

Achieving dual fault chassis tolerance for SS7 solutions

White paper

Ian Colville, Product Manager, Aculab

Introduction

When deploying equipment onto public telephony networks, operators and service providers demand a minimum level of service often expressed as ‘five nines’ availability. To meet that criterion for user service accessibility, systems have to be designed with reliability and fault tolerance in mind. Continuity of service is the paramount factor, such that a system must remain viable when a failure occurs – whether it’s associated with hardware or software.

This requirement is especially true for signalling system number 7 (SS7) networks, where the integrity of signalling components and preservation of an end-to-end signalling path is vital. In order to achieve a suitable level of resilience, network architects will have designed the infrastructure of the core network using multiple signalling paths to each end point and implementing multiple links (and link sets) between adjacent signalling points. In the case of application platforms interconnecting to an SS7 network, such as a prepaid calling card system, there exists a similar need for resilience and reliability.

A basic method is for the processing of SS7 terminations at a single signalling point to be distributed over a pair of links in a single E1 or T1 trunk. Taking this one stage further, the links are dispersed between separate trunks on a digital network access card. Card level redundancy is gained by splitting links between multiple cards in a single chassis – if one card fails, another card is available to maintain a signalling path.

A more fundamental approach is to duplicate the SS7 interface for a single signalling point in separate, but interconnected chassis in a dual, redundant capacity, which means one chassis continues to operate if the other fails. Additionally, separating applications from the signalling interface gives a further degree of resilience or fault tolerance and, as well as distributing or sharing the processing load, enables the effective and efficient scaling of call or transaction handling capacity.

Implementing an SS7 signalling point across two chassis

Implementing an SS7 signalling point over two chassis isolates discrete failures, allowing a system to continue to operate on failure of one of the chassis or on failure of a signalling route into the network.

The Aculab approach implements a dual MTP method, where a point code is shared between two signalling nodes, in separate chassis, each operating up to the MTP3. In this fault resilient architecture, the ISUP (and/or TCAP) application hosts are also physically separated from the signalling nodes – in multiple distributed chassis interconnected over a LAN.

Signalling links

Normal practice is for signalling links and link sets to be provisioned such that in the absence of failure they operate within 40% capacity. In this way, the ability of remaining links to carry a full traffic load, when some routes fail, is assured.

Aculab’s approach is to apportion multiple signalling links between two physically separate chassis. The links are split over physical E1 or T1 trunks on separate Aculab cards and resilience is achieved in that there is no reliance on a single delivery trunk from the network. This gives redundancy at link level between cards in a chassis and across the chassis pair. To achieve route resilience to a given destination, two or more link sets are configured, each separately connected to both nodes of a mated signal transfer point (STP) pair.

The essence of such a dual resilient set up is that in the event of a failure in one chassis, signalling is maintained by the remaining chassis. Procedures such as load balancing and load sharing are applied to ensure traffic is evenly distributed across links and link sets.

The dual MTP approach

With the MTP running separately on two interconnected chassis, there is a need for inter-chassis communication. Inter-chassis signalling links are established over a dual Ethernet LAN and are used for two key functions. Firstly, they are used to share routing information between MTP3 A and B – see figure 1. Secondly, and importantly, they are used to reroute outgoing transmit messages via the other MTP chassis in the event of failure of a link set between a chassis and its adjacent signalling point.

Beneficially, this architecture means there is no need for inter-chassis signalling links via physical trunk circuits. The benefit being that expensive trunk circuits can be fully utilised for network facing signalling and provisioning need not accommodate additional, costly and unnecessary trunks.

Aculab's dual MTP nodes operate on a peer basis rather than a typical master/slave configuration. Both MTP nodes are continually processing message traffic, instead of having one node lie quiescent until required to take over signalling. In normal operation, an incoming message is processed by the receiving MTP node and sent via the LAN to the appropriate application host where (assuming it's an ISUP host) a call is set up on the appropriate circuit (CIC). For transmit messages, the application directs the data to either MTP A or B, for balanced use of the system, using a simple configuration rule such as odd or even CICs.

Advantageously, there is no need for complex partitioning or sharing of state information and check pointing for transactions at either the MTP or user parts. And there are no tricky circuit group allocations for the developer to be concerned about. There is no need for an application to activate and deactivate mirrored circuit groups in the event of a failure as failover between nodes is automatic.

This design achieves link and link set resilience, or redundancy, without requiring any extra trunks for the inter-chassis link. It also achieves redundancy of the SS7 interface where, in the event of a failure, signalling can be maintained by the non-failing node. The separation, coupled with the distribution of application hosts, also enables software upgrades to be carried out with minimum disruption or loss of service. Beneficially, for the user, this implementation is fundamentally cost-effective and straightforward to implement.

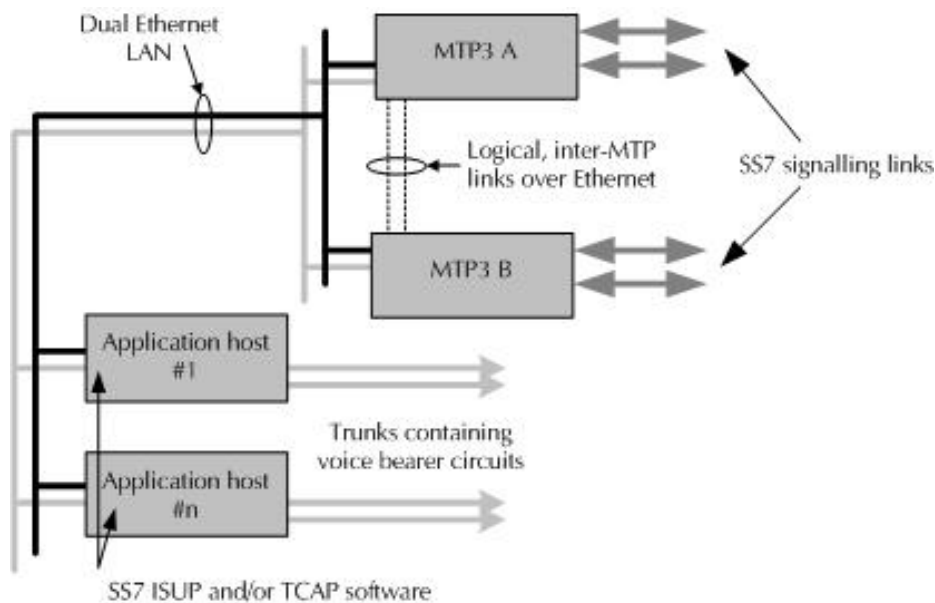


Figure 1 – dual MTP nodes showing signalling links and application hosts

Failure modes

The benefit of the dual MTP approach can be further illustrated by its operation in various 'what if' failure modes. Potential critical points of failure are:

- Failure of signalling links or link sets
- Failure of a single MTP chassis
- Failure of an application host
- Failure of an element of the LAN

Failure of signalling links or link sets

In normal operation, incoming signalling arrives at the MTP pair as if they were a single entity. They present a single point code to the network with load sharing and load balancing ensuring an equitable processing load. If all links to one side fail, through an alarm condition or congestion, for example, all messages from the network are handled by the other MTP chassis. For outgoing traffic, if all links to MTP3 B have failed, messages are routed over the Ethernet to MTP3 A, from where they are routed over an available signalling link to an adjacent signalling point. See figure 2.

The failover between MTP nodes is automatic; being handled by normal MTP operations on the network side, and the LAN-based intercommunication between the MTP nodes and the application hosts on the user side. However, this happens without the intervention of any ISUP or TCAP application. From a user application perspective, the changeover occurs 'under the hood'.

Uniquely, with Aculab's resilient design, the ISUP or TCAP runs independently on each application host and is only concerned with message traffic destined for its own CIC range or resident application sub-system. Either MTP node is capable of sending or receiving messages to/from any application host at any time and is not required to share call state information with its peer.

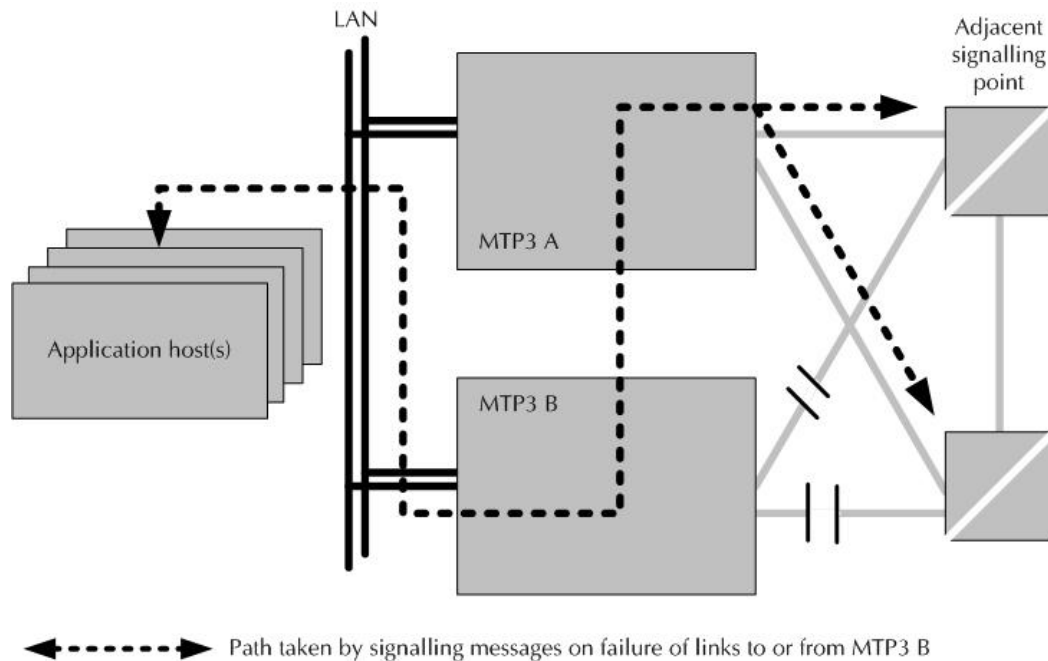


Figure 2 – message paths on failure of signalling links or link sets

Failure of a single MTP chassis

The essence of the dual set up is that in the event of a failure, signalling is maintained by the remaining unit. Incoming messages are automatically routed by the non-failing MTP node and passed directly to the appropriate application host. There is no hiatus during which an application has to enact a recovery process.

For outgoing traffic, if say MTP3 A has failed, the application simply routes the message to MTP3 B, from where it is sent to the adjacent signalling point. See figure 3. Once again, there is no recovery process for the application to perform in conjunction with the non-failing MTP chassis.

In this case, any ISUP calls on either of the application hosts shown in figure 3 are unaffected. This applies to calls in transient status as well as stable, connected calls. No calls are lost, because the ISUP activity associated with those calls occurs on the application host, where the circuits are located, rather than on the MTP chassis.

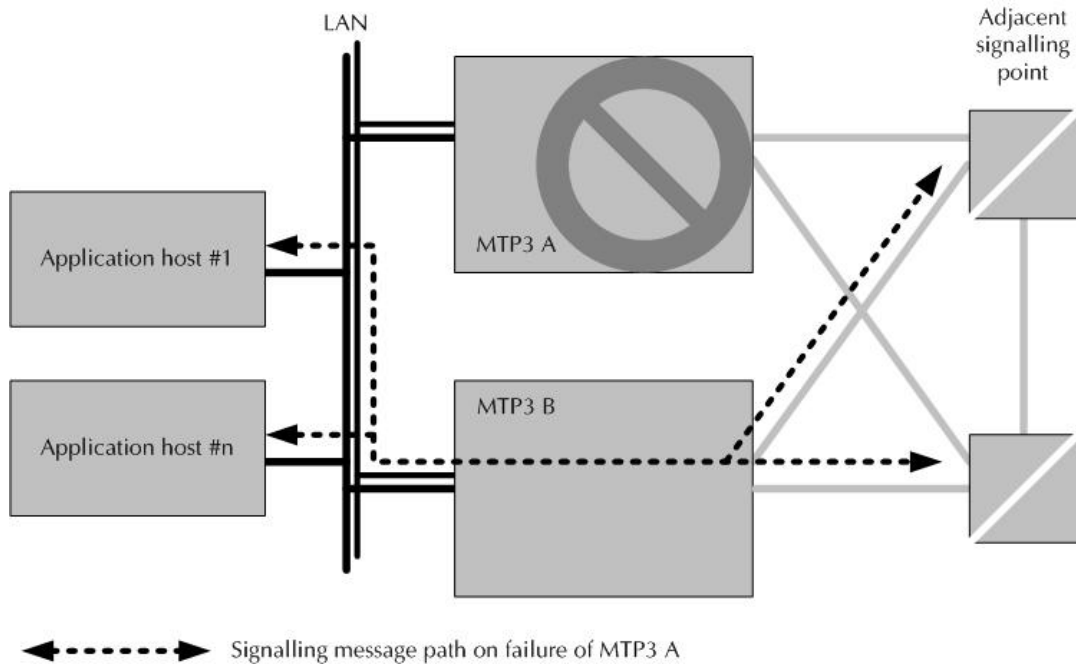


Figure 3 – message paths on failure of MTP3 A

Failure of an application host

The failure of an application host leads to the loss of a portion of system resources and also means effective loss of use of the physical trunk interfaces connected to that host. However, normal control operation will ensure that the affected circuits are blocked, with the result that the SS7 network will no longer attempt to initiate calls in the failing host, as they would naturally fail.

Note that these circuits cannot be readily transferred to another host, nor can signalling messages be sensibly redirected to an alternative host. Other application hosts in the system can have physical circuits to the same destination, however, circuits cannot have duplicate CICs or respond to an IAM for a CIC in a failed chassis. This is where the native resilience of the SS7 network comes into play, knowing a circuit is blocked, it sends an IAM to set up the call via a circuit in another application host.

Aculab's strategy of deploying an application on multiple hosts has several benefits, including that of ensuring that failure of a single application host doesn't mean total failure of the system. And unlike the idea of backup hosts, such distribution of application and physical bearers in multiple chassis provides a good degree of resilience at both application and trunk interface level. Trunks to the same destination should be prudently spread across several chassis to avoid failure of a single chassis taking out all circuit routes to a single destination.

Importantly, when using the distributed application approach, as the ISUP activity associated with a call occurs only on the chassis that hosts the applicable circuit, there is no possibility of other application hosts becoming affected, and problems with a single ISUP host will not affect the operation of the rest of the system.

Failure of an element of the LAN

For maximum protection and elimination of single points of failure, the LAN interconnecting the MTP nodes and the application hosts should be established over a redundant network using two switched LAN networks. Where two NICs are used in an MTP chassis, they must be configured such that each is on different subnets.

At start up each MTP system establishes a connection with its peer using a different route; see figure 4 – routing between points F, D and G, and points G, E and F respectively. The system should also be configured so that in normal operation an application host communicates with one of the MTP chassis by using one LAN segment, whilst traffic to and from the other MTP chassis uses the alternative LAN segment. Therefore, traffic from application host #2 takes the route between points B, D and F towards MTP3 A, and traffic takes a route between B, E and G towards MTP3 B.

If one link from application host #2 fails, say the route between points B and D, all outward bound traffic is routed to the node MTP3 B, and onward to the signalling links via a route between B, E and G. If a different link fails, for example, that between points D and F, then inbound traffic is routed onwards from MTP3 A to application host #2 via a route between points F, E and B.

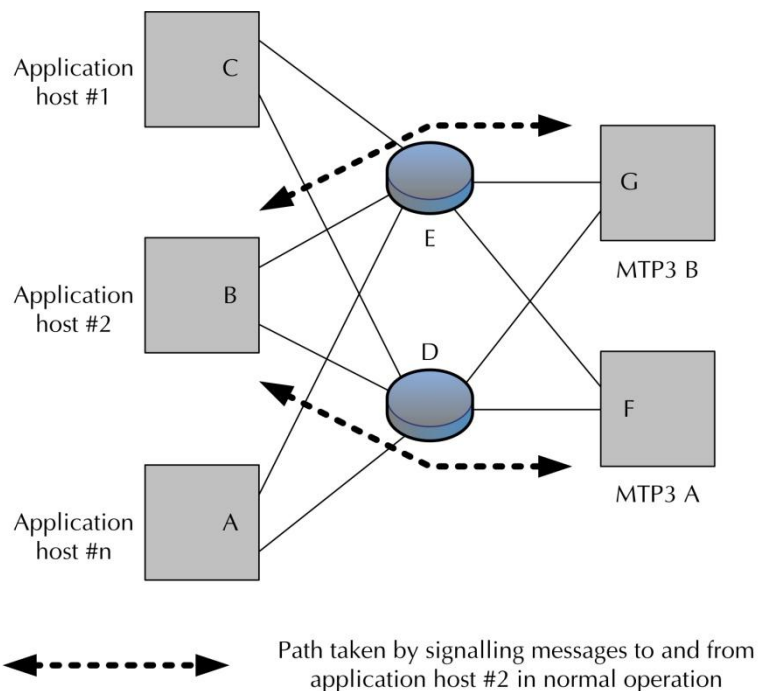


Figure 4 – message paths on failure of an element of the LAN

Conclusion

Implementing an SS7 signalling point over two chassis allows a system to continue to operate on failure of a major component. Whether the failure concerns a signalling link, a link set, a signalling route, a trunk interface, a digital network access card, a single chassis, or an interconnecting LAN element, the resilient architecture offered by Aculab provides a high degree of protection. For further resilience, application hosts can also be physically detached from the signalling nodes – in separate, distributed chassis.

Multiple signalling links are spread between two separate chassis, and the two appear as a solitary network element, behaving as a single point code. Links and link sets are split over separate E1 or T1 trunk interfaces and resilience is achieved in that there is no reliance on a single delivery trunk from the network.

There are several design pointers, which indicate that the dual MTP approach has distinct advantages. These are summarised as follows:

- Cost-effective design – does not require the wasteful provisioning of additional, expensive trunk connections for the inter-chassis links
- Lower cost solution – avoids the extra cost of licensing ISUP software on the signalling nodes (note that Aculab's SS7 software is offered under a cost free licence when used with Aculab's cards)
- Easier integration with user applications – no need for complicated application involvement to swap-over active and inactive circuit groups; failover occurs 'under the hood'
- Safer system operation – no potential for conflict over which circuits are being handled by each half of the signalling node
- Better performance – assured completion of calls that are being set-up, as well as calls in a transient state, if one MTP chassis fails
- Better performance – problems with a single ISUP host will not affect the operation of the rest of the system
- Hitless software upgrades – software upgrades are possible with minimum disruption or loss of service
- No single point of failure – all areas of potential failure: links; link sets; routes; trunks; cards; chassis; or LAN elements are protected



About Aculab

Aculab enables developers and systems integrators to produce a variety of high performance communications solutions. Aculab's portfolio offers an exceptional mix of capabilities that are easy to integrate and bring real value – reduced costs, increased customer satisfaction and competitive advantage. A complete range of open standards building block technologies for use within telco or enterprise environments, as well as essential support services, are offered. Products include media resources, digital network access, VoIP, fax, speech processing and conferencing. Support is available to help developers through each stage of their product's life cycle including pre-sales consultancy, technical support, training and marketing.

Copyright and other notices

© 2007 Aculab plc. All rights reserved.

Aculab, The Groomer, Prosody, "Connected With Aculab", Aculab Worldwide Digital Connectivity and Aculab Your Connection To The Future are registered trademarks of Aculab plc. Prosody S, TiNG and GroomerII are trademarks of Aculab plc. All other product or company references or registered and unregistered trademarks are the sole property of their respective owners.

For more information, visit Aculab's website at www.aculab.com

Aculab plc, Lakeside, Bramley Road, Mount Farm, Milton Keynes. Buckinghamshire. MK1 1PT. United Kingdom. Tel: +44 (0) 1908 273 800

The information in this white paper is provided by Aculab plc, and is for informational purposes only. Nothing in this publication forms any part of any contract.

The information contained herein is based on material, which Aculab, based on its best efforts, believes to be reliable, but no representation is made as to its completeness or accuracy. Aculab make no warranties, express or implied, in this document. E&OE