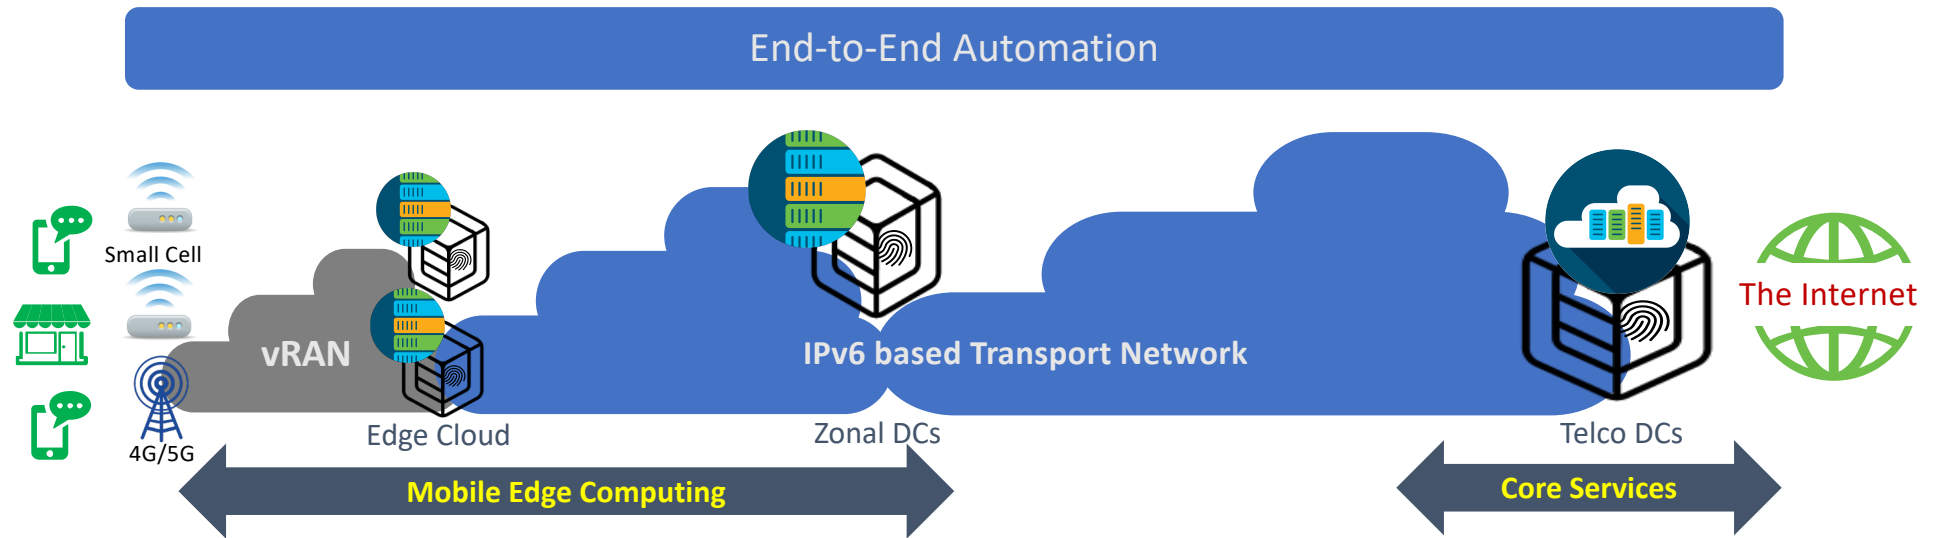
SANDOG

An Outage Incident in a Fully Virtualized Mobile Network and the Learnings from it

Santanu Dasgupta

Distinguished Systems Architect, Cisco Systems

About the Operator



- A fully virtualized (and automated) mobile operator in Asia-Pacific region
 - Network Virtualization from Radio Access (RAN) to Core
 - End-to-End / multi-domain orchestration
- End-to-end IPv6 based

The Incident

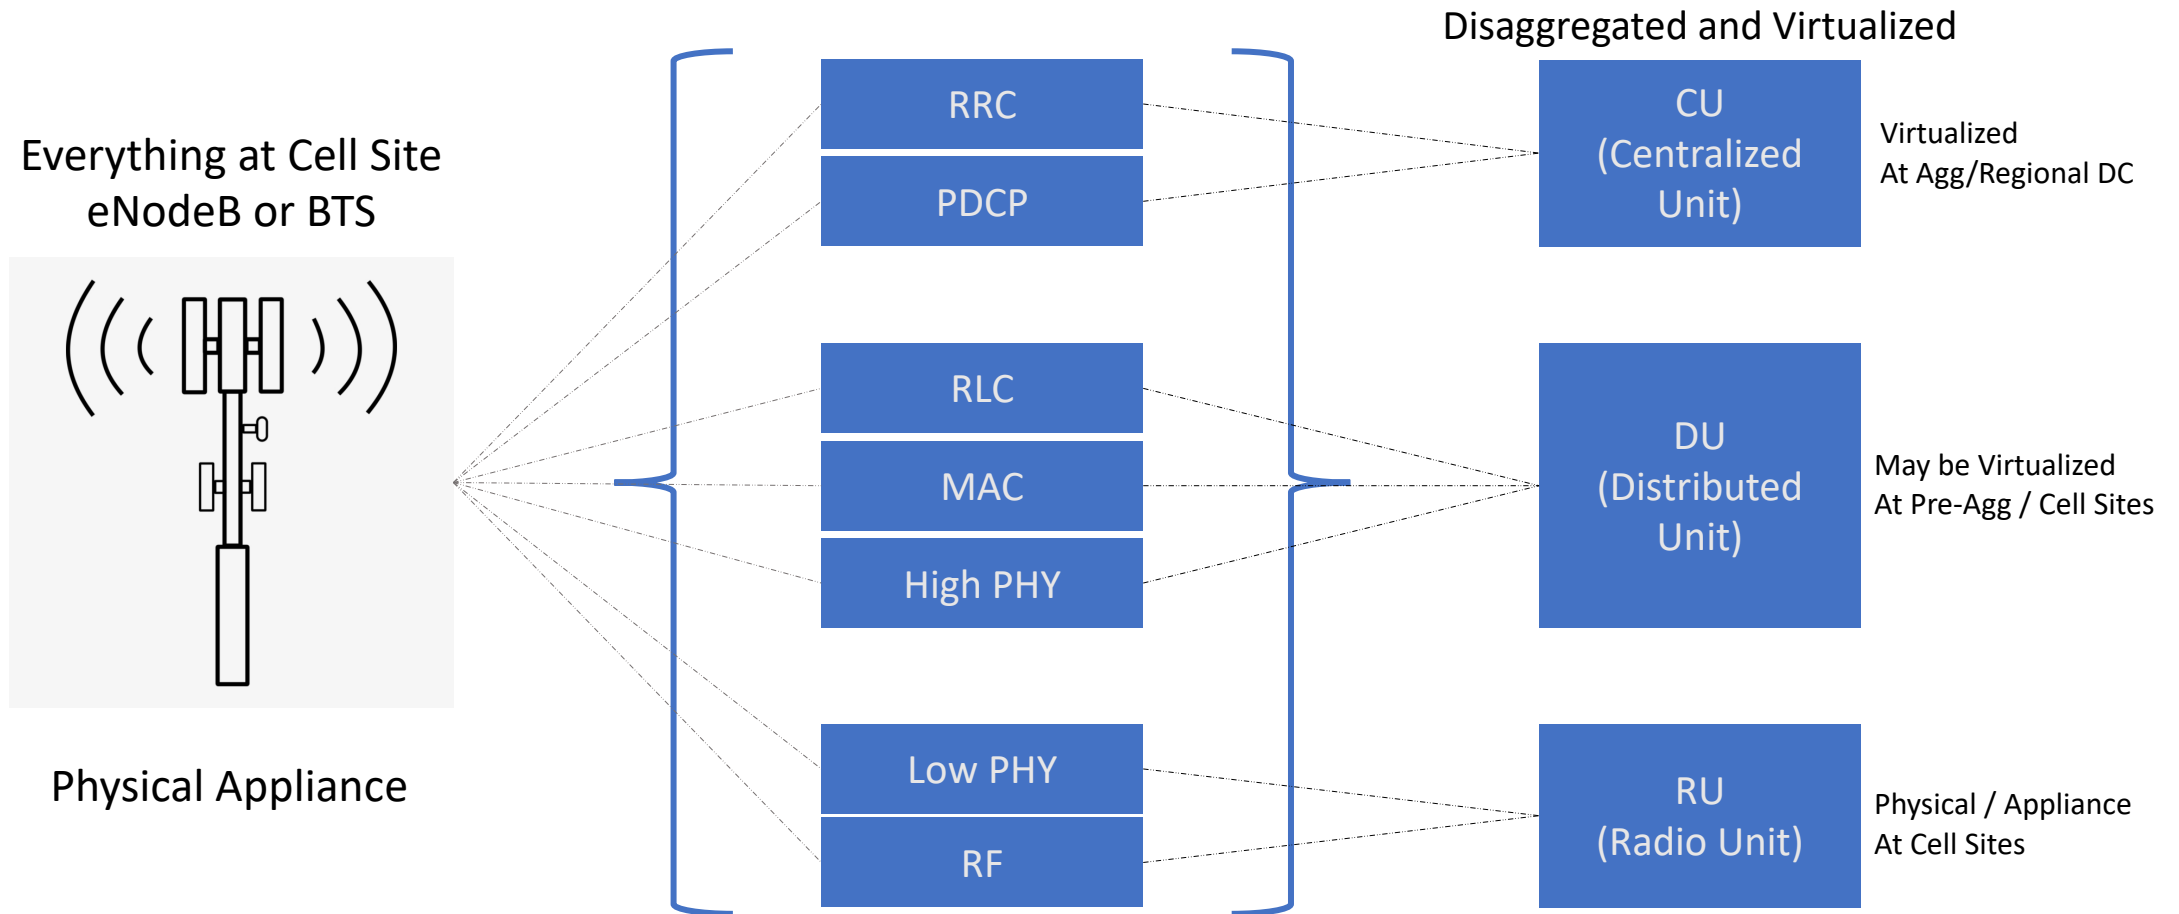
- Happened during User Trial Phase (before commercial launch .. phew)
 - Less than 1000 cell sites were operational at that time
- Once in every 2-4 days, upto 50% of the cell sites were going down
 - Random time of the day (and day of the week)
 - Starts with a few, then upto 50% sites goes down one by one, all within ~15 min
- Major outage – imagine half of the mobile network in your city going down ;-)
 - Became a Sev1 escalation in almost no time

Some Basics – Radio Access Network

Traditional Deployment Model

Radio Access Network Stack

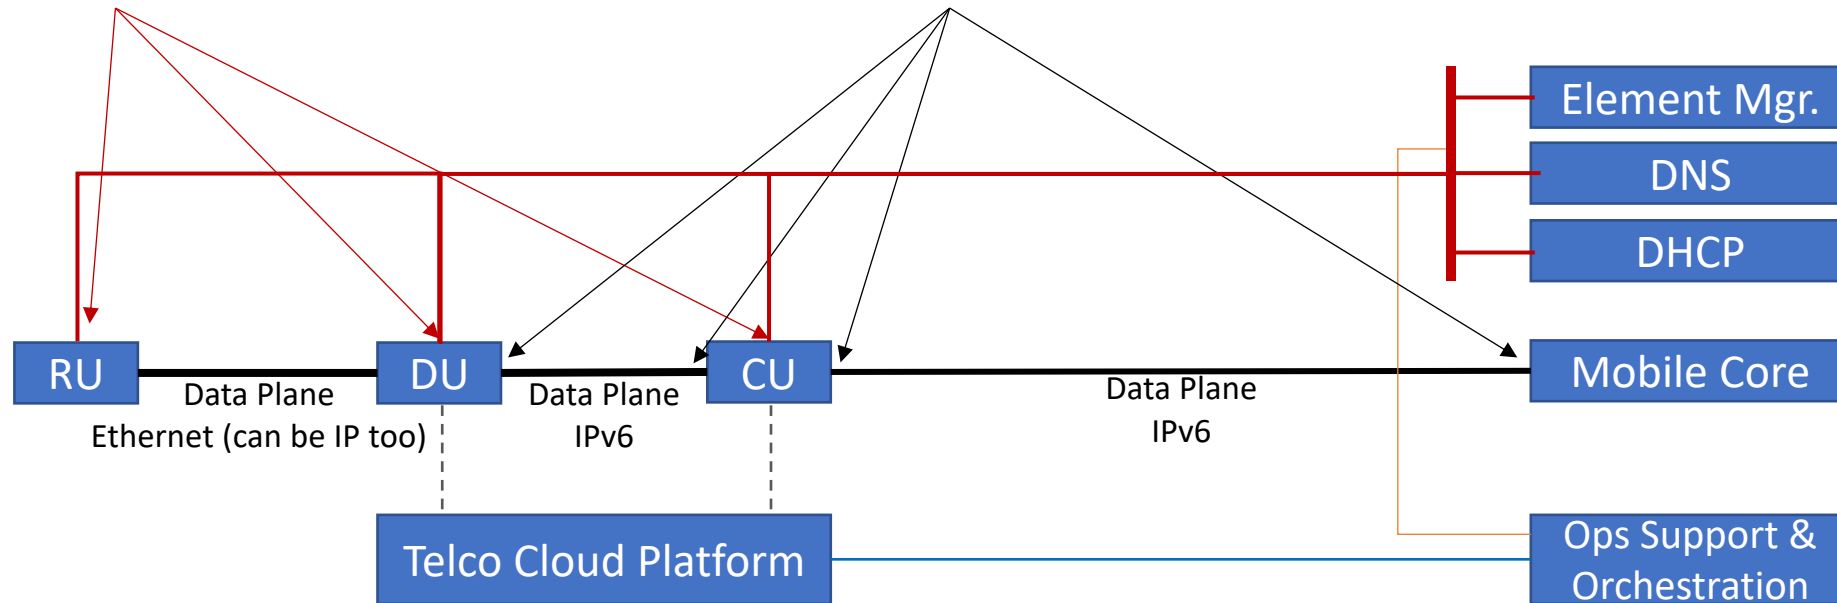
Evolving Deployment Model



Open/Virtual RAN Deployment

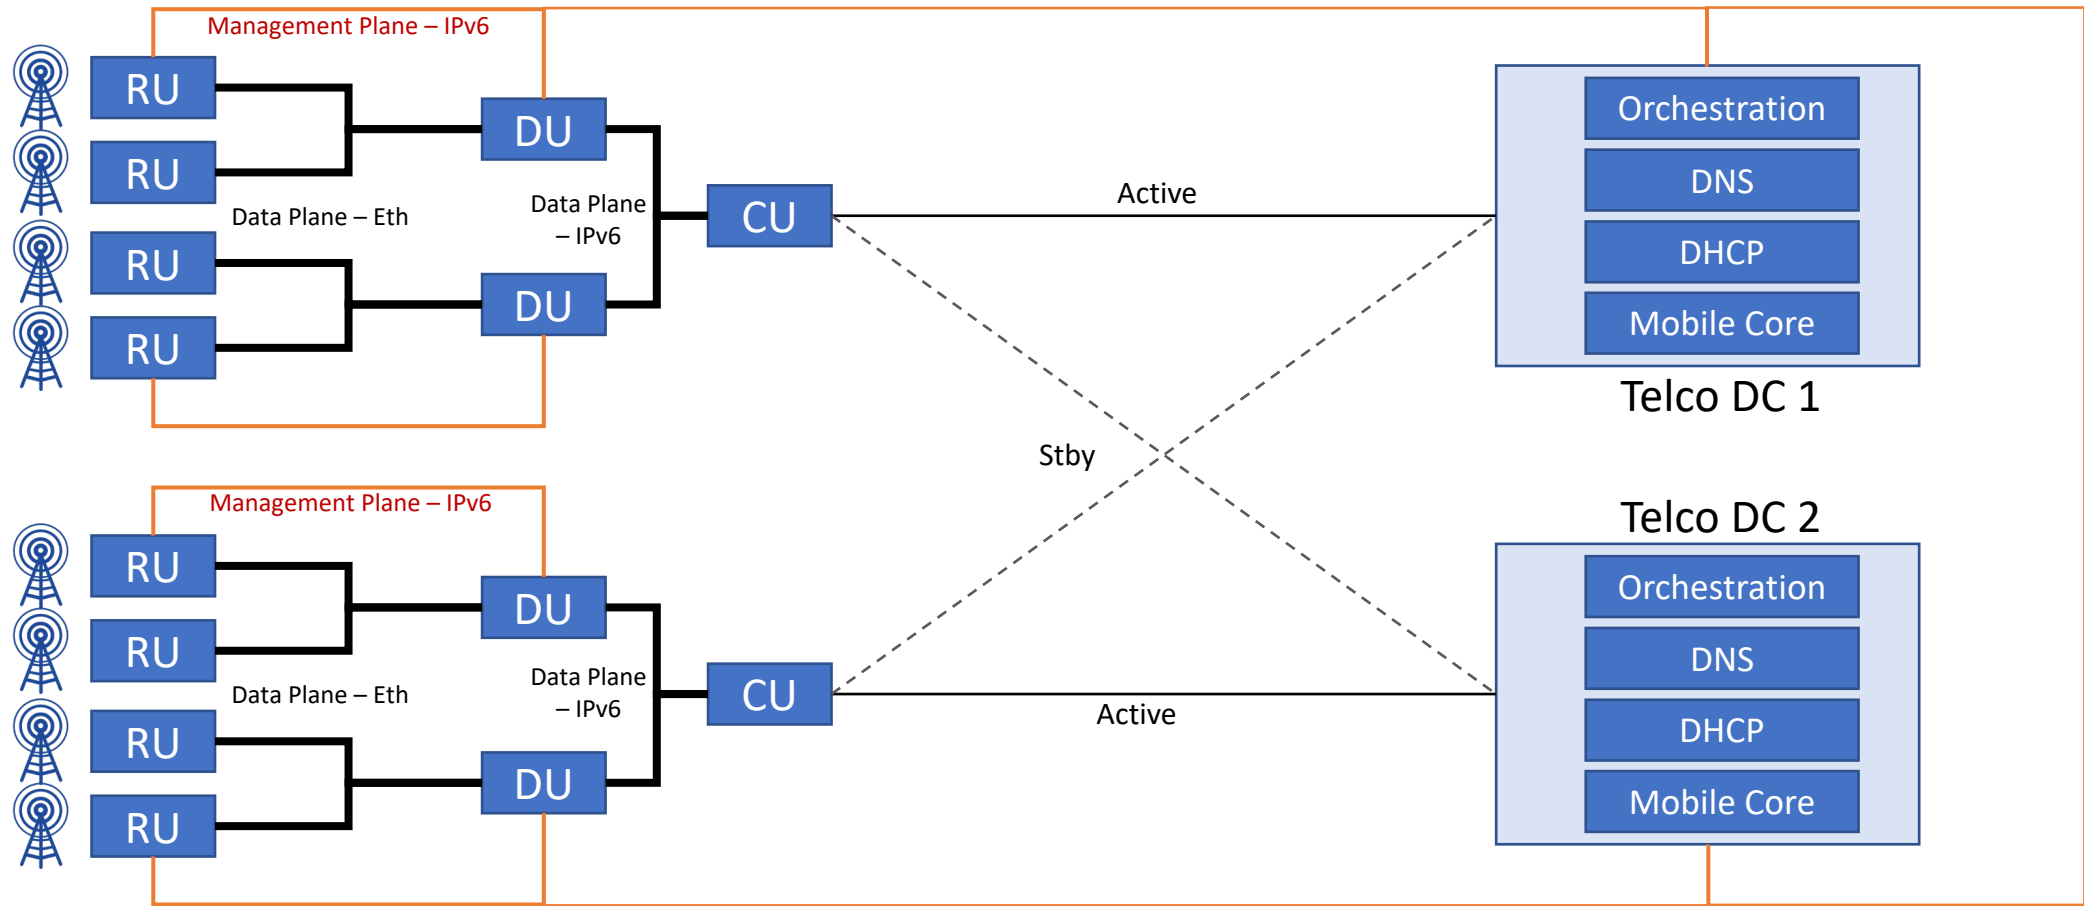
These Management Plane Interfaces has IPv6 addresses

These Data Plane Interfaces has IPv6 addresses



A disaggregated and virtualized system architecture is typically more complex

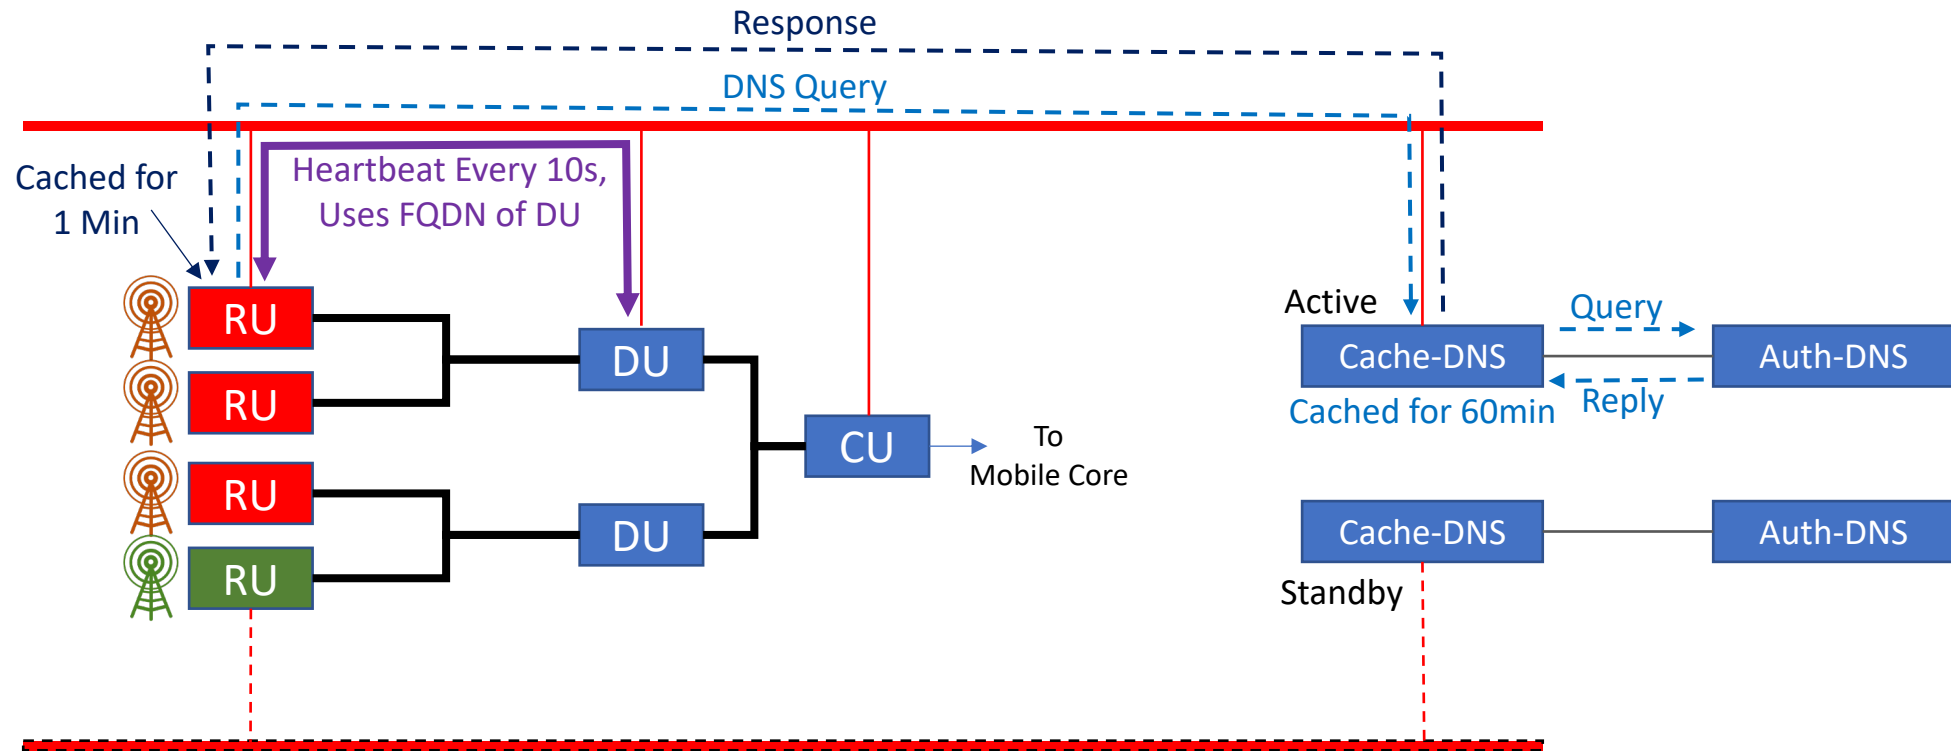
Resiliency of Control and Core Services



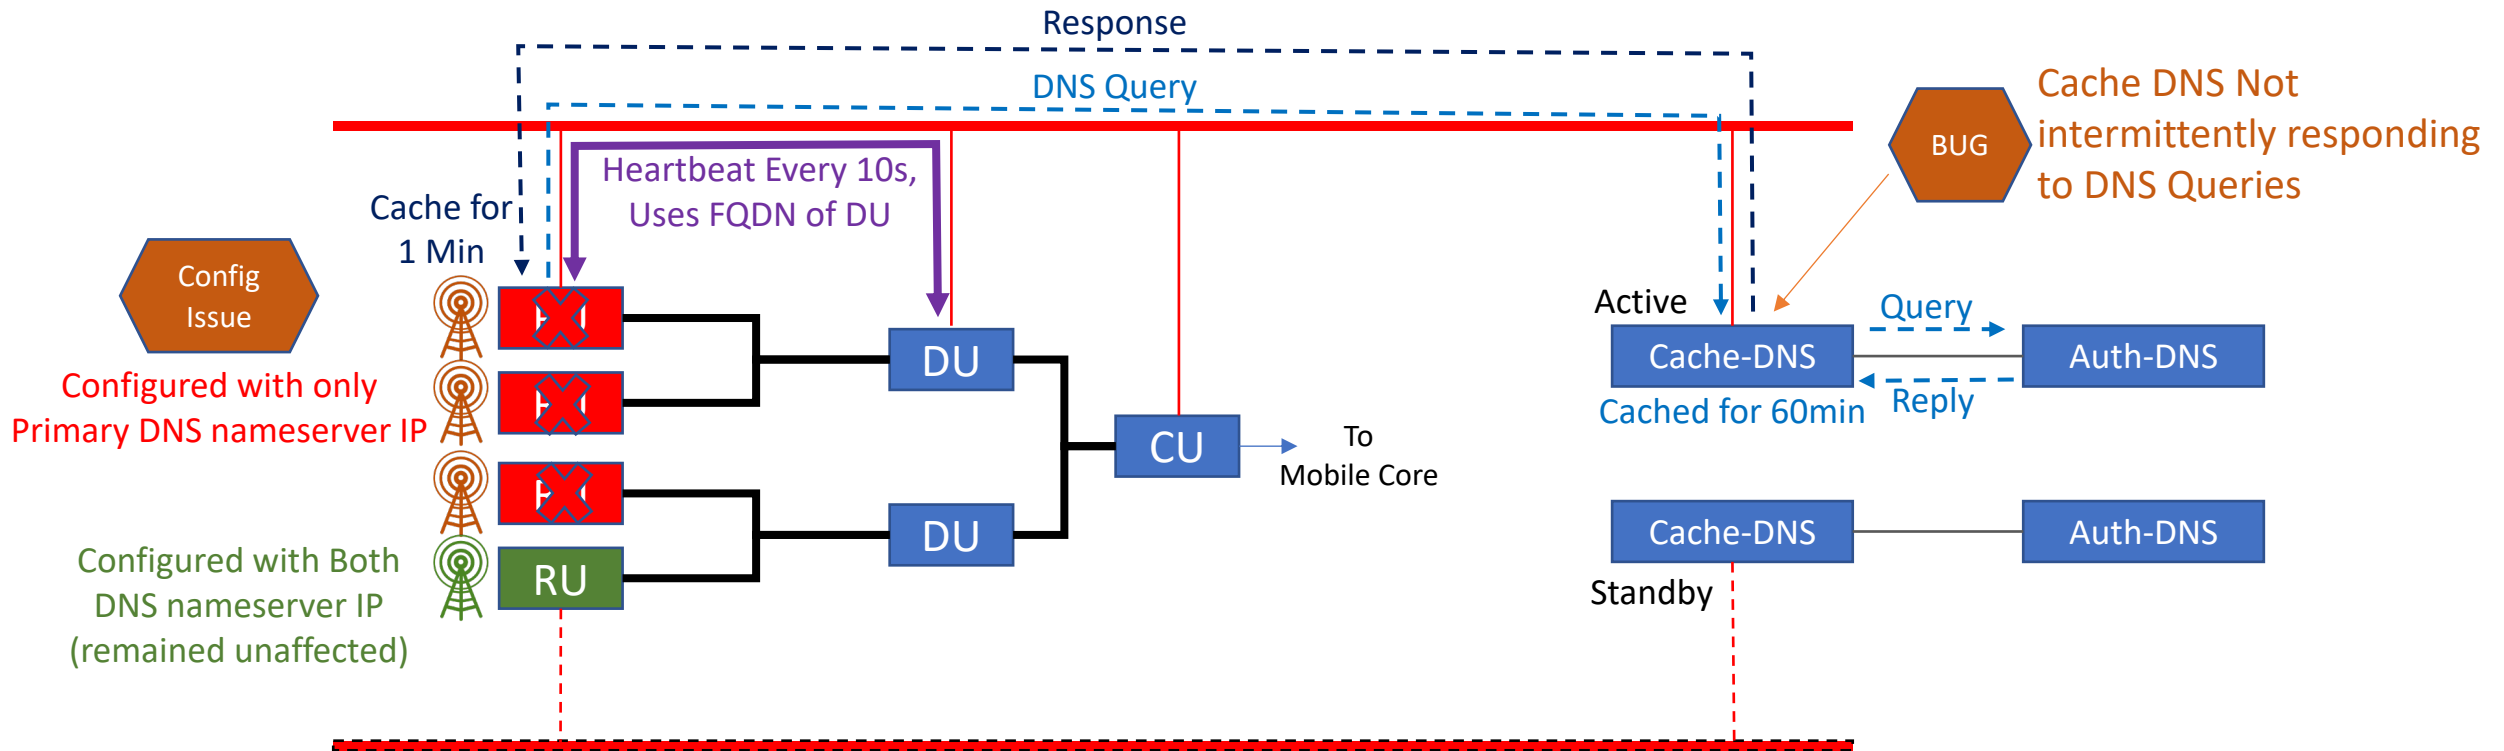
Troubleshooting to Identify Root Cause

- Initial troubleshooting took almost 72 hours
 - Was a manual process, as we lacked a proper tool chain/instrumentation (much better now)
 - We identified (almost) all possible elements that can cause failures and checked on them
- Primary debug found no major red flag on initial checks / log scrubbing
 - Routing / Connectivity, Mobile Core, RAN software, Telco Cloud platform – all seemed ok
 - We had no choice apart from waiting for the next outage to happen to collect more data
- During next outage, we collected more logs from all potential elements
 - Upon a comprehensive log scrubbing, we found some DNS timeouts on the RU (Radio Unit)
 - All elements including RU's use FQDN to reach out to their DU & CU
 - We picked random RU's from failed lots, put static entries in /etc/hosts to avoid DNS queries
- During the following outage, all those RUs with static entries remained up!

Root Cause – What Was Going Wrong with DNS



Root Cause – What Was Going Wrong with DNS



Key Takeaways

1. Strong knowledge of networking and fundamentals are still foundational – solidify that foundation and build on top of it
2. Think about Operations in every stage of design and deployment
3. Having a right tool chain and instrumentation are extremely critical for ops
4. Decomposed architectures tend to bring more complexity, we need to be ready
5. Pay close attention to cloud & virtualization in Networking – it is the future!

Thank You!