

II. The IGS WAN Network

The main task of the IGS WAN services is to transfer between the IGS sites, different application flows each of them requiring an appropriate Quality of Service (QoS).

The IGS WAN services include the following components:

- **Wide Area Network (WAN)**, built around MPLS and ISDN services. ISDN is used for out-of-band management of remote equipment and as a second backup for ATV-CC TM/TC traffic
- **Local Loop**: to interconnect Telecommunications Provider (TP) or Service Provider (SP) premises with user premises
- **Network Termination or Customer Edge (CE)**: fixes the SP reference point inside the user premises (IGS site). The customer access ports of the CEs represent the line of demarcation for the IGS WAN end-to-end services as provided by the SP.

The following figure depicts all international partner sites and USOCs that are part of the IGS-WAN. The main - mission critical- sites are designed with redundant local loops and CE equipment. The dimensioning of the local loops is based on the traffic requirements of the specific site while the redundancy requirements call for physically diverse routing with two physically separated and independent local loops (communication lines) terminated in two different buildings using dedicated CE equipment.




Figure 2. IGS WAN Network Plan

It was chosen to have one global service provider (T-Systems) responsible for the end-to-end network connectivity and service availability. The service provider uses beside his own telecommunication platform also communication lines belonging to other TPs, nevertheless he is responsible for the fulfillment of the Key Performance Indicators (KPIs) of the global IGS-WAN Service Level Agreement (SLA).

The SLA foresees a 99,93% availability of the communication services for a redundant site. This allows for a maximum of 31 minutes service outage in a timeframe of one month (30 days). This value sounds fairly reasonable taking into account site access and equipment exchange times and even better when considering how long does it take to repair accidentally cut fiber cables. On the other hand in the “space business” there are mission critical phases (e.g.: docking maneuvers, first acquisition etc.) where a communication outage of a few minutes can have

quite dramatic consequences. In order to be prepared for such phases extra measures were made in the IGS-WAN achieving even higher availability and service restoration times than the ones specified in the SLA. These measures include parallel routing of specific data streams (i.e.: TM/TC), network monitoring tools with regular polling (1 min), SNMP traps combined with audio/visual notification for the ground controllers.

In spite all these measures the IGS network engineers together with the Service Provider had to experience the existence of network and communication service outages that can not be seen/localized by network monitoring tools and can completely stop the communication of a given (even redundant) site.

The article addresses this problem: complete and “silent” network outage - also known as: “communication black hole”.

III. Communication Black Holes

A. Description and consequences

Communication black holes occur mostly in the Service Providers MPLS backbone network. They manifest in a complete communication outage between two (or more) sites without any recognizable reason: the communication lines (local loops) are all up and running, CE equipment is also all up and running (all units are reachable via outband management/i.e. using ISDN lines).

This is a quite undesirable situation: all the redundant network build-up is completely useless and in regular case the first level support cannot help either: all telecommunication equipment and communication lines that they see, monitor and manage are in principle working; “just” the end-to-end communication is not possible. A trouble ticket needs to be opened and escalated to level 2 support for investigation. Problem escalation, investigation and manual intervention is always time-consuming – such a communication problem is not likely to be solved in a few minutes. Under certain circumstances (i.e. mission critical operation) such an event can have grave consequences.

In order to avoid similar events in-depth investigation, problem detection, mitigation and recovery procedures are needed.

B. Technical background

Today’s operational backbones and high speed IP networks are layered: overlaid often on optical networks. An IP link is often implemented as a path through a set of optical components, some of which are shared across multiple IP links. Similarly, there may be an intermediate MPLS layer, wherein IP packets are transported via Label-Switched Paths (LSPs), which in turn are established using IP routing protocols (such as OSPF). Moreover, multiple Virtual Private Networks (VPNs) may be overlaid on top of the MPLS topologies. Given the complex, cross-layer interactions between the different control planes, complex failure modes and fault scenarios arise.

For example, in cases when LSPs are established following the shortest-path routing using OSPF, one failure scenario that has been observed in practice is when OSPF re-routes due to a problem, but MPLS does not follow it.

MPLS is based on dividing the forwarding space into Forwarding Equivalence Classes (FECs) and establishing corresponding Label Switched Paths (LSPs) from sources to destinations in the network. For each FEC, labels are exchanged between adjacent nodes in the network and used to forward packets along the LSPs. A protocol called the Label Distribution Protocol (LDP) has been explicitly defined for distributing labels. When LDP fails while all lower-level protocols and physical connectivity are operational, IP Interior Gateway Routing Protocols (IGPs) may remain unaware of the failure in LDP. In this situation the IGP continues to make routing decisions regardless of whether an LDP session is established or has failed, and whether or not the LDP labels are correct.

In other failure instances, MPLS control plane is working properly (hence no alarm), yet there is corruption in the forwarding plane due to poor implementations and/or configuration errors.

Such an event (black hole) can be silent in nature, with no router alarm indicating that end-to-end communication is broken. Due to the ever-increasing complexity of backbone networks, it is highly unlikely they will ever disappear entirely. While such silent failures are rare, they can have a large impact; in many cases, a complete loss in connectivity. These failures are extremely time-consuming to localize (order of hours to days) because there are no alerts/alarms to guide operators to the location of the failure. Hence, from an operational standpoint, it is extremely important to design a mechanism that can quickly detect such failures, and, moreover when possible re-establish communication.

IV. The IGS WAN experience

A. Network operation

The IGS Network is based on the network transport services of the Service Provider's (T-Systems) MPLS backbone. In order to establish MPLS black-holes detection mechanisms and avoidance procedures for the IGS a close collaboration was needed between the two engineering teams (Service Provider and IGS).

The IGS Network is built using a star topology, all sites connecting to the Central Node: COL-CC. COL-CC is a redundant site with 2 CE routers located on-site and 2 MPLS network access lines (local loops). Additional to the CE routers in COL-CC there are two SAA-routers installed (SAA= Service Assurance Agent), one in the Prime and one in the Backup facility. The SAA routers monitor and validate the IP connectivity to all remote sites (performing availability and QoS measurements).

For detecting/avoiding MPLS black holes in the IGS network an SAA-like monitoring and tracking procedure was implemented but this time on the remote CE routers checking the connectivity to the central node. The communication black-hole tracking procedure was chosen to be implemented only at the redundant/mission critical sites.

The redundant sites, similar to COL-CC, have 2 CE routers located on-site, terminating 2 physically separated and independent MPLS network access lines. The two access lines connect on their other end to two Provider Edge (PE) routers. The PE routers represent the MPLS point of presence (access point) into the MPLS backbone network of the service provider. The PE routers (PoPs) of the redundant sites are usually located in geographically distant areas. For example in case of the site MCC-H (Mission Control Centre - Houston) one PE router is located in Miami (Florida, USA), while the second in Dallas (Texas, USA). Border Gateway routing Protocol (BGP) is running between the CE and PE routers: based on predefined routing metrics the traffic of a redundant site is routed primarily over one of the local loops (the one set with the higher BGP routing metric). Shall this local loop fail the second local loop will automatically take over (a reroute will occur). A local loop traffic reroute - using geographically distant MPLS PoPs - will most likely result in completely different and new path through the MPLS backbone. A straightforward idea is that such a reroute could be beneficial in case of MPLS Backbone communication black-holes: unfortunately the BGP routing protocol running between the PE and CE routers will never receive automatically a notification that a "black-hole" situation occurred. Therefore the implemented solution tracks such events and "notifies" the BGP routing protocol to perform a reroute.

B. Tracking/recovery concept

The principles used in the IGS for discovering and recovering from an MPLS black-hole situation are:

- 1) Monitor and track end-to-end connectivity between two sites (due to the star topology between remote site and COL-CC)
- 2) In case of a black-hole event, where no communication between the two sites is possible, trigger a local loop reroute event in order to run over a new (and most likely different) MPLS path

C. Solution implementation

- 1) Prerequisites

The CISCO CE routers performing the detection/avoidance procedure shall have at least an IOS Release-Version 12.4 (22)T installed. It is also important to assure that the source addresses of the SAA-Probes are not routed over the crosslink between the CE routers at the remote sites.

- 2) SAA-Probes

On the Prime CE-Router two static ICMP-Echo measurements/probes are implemented. These probes check the availability of the loopback address of the destination CE router (here Prime and Backup CE in COL-CC). Every 5 seconds an ICMP-request is being sent to the destination address using the own loopback address as the source address. The measurements can be run in a specific VPN/VRF and inside a specific traffic class.

Probe properties:

- Every 5 seconds
- Datagram size: 100 Byte
- Option: checksum can be controlled
- Option: define Class of Service
- Option: define VRF

3) Tracking

The SAA-Probes (for Prime and Backup connectivity) are being monitored using a “Tracking” function and they are set in a logical “OR” condition to each other. The intention here is to prevent a status change if only one of the remote connectivity fails. A complete outage shall be tracked, where both prime and backup connectivity is out of order. In the probe tracking function an additional delay-condition is included to avoid network and status flapping situations. The status of the aggregate tracking is UP or DOWN

Tracking logic:

- Tracking status Prime OR Tracking status Backup = Aggregate tracking status
- Delay DOWN = 15 Seconds
- Delay UP = 90 Seconds

4) Event Manager

The “Event-Manager” (from an IOS-Version 12.4 (22)T or higher) is able to interpret the status of a “Tracking” as an “event”. In our case the UP and DOWN event of the aggregate tracking status is triggering a series of actions which result in the modification of the configuration of the routing protocols on the Prime CE router. The following steps are followed during the modification:

- Change of the Route-map relationship for outgoing routing updates (for the relevant VPNs)
- For status UP: the route-map with the community 20570:100 will be used. The metric of the routing protocol towards the on-site IP routing components will be changed (Example: OSPF metric will be set to 100)
- For status DOWN: the route-map with the community 20570:80 will be used. The metric of the routing protocol towards the on-site IP routing components will be changed (Example: OSPF metric will be set to 170)
- The BGP process (after modifying the route-map relationship configuration) will be “soft” reset in order to notify the neighbors about the latest routing changes

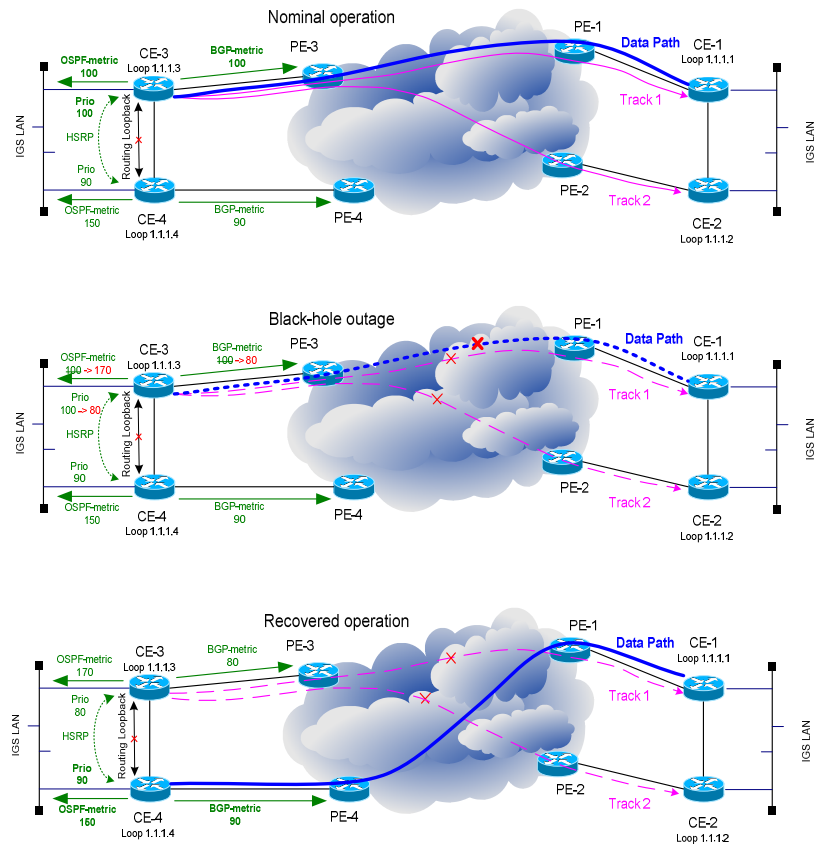


Figure 3. Black hole tracking and recovery

5) HSRP

When HSRP is being used at a remote redundant site the HSRP tracking status of the WAN interfaces can be changed using the consolidated SAA-Probe tracking status.

V. Summary / Conclusion

The presented solution was implemented after 1 year of IGS WAN over MPLS operation. In the one year without the tracking and recovery solution two “black-hole” events were experienced – luckily they took no longer than approximately 15 minutes each as the Service Provider noticed relatively quickly the issue and each time they could correlate the event to ongoing configuration changes/maintenance activities in their core network and managed to reroute the affected traffic. Nevertheless the outages called for the implementation of the presented procedures.

Following the validation tests and implementation we have now 2 years of operation without such an event – we experienced several times the tracking and event manager performing a reroute which resulted in a short communication outage (in range of seconds) of the affected site but no operational impact.

We can conclude that the implemented tracking and recovery concept successfully chases down communication black-hole events between a remote site and COL-CC in the IGS WAN network.

The presented solution and ideas can be definitely implemented in other ground segment MPLS-based communication networks.