Real-Time Communications
Without Boundaries

Using the Cloud to Transform SS7

How to Take Advantage of the Cloud to
Change Existing SS7 Architecture

White Paper

Contents

Introduction 03
Traditional SS7 Network Architecture 03
- Redundancy within a mated STP 03
Features of Cloud Architecture 04
- Supporting TDM in an IP-based Cloud deployment 04
- Resiliency 05
- Geography 05
- MPLS Network 05
- Elasticity 05
New Architecture 06
- The case for mated pairs when replacing legacy architectures 07
Conclusions 07
About Ribbon Communications 07

Introduction

The Signaling System #7 (SS7) network was conceived in a hardware-based, pre-IP world, but this has been changing for years. The industry has fully embraced SIGTRAN to transport SS7 over IP, and is now offering SS7 on cloud-based technology. In this paper, we examine moving the architectural decisions from a hardware-based mentality to take advantage of a cloud-based infrastructure. Specifically, we will be looking at what changes we can make in network architecture based on the increased resiliency and location independence that the cloud offers.

Traditional SS7 Network Architecture

In Figure 1 below, we simplify the network to show a Public Land Mobile Network (PLMN) with two Signaling Transfer Points (STPs) acting as both edge and core STPs, and one application. Clearly any real PLMN will have more than one application, but this simplification will help our discussion. We also show two external STPs, representing SS7 network connectivity to a partner carrier or IPX provider.

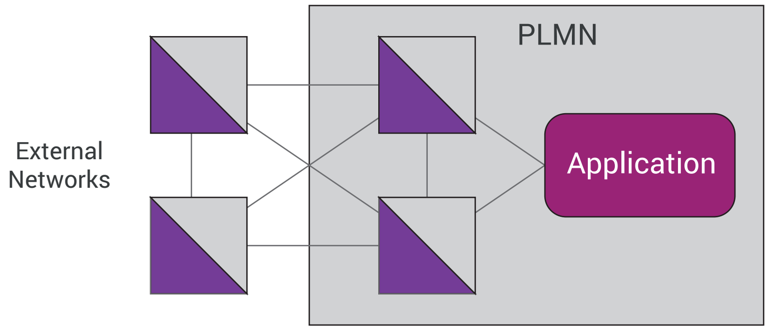

Figure 1: Basic SS7 Network Diagram

The two STPs in the PLMN are mated. By mated, we mean that they are expected to have similar connections and similar translations provisioning. These STPs are typically provisioned independently, although some applications such as Gateway Screening or Global Title Translation (GTT) are often jointly provisioned, because they contain the exact same data. The mates do not have the same Local Point Code and are connected by C-Link to their partners, thus Message Transfer Part Level 3 (MTP3) provisioning will be slightly different on each mate.

The mates are assumed to be geographically separate. The geographical locations of those mates in traditional architecture are determined in consideration of the following items:

- Geographically separate enough to isolate units from natural or other disasters.
- Separate electricity and network sources.
- The cost of backhaul of TDM links across distance.
- Allowing staff to conveniently access both systems.
- Where the PLMN has a point of presence.

Redundancy within a mated STP

Each STP vendor has a slightly different architecture, so we will again present a simplified view, which can be expanded as needed. Each mated STP should be redundant to at least 99.999% availability. In a hardware-based system, this implies redundant storage, Central Processing Unit (CPU) and TDM Input/Output (IO). (Note: we will intentionally leave OAM and IP switches out of this simplified view.) In Figure 2, we are showing CPU cards that perform both SIGTRAN and SS7 processing. Presumably more IO cards or more CPU cards would be added for extra capacity, but at least two are required for redundancy. Again, not all vendors will have this simplified architecture. We also assume that all components run in an Active/Active manner. The following on the next page shows the simplified view.

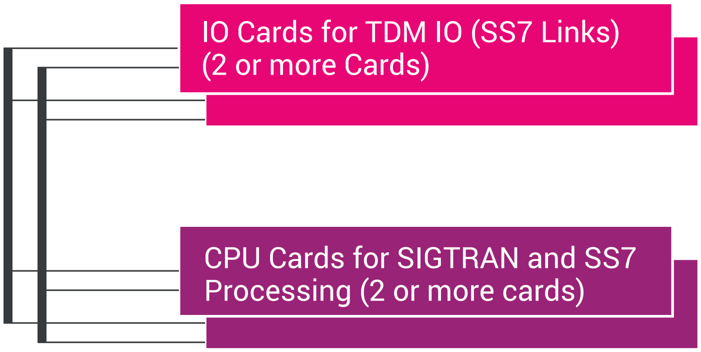

Figure 2: Simplified Internal View of an STP

In the above diagram, we show links from each type of card. There is an assumption in traditional architectures of having dual links to each external location. This allows use of both cards to reach the destination, and assuming upgrades are on a card-by-card basis, allows for connectivity to a destination to continue during software upgrades while a specific card is being upgraded, or otherwise fails.

Additionally, we expect to see C-link connectivity to the mated STP. This provides yet more paths to a given destination. The C-link is likely a dual link, so this can be thought of as two more paths, not just one.

Features of Cloud Architecture

We now look at some features of the cloud that will bring us advantages in looking at new architectures from the above.

Supporting TDM in an IP-based Cloud deployment

TDM links are still a requirement in most SS7 networks. M2UA is a SIGTRAN protocol used to provide backhaul of MTP Level 2 user signaling messages and service interface across an IP network. This means M2UA is the method to provide support for TDM links with remote I/O cards. M2UA is not a cloud feature per se, but we include it in the discussion, because it is required where TDM links are required. M2UA as a protocol only supports Low Speed Links, and does not support ATM and Annex A.

Figure 3 shows M2UA links in relation to other SIGTRAN link types.

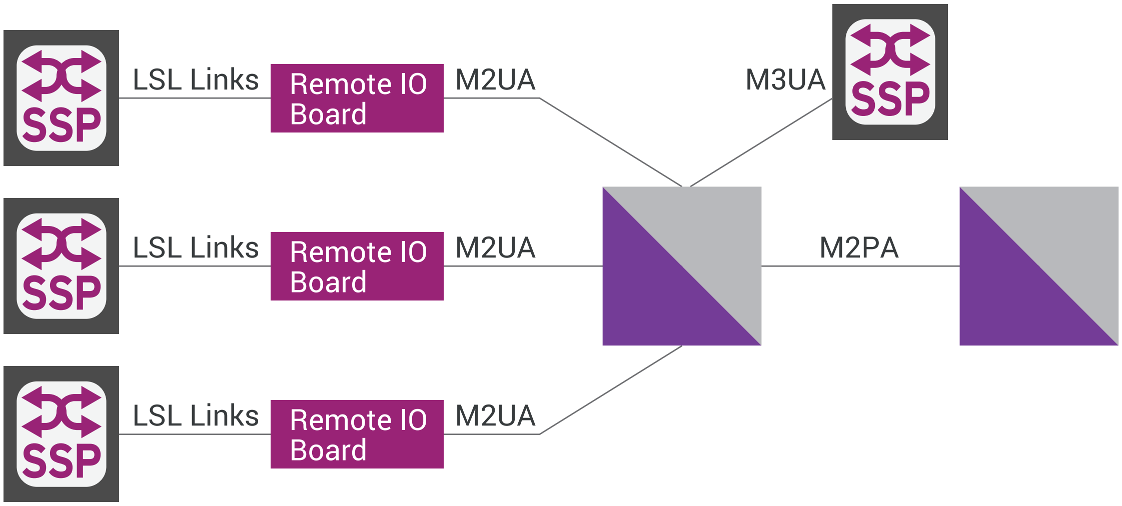

Figure 3: M2UA and Other SIGTRAN Link Types

Resiliency

In traditional hardware-based systems, loss of a card was a major event requiring time-consuming staff intervention. After a period of troubleshooting, the card may have had to be manually replaced, and a backup of a previous image may have been required. Thankfully, with modern hardware, card failures are extremely rare, but can still occur.

In a cloud environment, we expect the cloud infrastructure to be able to bring up a new Virtual Machine (VM) instances of specific STP functionality with the exact same configuration very quickly. We expect to see the same IP addresses, even MAC addresses, software versions, and provisioned data.

In an industry where we measure downtime in the order of seconds over years, this is a very welcome change to the way we work.

Geography

Where to put the mates of an STP pair in hardware-based architecture was the subject of much study, as discussed above. In a simplified view of the cloud, an application, such as an STP, is going to assume the underlying geographic redundancy has been taken care of. However, in these early days of cloud implementation, that faith may be overly optimistic. As an application in a cloud deployment, it is wise to be vocal and ask the cloud administrators to ensure geographic redundancy.

MPLS Network

Using the cloud can solve other geographic problems. In looking at where to place mates in a traditional network, it was important to consider the cost of backhaul of the TDM links. However, with ubiquitous IP access due to modern Ethernet and MPLS networks supporting cloud connectivity, we can assume reliable and fast connectivity for any location for the STPs. We will also use this assumption for the placement of M2UA-enabled TDM IO boards in remote locations to support legacy TDM switches.

Elasticity

Expanding the capacity of an STP deployed in the cloud should be as simple as applying more resources to a VM, or allowing the system to expand the number of VMs. Given the static nature of SS7 SIGTRAN links to specific IP addresses and the CPUs behind them, elasticity may still today require some human intervention in deciding where links should be terminated.

The diagram in Figure 4 shows this elasticity. The diagram shows COTS servers, but it could equally apply to the cloud. It shows an STP pair still in a mated configuration, but each has been expanded. In the top row, we have two CPU blades in the same server for each mate, for a total of four CPU blades. The "M+R" VMs are doing both Management and Routing, and the "R" VMs are doing Routing only. At this level, we have redundancy and, assuming the CPUs have enough core processing, we have expanded capacity.

In the second row, each mate has had its capacity and redundancy further augmented. The "M+R" VMs have not been added at this level, presumably because a busy STP does not require more management power, just more routing power. In this case the new VMs are on new blades and new COTS servers.

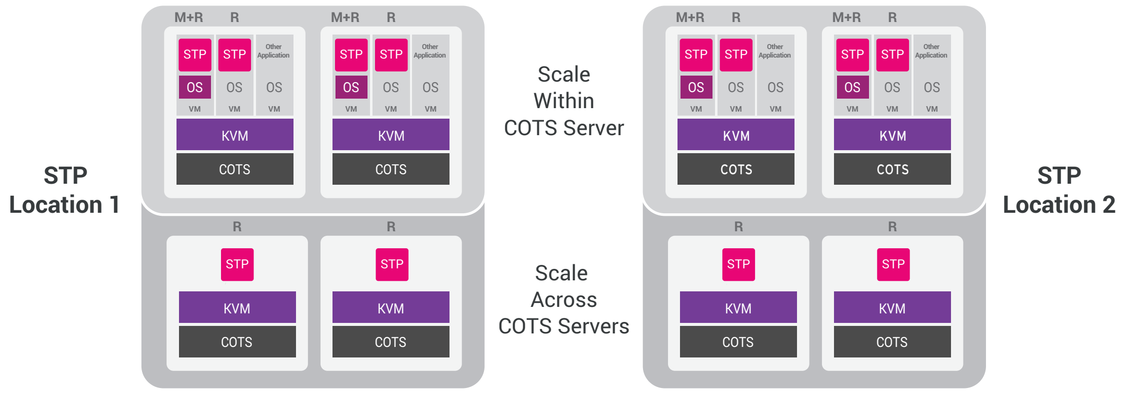

Figure 4: Scaling of STP functions across Hardware or within the Same Hardware

New Architecture

Given the above cloud capabilities, we must start to ask ourselves, why would we implement the same old mated pair hardware-conceived architecture in a cloud environment? Instead of Figure 1 above, we present an alternative architecture as shown in Figure 5. In the diagram, we show one STP with one LPC. It happens that this STP contains four VMs rather than two CPUs.

When working with the cloud, we think of it as one unit, but in reality, it may be over several locations and certainly over several servers. The diagram below shows our desire to have this STP spread over geographic and redundantly separate systems. We show this on the diagram as four VMs in two locations.

Overall, this provides the same number of VMs as the traditional architecture of Figure 2. It has the same level of redundancy, but has improved recoverability and elasticity. It is also simpler to operate and maintain, because rather than two STPs being separately provisioned, the four VMs make up one STP.

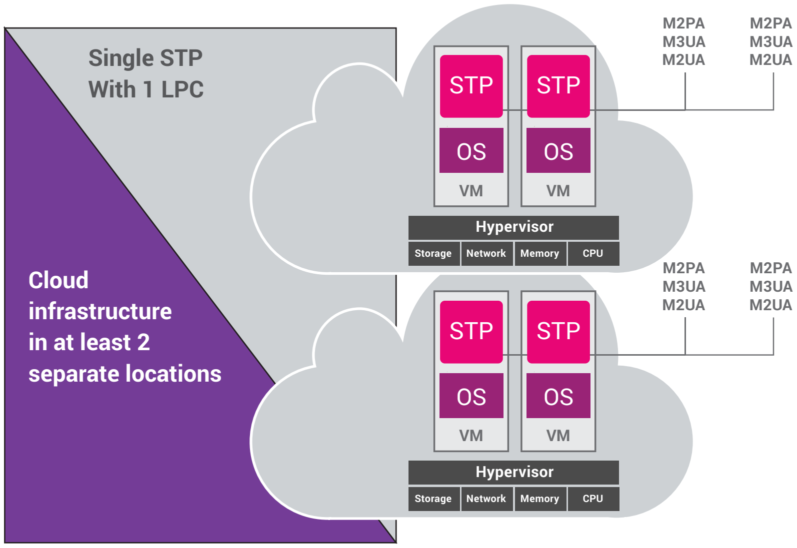

Figure 5: Cloud-based STP Architecture

Other advantages the cloud brings us are when we take the STP internal structures and start to design them to use the tools of the cloud. The cloud enables us to do more work with load balancers and clusters of servers. In Figure 6, we present a possible architecture combining SS7 and Diameter routing with number portability and other servers.

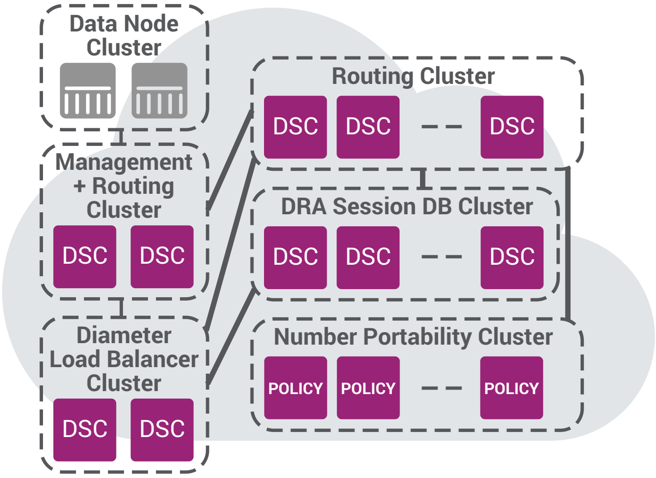

Figure 6: Cloud-based STP and Diameter Architecture

The Case for Mated Pairs When Replacing Legacy Architectures

There is one common situation where a carrier may wish to retain the mated pair configuration even when using a cloud architecture. This is in the case of REPLACING legacy STPs. Legacy STPs will almost always be already in a mated pair configuration. If the carrier's goal is to simply transition from legacy STPs to cloud-based STPs, it may be simpler during that transition to keep the mated pair configuration around which the current architecture and all its translations and links are currently based. Operators need to be sure they understand the trade-off for this option, where they may minimize migration issues versus the inability to take full advantage of the cloud-optimized solution shown in Figure 6.

Conclusion

The new features and functions that a cloud architecture provides allow us as an industry to redefine the deployment of the Signaling Transfer Points. This applies more for greenfield STP deployments than for legacy mated pair replacement projects. During these early stages of cloud deployment, this will likely apply to greenfield STP deployments or expansion of STP capacity at an existing STP site. Each carrier needs to examine their own architectures, which reflect their network requirements and the capabilities of their cloud environments, to decide how to deploy STPs in the cloud.

About Ribbon Communications

Ribbon is a company with two decades of leadership in real-time communications. Built on world class technology and intellectual property, Ribbon delivers intelligent, secure, embedded real-time communications for today's world. The company transforms fixed, mobile and enterprise networks from legacy environments to secure IP and cloud-based architectures, enabling highly productive communications for consumers and businesses. With locations in 28 countries around the globe, Ribbon's innovative, market-leading portfolio empowers service providers and enterprises with rapid service creation in a fully virtualized environment. The company's Kandy Communications Platform as a Service (CPaaS) delivers a comprehensive set of advanced embedded communications capabilities that enables this transformation.

To learn more visit RibbonCommunications.com