

BIT BV  
Galileilaan 19 - 6716 BP Ede  
Postbus 536 - 6710 BM Ede  
T: +31 318 648688  
F: +31 318 643334

**bit**  
internet technology

E: [info@bit.nl](mailto:info@bit.nl) - <http://www.bit.nl>  
KvK-nr : 09090351

## Incident report, network error March 7th 2016

### Summary

A high CPU load on one of the core switches caused a loop in the core network on Monday March 7th 2016. As a result, the core routers had to alter their routing and forwarding tables, but they could not handle the load that was put on the processors. This meant that the BIT network was not accessible to the internet and the internet was not accessible from the BIT network.

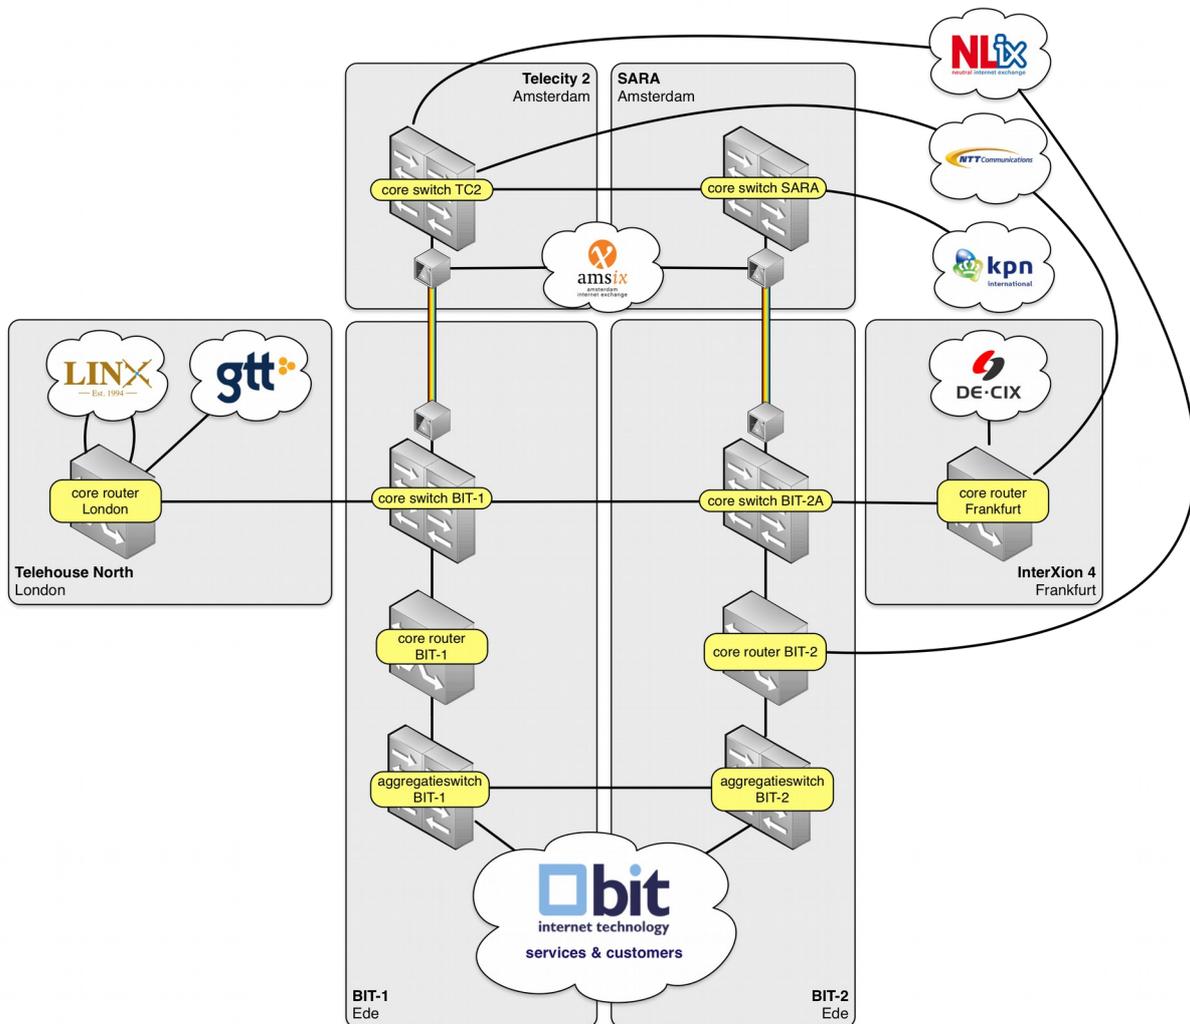
### Details

- 14.23 Several ports on the core switch in BIT-2A become unusable. At the same time there are reports coming in from the monitoring systems that the core ring has been interrupted. The last logs and counters of the switch cannot be sent to the central log services due to the high CPU load. A loop occurs in the core switch, making all connections to the switch unreachable. As a result of the interruption of several connections that enter BIT-2A over the core switch, a high CPU load is put on the core routers in BIT-2A and BIT-1. This high load means that the routing and forwarding tables of both routers get out of sync, which leads to blackholing of all traffic. This also means that BIT's regular phone number is also not reachable.
- 14.24 A BIT engineer is present in the data center to look at the core switch at BIT-2A. The switch is unreachable through serial console.
- 14.28 The switch is rebooted. This action results in the loss of local logging and counters of the switch. Rebooting the switch does not have the desired result. The switch is then turned off to stabilise the rest of the network.
- 14.30 – 15.00 There are several routing adaptations in an attempt to lower the CPU load of both routers and hence to sync the routing and forwarding tables again.
- 14.46 As a result of the network problems, placing an incident report on [www.bit.nl](http://www.bit.nl) and [www.bit.org](http://www.bit.org) is problematic. That is why <https://twitter.com/bitnl> and <https://facebook.com/bitnl> are used to send out an incident report.
- 15.02 The incident is reported on [www.bit.nl](http://www.bit.nl).
- 15.05 Due to the continuing problems, all connections with the core switch in BIT-2A are disconnected and the switch is turned on again. All traffic to the core router in BIT-2A is rerouted to the core router in BIT-1. This means that the router in BIT-2A is not handling any traffic anymore. The network is stabilising.
- 15.06 – 16.00 The core switch in BIT-2A is checked on hardware defects. No defects are detected. Several tests are run on the switch.
- 15.15 The incident is reported on [www.bit.org](http://www.bit.org) since [www.bit.nl](http://www.bit.nl) is not widely accessible. Further updates on the incident are posted on [www.bit.org](http://www.bit.org).
- 15.21 – 17.00 The core router in BIT-2A is fed with routing and forwarding tables of BIT-1's core router. The core router in BIT-2A is not yet put to use for traffic going through the core switch in BIT-2A.
- 15.24 BIT's regular phone number is back in use.
- 15.34 IP routing is stable, the majority of the services is hereby restored. Services using VLAN's between BIT-2A and Amsterdam are not yet functioning.
- 16.00 The ring between Ede and Amsterdam is put back online without allowing services on it yet. Several tests are successfully executed on the ring.
- 16.10 – 16.30 Several ports on the core switch BIT-2A are online again, under supervision, making all services available again. Core switch redundancy is hereby restored.
- 16.30 All services are restored.
- 16.48 – 16.55 The router in BIT-2A is put back into use. Redundancy on core routing is hereby fully restored.

### Conclusion

For a correct image of the incident, it is important to have a clear idea of BIT's network architecture. The diagram below depicts the most important components in BIT's core network. The components in BIT-1 and BIT-2 are simultaneously active, failure of one of the components automatically means that all traffic

that should go through that component, is rerouted. This goes for both the switching and the routing. There is a core ring with loop protection between the data centers in Ede and Amsterdam. Failure of one of the core switches or one of the paths will lead to the use of an alternative path.



High CPU load on the core switch in BIT-2A leads to a loop on the core network. As a result of that high CPU load, loop protection on the switch is not functioning as it is supposed to. In turn that results in all traffic on that switch being sent to all ports and the disabling of the connections on the switch. The paths between BIT-2A and Frankfurt and SARA (Amsterdam) will not be available if this happens. This only has direct consequences for a number of VLAN's between BIT-2A and Amsterdam and for a part of the access services provided over the Tele2 network.

Because of the changes in the core network, the core routers have to alter a great number of routes. The routers' CPUs appear to be a bottleneck here, leading to discrepancy between the routing and forwarding tables of the routers for a considerably long time. This means that the traffic is not arriving at its destinations, but that it gets lost and that the rest of the internet is unable to reach the BIT network anymore. Without the CPU problems of the routers, the impact would have been contained to a possible short interruption of existing connections.

A similar problem occurred during the network incident on February 25th, 2016 (<https://www.bit.nl/news/...>). This incident was immediately reported to both the supplier of the router and the vendor. These parties are looking for a solution, but they have not found one yet. The routers that are having problems, have not even been online for a year and should have more CPU capacity at their disposal than the old routers.

As a result of the strict access procedures, it took a long time for the incident to be reported on [www.bit.org](http://www.bit.org). This host is situated outside the BIT network and data centers to communicate incident reports when [www.bit.nl](http://www.bit.nl) has limited or no availability.

### **Adjustments**

- The access procedures to [www.bit.org](http://www.bit.org) have been altered to allow faster access to be able to report incidents quicker.

### **Adjustments - Update 11-03-2016**

- There has been a lot of contact between our distributor and the fabricator of the routers to place routers with much more CPU capacity on very short notice. An incident analysis has been conducted with the fabricator's specialists to determine the exact cause of the CPU load.
- Partly due to the fabricator's recommendations, several adjustments have been made in the network configurations to lower the load on the core routers without compromising quality and availability of the network. Even though these adjustments lower the load on the routers, it is not a permanent solution. We expect more of these adjustments in the coming weeks.

### **Continuation**

- Research is being done on possibilities to keep BIT's regular phone number reachable in case of (big) network incidents.
- Further research is being done on the cause of the high CPU load on the core switch that triggered the incident.
- Further research is being done on improving the loop protection of the core network.

### **Continuation - Update 11-03-2016**

- Based on the incident analysis, the fabricator suggested replacing routers with more and more powerful CPUs. The fabricator's simulations showed that this new type of router performs much better. However, these routers are not immediately available. The fabricator expects to be able to supply temporary routers of the same type within two weeks.
- In expectance of these routers, BIT's network engineers are preparing configurations.
- As soon as the routers can be delivered, BIT's network engineers will work on testing and configuring them with the highest priority. They expect to need two workdays to do this. Emergency maintenance will be announced to replace the two core routers one by one. A rollback scenario will be developed to switch back to the old routers in case of unforeseen problems.

- The delivery of the final routers is expected to be eight weeks. The replacement of the temporary routers with the new, final routers will also be announced with a maintenance notification.

## **FAQ – Update 11-03-2016**

*BIT's network is set up redundantly. Why did problems occur anyway?*

BIT's network consists of two core routers, geographically separated in BIT-1 and BIT-2. These core routers exchange all routes to the internet they are being taught by transit suppliers and on internet exchanges, so the other router knows these routes are available to be used. These two routers together are responsible for routing the traffic towards client networks: if one of the two routers fails, the other router can immediately – and unnoticeably – take over. As described in an earlier paragraph, the problems started on one of BIT's core routers. The failure of the core switch meant that the core router in BIT-2 lost its connection to the router in Frankfurt (where BIT is connected to the DE-CIX and a transit supplier) and another transit connection. A large number of routes were rendered unusable. All these routes, however, were also communicated to the core router in BIT-1. Both routers could not process the large number of changes in the routing tables quickly enough.

*Why has BIT chosen for this type of routers that seems to cause problems?*

While purchasing these routers, over a year and a half ago, BIT has had extensive contact with specialists of the distributor on which type of router would be most suitable for BIT's aims. Statistics and network designs were also shared with the distributor. Based on the available information on the intended model of routers, it was concluded that they fit the demands, needs and growth goals of BIT. The fabricator has not given any signals that this type of router was not suitable. Based on the performance statistics supplied by the fabricator and information from other ISPs, we concluded that the intended new models will be able to deliver the necessary performance.

## **Contact**

For any questions relating to this RFO, you can contact our customer care department on +31 (0)318-648 688 or [support@bit.nl](mailto:support@bit.nl)