SLE Routing – Simplified Station Access for Mission Operations

M. Arza[1] and H. Dreihahn[2]
European Space Agency, European Space Operations Centre, Robert Bosch Str. 5, D-64293

In the recent years ESA has adopted the CCSDS Recommendations for Space Link Extension (SLE) transfer services [1] as a standard interface for their TT&C station network (ESTRACK). Both ESA and external space missions interface ESTRACK Network via SLE services, which decouple the ESTRACK Users to a good extent from network implementation details. However, as of today, the SLE User (mission) requires some knowledge of the ground station configuration in order to establish the SLE links to the active communication chain. Furthermore, in case of a failure in the ground station active chain, mission and network operators need to coordinate the eventually required swap to the back-up chain. To overcome this relatively complex process, a discussion to simplify the current approach for the ESA Tracking Network has been started. In this context the idea of SLE Routing has been proposed and analyzed. SLE Routing shall basically route SLE Users' connection to the ground station active chain without requiring the User to know a priori the station configuration; effectively the SLE User connects to the station only.
This paper discusses in detail the requirements for the SLE Routing. Reliability and backwards compatibility, as well as failover scenarios and interoperability considerations are covered. We will discuss these requirements in the light of several implementation options. The importance of minimum impact on the current infrastructure to enable a robust, cost- and schedule-efficient solution is highlighted. The paper concludes with the trade-off of the considered options, which range from solutions based on connection routing to dynamic DNS (Domain Name Service) based approaches, and the preferred candidate solution. Finally, the current status is presented.

I. Introduction

Today each spacecraft pass using a ground station of the ESA Tracking Network (ESTRACK) requires an active decision when the SLE services for telemetry and telecommanding are established. The SLE user, which is located at the mission control center, has to decide to which of the redundant SLE Providers in the ground station the SLE links shall be established. This decision is based on the current ground station configuration, or more specifically, which communication chain configured as prime chain for the current spacecraft support.

Traditionally, the SLE routing was performed, within the mission control center, by the ESTRACK Control Center (ECC) which is in charge of operating the ESTRACK ground stations for the scheduled spacecraft passes as shown in the figure below. The ECC Operators monitor the automatic ground station setup and establish manually the SLE links from the SLE User to the active SLE Provider at the stations. In essence, the SLE User provides Spacecraft Controllers with a single point of contact to all ESTRACK stations, because the Mission Control System (MCS) is always connected to the same SLE User.

In recent years, the gradual collocation of SLE User and Mission Control System has been envisaged. Such a collocation has been foreseen both in terms of technical integration, which has been adopted in the systems evolution concept, but also in terms of operating responsibility. This new approach, illustrated in the figure below, translates the ESTRACK service boundary to the level of SLE interfaces and shifts the responsibility to operate the SLE User from the ESTRACK Control Centre to the Spacecraft Controllers. As a consequence, and in order to establish the SLE links to the active SLE Provider, the Spacecraft Controller would need to know which station communication chain has been configured as prime.

[1] Ground Operations Engineer, ESA/ESOC HSO/ONO, maite.arza 'at' esa.int
[2] Software Engineer, ESA/ESOC HSO/GI, holger.dreihahn 'at' esa.int

Today, during the station pre-pass configuration, the active SLE Provider is defaulted to the prime chain unless it is unavailable due to scheduled maintenance or unexpected errors / incidents in the station prime chain. A redefinition of the prime chain, during the pre-pass configuration, is un-critical for the telemetry services. These return services can be provided in parallel by prime and backup chains. For the telecommanding the selection of the communication chain is more critical, only the prime chain with the active SLE Provider is able to uplink telecommands to the spacecraft.

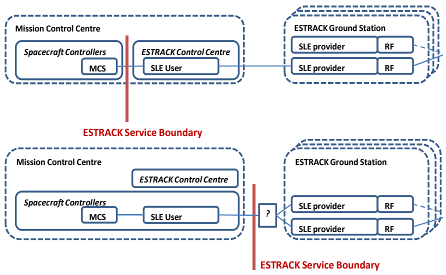

Therefore, the operator of the SLE User shall, via voice communication, confirm with the ECC Operator the prime chain for the pass, and use the set of SLE Service Instances (SI) corresponding to that particular communication chain in the ground station. Similarly, error-handling requires coordinated manual intervention between the two entities responsible for the operations of SLE User and SLE Provider to ensure SLE links are consistently re-routed.

As a consequence, this approach presents a limitation for the introduction of automated SLE User operations, i.e. unmanned operations, being in particular unsuitable for deep space missions operations characterized by large latency in the confirmation of telecommand via telemetry due to the large round trip light times involved[3].

II. SLE Routing – The Requirement

The BepiColombo mission is driving the requirement to provide users with SLE access to the ground stations without requiring knowledge of the active communication chain. This shall allow end-to-end consistency of the SLE routing from SLE User to SLE Provider without requiring voice coordination between the two entities responsible for the operations of SLE User (Spacecraft Controllers) and Provider (ECC Operator) and manual intervention. As a result automation is easier to implement if access to the ground station does not require awareness of the prime communication chain.

In order to address this mission requirement, the concept of SLE Routing has been conceived. The SLE Routing shall constitute a single entry point for SLE links established from users and route them to the appropriate SLE provider, i.e. the selected prime / backup communication chain at the station. The following high level requirements have been defined:

- REQ 1 **End to End Consistency** of the SLE routing, from the SLE User to the active SLE Provider at the ground station, and ground station configuration shall be ensured. The SLE routing concept shall prove robustness.
- REQ 2 **Redundancy** capability for SLE return services shall be maintained.
- REQ 3 **Interoperability** for cross-support shall be supported.
- REQ 4 **Backwards Compatibility** shall be maintained, to allow a seamless deployment and avoid impact on flying mission.

[3] In fact the BepiColombo mission is one of the drivers of the SLE Routing. However, this requirement is of general nature and will be adopted by other missions. Even flying missions consider adoption of SLE Routing, mainly to simplify automation.

REQ 5 **Integration** – The introduction and integration of SLE Routing shall impose a minimal impact on the infrastructure and the current operational station management

In addition, the SLE routing concept shall be feasible both from a schedule and cost efficiency view-points. Especially for a cost effective solution both the initial implementation and the maintenance costs are a key factor.

In order to avoid an over-complex implementation, the following assumptions have been verified with the users:

A 1 In line with the current operational approach, if the prime chain of a ground station shall be changed during a pass, it is acceptable that the SLE user has to re-establish the SLE communication links for forward services to the ground station.
A 2 For hot redundancy also the current approach is kept: The SLE user has to establish SLE links to each communication chain in the ground station, the prime being the active SLE connection and the second in hot redundancy.

III. Implementation Options and Trade-Off

In the following, we present a number of implementation options for SLE routing:

- SLE User IP Address List
- Routing on SLE Level
- Routing on TCP Level
- Intelligent DNS

Each option will be described and assessed according to the presented requirements.

SLE User IP Address List

ESA's SLE user systems allow the configuration of an IP address list for each each foreign (SLE) port. At time of an SÖE BIND the SLE user[4] tries to connect all IP addresses in parallel and sends the BIND invocation via the first established TCP/IP connection to the responding SLE provider.

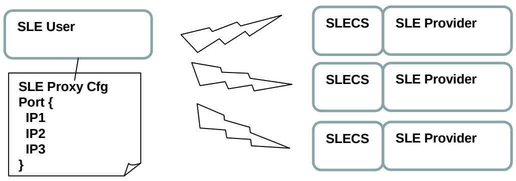

Evaluation

End to End Consistency	Robustness and end-to-end consistency can only be provided at expenses of operational complexity:SIs have to be loaded only on the active SLE provider to ensure the SLE User connections to the redundant Providers are rejected. SIs from the previous or next pass can interfere, because those SIs use the same TCP ports. The common TCP ports are required by the QoS configuration on TCP network level. For QoS purposes the different SLE delivery modes 'online timely', 'online complete' and 'offline' are mapped to certain agreed TCP ports.
Hot Redundancy	Not available as of today: SIs can only be loaded on one provider
Interoperability	OK, the approach is used by NASA / DSN as a provider to ESA missions
Backward Compatibility	The SLE IP Address List approach can coexist with the current approach. It requires an updated SLE User configuration and, for certain passes, an updated SI management.
Station Integration	OK but requires a major update of SI management as performed by the Station M&C

[4] Actually this feature is implemented in the ESA SLE API software package used by all of ESA's SLE applications. Note the SLE API is also used by a number of other agencies.

A common SLE Communication Server – Routing on SLE level

Another possibility to implement SLE routing is to use an SLE Communication Server (SLECS) like the one provided by the ESA SLE API in an extended scope. Instead of running on one machine, the SLECS would be used for all SLE provider systems in a station. The SLECS accepts the incoming TCP connections from the SLE user and routes the connection to the SLE provider, which has registered the Service Instance ID (SI ID) of the incoming BIND.

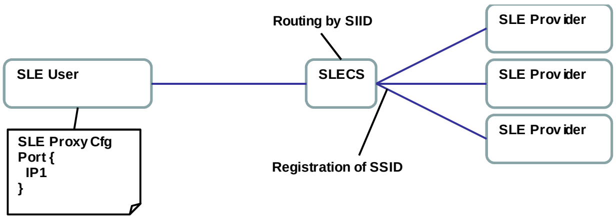

Evaluation

End to End Consistency	End to end consistency can be ensured by loading the prime SIs on the current prime SLE Provider and the backup SIs on the current Backup Provider. However, the set of SIs on the SLE provider systems cannot be the same.
Hot Redundancy	Parallel SLE associations to different SLE provider can be established, as long as the SI IDs are different, which is on line with the current approach. A theoretical limitation is the fact that SI IDs have to be unique among all SLE provider systems at the station. A *severe limitation* is the single point of failure of the SLECS, which is introduced by the concept.
Interoperability	Very good, routing is based on the SLE SI ID.
Backward Compatibility	Yes, fully compatible with current approach.
Station Integration	Would potentially require machines per station to host the SLECS.

Connection Routing on TCP level

A further option to make access to prime communication chain transparent would be the routing of connections on a TCP level: An incoming connection is always accepted at the same IP address on defined TCP port and depending on the current setting it is forwarded to another IP address and TCP port. This can be realized with a commercial network router or with a dedicated application.

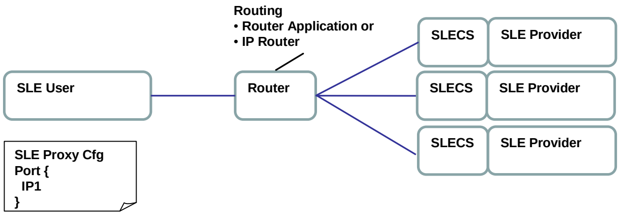

Evaluation

End to End Consistency	Can be ensured very well – routing settings can be precisely adjusted and monitored. Prime and backup SIs can be loaded in parallel on each SLE provider. In addition the same set of SIs can be loaded on each SLE Provider.
Hot Redundancy	This concept offers several possibilities, In addition to the current approach to have parallel SLE associations from the SLE User to several SLE providers, even for a single SLE association a hot redundancy can be provided by mirroring one SLE association to a second (third) SLE provider. This might be attractive as soon as the available bandwidth to the station does not allow parallel connections.
Interoperability	The approach supports interoperability very well. Routing is based in SI IDs and only one IP address is required per ground station.
Backward Compatibility	The approach is fully compatible with current approach.
Station Integration	The router has to be integrated into the ground station M&C system, which uses a private M&C protocol. For a commercial router this probably means the development of a proxy agent. This is rated as quite some effort in terms of complexity and costs.

Intelligent DNS

An appealing approach is the use a dynamic DNS lookup in the current context. Instead of configuring the IP addresses in dotted notation at the SLE user, one would configure the hostnames instead. The idea is now that for an SLE bind the lookup of the hostname is done via DNS, which resolves the hostname to the IP address of the SLE provider which shall be used. The DNS resolution is dynamically updated in order to resolve a name to the intended address. There are several options where this DNS server can be located, but the name lookup must be possibly adjusted on a per pass basis (or even in the middle of a pass to switch to another SLE provider). So in any case the DNS server must be in a way under the control of the various station M&C systems.

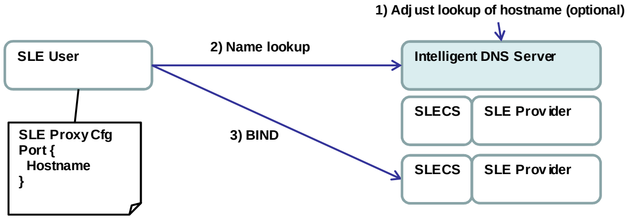

Evaluation

End to End Consistency	End to end consistency can be ensured, if the DNS settings are updated with immediate effect from the Station M&C system. Prime and backup SIs can be loaded in parallel on each SLE provider. In addition the same set of SIs can be loaded on each SLE Provider.
Hot Redundancy	Parallel SLE associations to different SLE provider can be established.
Interoperability	Requires external users to use a DNS server of the SLE provider network. This approach is in use with external partners for other services.
Backward Compatibility	The approach is fully compatible with current approach
Station Integration	Assuming a central DNS at ESOC, the station integration requires that each Station M&C system can update the central DNS server to reflect the selection of the prime chain in the station.

IV. Detailed Evaluation of the Current Candidate Solution

Based on the assessment of the SLE routing high level requirements, the Intelligent DNS is the preferred candidate. It offers a good compromise between the provided features, implementation costs and effort. Furthermore several DNS implementations are available as commercial and open source iplementations. The Intelligent DNS minimizes the station integration effort, since it only impacts the station M&C system. Also the current way to operate SLE provider in terms of installing and archiving SIs can remain as is. However, the use of a new technology (DNS) implies some risks (failed DNS lookup, long timeouts, erroneous responses) and adds a new sort of task with some complexity, namely the dynamic modification of DNS lookup to guarantee the required end-to-end configuration consistency.

To address the concerns stated above, the following requirements are defined:

- DREQ_01 The DNS server shall respond to queries with minimal delay (<100ms), especially if addresses cannot be resolved.
- DREQ_02 The DNS server shall be controlled by the ESTRACK operators, i.e. the ESTRACK Operator shall be in position to change the name resolution 'on the fly' during operations to e.g. switch the prime communication chain. Similarly, in order to ensure operability, a simple configuration mechanism shall be provided.
- DREQ_03 The DNS reconfiguration mechanism to change the prime communication chain shall be taken into account immediately
- DREQ_04 The DNS server shall support a centralized architecture.
- DREQ_05 The SLE user shall limit caching of DNS query results.

As DNS supports from a theoretical point of view all of the above requirements, implementation options for DNS server have been analyzed.

1. Options based on the Linux operating system, which could be operated and maintained as an 'application' like other applications of the ground station domain. On Linux a lightweight DNS forwarder 'dnsmasq' is available and seems to fit the needs from a technical point of view.
2. A solution called 'Intelligent DNS' based on a commercial DNS server, which is in use at ESA/ESOC. Namely this solution is in use at ESOC for the Navigations Office and the Envisat mission to allow routing of external parties to prime and backup data distribution servers at ESOC.

In particular, the latter option is considered suitable, as it has a proven track record for internal and external users, it is scalable and has been developed to support services with requirements very similar to those of SLE routing. In a way it is an extended use of a DNS solution already in use at ESOC, last but not least this options promotes a homogeneous DNS implementation.

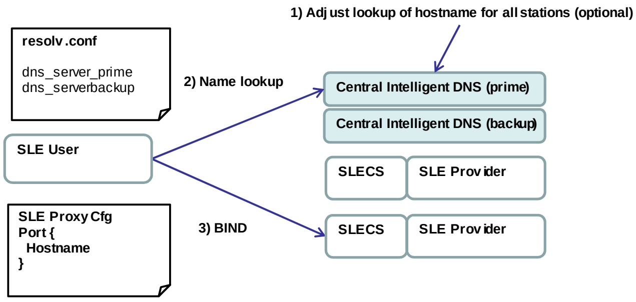

The figure above shows the basic architecture, which can easily be extended for external users. At any point in time the ESTRACK operators can adjust the name resolution for the prime and backup chain. When an SLE connection is established, the hostname is looked up via the 'Intelligent DNS' and the SLE connections are established to the correct prime and backup chains.

The end to end consistency requirement is fulfilled by extending the ground station configuration, via the station M&C system, to the “Intelligent DNS”. For the routine update of the name resolution the ‘Intelligent DNS’ offers the possibility to probe (poll) TCP ports to detect if a ‘service’ is available or not. The configured name is resolved according to ‘availability of services’ (TCP ports in state LISTEN). For the use in ESTRACK, i.e. the capability to update the ‘Intelligent DNS’ name resolution from the Station M&C system, this seems to be straight forward: To activate the prime / backup chain, the Station M&C system can enable / disable (dummy) services e.g. on the SLE provider, which listen on a particular TCP port. The ‘Intelligent DNS’ detects the availability of services by probing (interval e.g. 2s) and adjusts the name lookup accordingly. The activation / deactivation of (dummy) services by Station M&C can be done via shell scripts and is as such a matter of configuration.

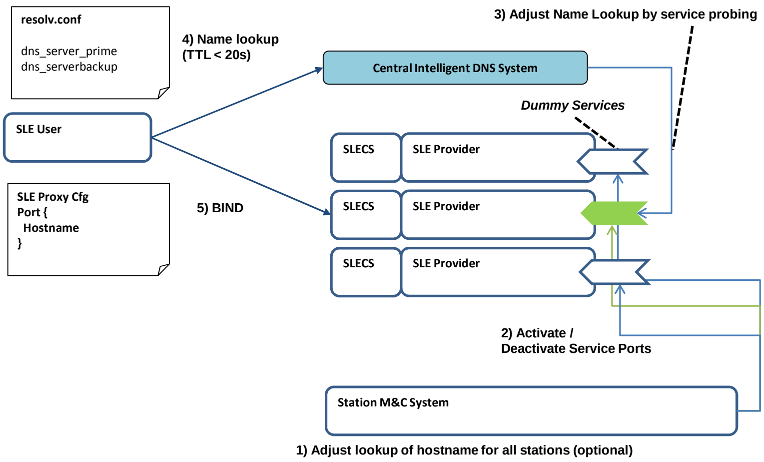

An important aspect in the assessment of DNS routing, is the name caching, which potentially affects (being problematic) external clients. The ‘Intelligent DNS’ can be configured to answer name resolution queries with a particular the time to live (TTL). This can be set to relatively small values, but one should avoid that a resolved name reaches the client with an expired TTL[5]. In the context of SLE routing the TTL imposes a ‘forbidden period’, in a sense that an SLE re-connection attempt must wait TTL seconds to ensure correct name resolution[6]. An initial discussion suggested that a TTL of 2 seconds is feasible; for SLE routing this translates to the fact that when a communication chain is switched from backup to prime, the SLE user has to wait the ‘forbidden period’ (TTL, e.g. 2s) before the new connection can be established. This wait period ensures that a new name lookup is triggered as the TTL has expired.

Logging of DNS setting

In order to allow analysis of the used DNS settings for a pass, it is recommended to record the actual name resolution. This can be done by adding a log to the involved systems like the SLE user or the station computer or the SLE user, which logs the name resolution to IP addresses. Even a cron job can be setup to record name (changes) resolution in defined intervals. This idea is to keep records for troubleshooting and service level agreement documentation.

Dynamic DNS and Offline Data Retrieval

It shall be noted that although SLE routing can be used also for offline retrieval, it is expected that this will not be necessarily the case. In contrast to online operations, which are planned and scheduled activities, offline operations are typically unscheduled operations. In combination with the anticipated requirement to access *any* of the

[5] A TTL of zero, which means ‘don’t cache’ for name resolution replies seems to be problematic for some DNS implementations.
[6] The need to wait should be only present if the communication chain has to be switched due to a failure during a pass. Nominal changes of the configuration chain are done outside of passes and leave therefore enough time for name change propagation.

independent and redundant telemetry storages at *any* time, there is the option to use dedicated IP addresses or names as of today for offline retrieval. As an alternative the DNS based approach could be used also here, but it is then left to the SLE user (mission) to check the current name resolution and to change it e.g. via voice loop if needed.

Dynamic DNS and Multiple Spacecrafts per Aperture – MSPA

The presented Dynamic DNS lookup solution has been analyzed in the light of MSPA. For MSPA two or more spacecraft, which shall both be covered by the antenna half power beamwidth, use the same ground station antenna in parallel to communicate with a spacecraft. For the current ESTRACK ground receiver (IFMS) this requires a dedicated IFMS per spacecraft. For the uplink, time multiplexing is currently assumed, i.e. only one mission is allowed to actively uplink from the ground station.

Basically two options exist to extend the SLE routing concept:

1. Both SLE Users connect to the same SLE Provider and commanding is only possible for one spacecraft per pass. This is ensured by ESA's SLE provider implementation, which allows only one active forward session, but up to four concurrent independent return sessions. Each return session can connect to a different IFMS.
2. Both ground receiver (IFMS) and SLE Provider are dedicated to each supported spacecraft. Each SLE User connects to a different SLE provider. The time sharing for commanding has to be coordinated either by agreed timeline and / or via voice loop since this option implies 'swapping' of the commanding chain.

V. Conclusion

The concept of SLE routing has been presented as a response to the mission requirement to have transparent SLE access to ground stations. SLE routing eliminates the need of the SLE user / Spacecraft Controller to know the active configuration (prime / backup chains selection) at the ground stations and simplifies automation.

A number of implementation options have been presented and assessed in light of the main requirements; currently the approach of an intelligent DNS lookup is considered to offer the best compromise among functionality, robustness and cost effectiveness. Intelligent DNS lookup combines efficiently established internet standards like DNS with CCSDS compliant SLE implementations as well as systems dedicated to ESTRACK operations. It is expected that the combination of proven standards with reliable implementations provide advantages over dedicated implementations. Currently the first tests are being conducted with missions in order to gain experience and to prove the concept.

Acronym List

DNS	Domain Name System
SLE	Space Link Extension Services
SLECS	SLE Communication Server
SI	Service Instance
SI ID	Service Instance Identifier
M&C	Monitoring and Control

References

[1] CCSDS 912.1-B-3, FORWARD CLTU SERVICE SPECIFICATION, Blue Book, July 2010
CCSDS 912.3-B-2, FORWARD SPACE PACKET SERVICE SPECIFICATION, Blue Book, July 2010
CCSDS 911.1-B-3, RETURN ALL FRAMES SERVICE SPECIFICATION, Blue Book, January 2010
CCSDS 911.2-B2, RETURN CHANNEL FRAMES SERVICE SPECIFICATION, Blue Book, January 2010
CCSDS 911.5-B-2, RETURN OPERATIONAL CONTROL FIELDS SERVICE SPECIFICATION, Blue Book, January 2010

[3] Y. Doat, M. di Giulio, G.P. Calzolari, "ESA Telemetry and Telecommand System (TMTCS)", SpaceOps 2004

[4] T. Ulriksen, G. Scotti, C. Mueller, Dr. C. R. Haddow, "Network Interface System ESA's Next Generation CCSDS SLE Interface System", SpaceOps 2008

[5] T. Beck, N. Bobrinsky, A. Pena, J. Reynolds, "Remote Operations of ESA's Ground Station Network", SpaceOps 2002