

TNC2010 Extended Abstract:

ANOMALY DETECTION IN BACKBONE NETWORKS: BUILDING A SECURITY SERVICE UPON AN INNOVATIVE TOOL

Author and Author Affiliation

Maurizio Molina

DANTE, City House, 126-130 Hills Rd, Cambridge, CB2 1PQ, United Kingdom
Email: maurizio.molina@dante.net

Wayne Routly

DANTE, City House, 126-130 Hills Rd, Cambridge, CB2 1PQ, United Kingdom
Email: wayne.routly@dante.net

Ignasi Paredes-Oliva

Universitat Politècnica de Catalunya (UPC), Computer Architecture Dept., Jordi Girona 1-3
(Campus Nord D6), Barcelona 08034, Spain
Email: iparedes@ac.upc.edu

Ashish Jain

Guavus, Inc., Mission Tower One, 3975 Freedom Circle, Suite 100, Santa Clara, CA 95054,
USA
Email: ashish.jain@guavus.com

Keywords

Network Anomaly Detection; Anomaly Validation; Security; GÉANT.

Abstract

DANTE first started an initiative to detect security-related anomalies in the GÉANT backbone network several years ago, leveraging the sampled NetFlow v5 data that was already being collected from the GÉANT Juniper core routers rather than deploying expensive, dedicated probes on every link.

Currently, there are 18 T or M series routers spread across Europe. NetFlow is collected on every interface with an external peering network, in the incoming direction, which, as GÉANT is mainly a transit network, is sufficient to account all the traffic. NetFlow data is exported by the routers to a single fanout box, which then duplicates the flows to multiple NetFlow collection and analysis tools. The NetFlow sampling rate was initially 1/1,000 and has recently been increased to 1/100.

The implementation of some proof of concept plugins to the open source tool NfSen [1] successfully detected scanning and Distributed Denial of Service (DDos) events, as reported in [2], and encouraged us to look at commercial tools for setting up an anomaly alert service for the benefit of customers of the GÉANT network. DANTE's goal was to establish a service to report relevant security anomalies to the Computer Emergency Response Teams (CERTs) of the National Research and Education Networks (NRENs) connected to GÉANT.

The complexity of the GÉANT network – its hybrid nature, the type and volume of traffic, and its global connectivity – presented a challenge to tool selection. As a minimum, the commercial tools had to detect the same types of anomalies as those identified by the proof of concept plugins, plus those specified by the partner NREN CERTs in requirements definition discussions.

We started a trial with three distinct commercial tools – PeakFlow SP [3], Stealthwatch [4] and NetReflex [5]) – which were selected for testing because of the different mechanisms they used to detect network anomalies.

For 13 selected days during autumn 2008, we meticulously checked the security anomalies highlighted by each tool (more than one thousand altogether), using the capabilities of NfSen [1] to analyse at a raw NetFlow level, and also some manual techniques. This resulted in a true/false positives ratio for each tool, as well as a valuable count of anomalies that were unseen by the other two tools. We showed a subset of the anomalies to some of the NREN CERTS, to obtain independent verification of the anomalies and confirmation of our true/false positives ratio. Further details of the trial can be found in [6].

Based on the results derived from analysing +1000 anomalies, NetReflex was the tool chosen. It showed a more balanced capability to detect both (D)DoS and scanning events, can detect low-volume events as well as large-volume ones, and provided a more even “spatial” distribution of the origin of malicious events, irrespective of the peering “type”.

The deployment of NetReflex in the production environment started in summer 2009. The main technical problem we faced was due to increasing the sampling rate from 1/1,000 (used during the trials) to 1/100 (the rate we want to use for the production service): the Juniper router hardware needed to support the increased sampling rate reported duration incorrectly. (This anomaly is described in [7].) However, the NetReflex vendor Guavus successfully implemented a workaround.

NetReflex supports the detection of the following types of security anomalies: Denial of Service (DoS) and Distributed Denial of Service (DDoS), both with sub-types UDP Flood and TCP SYN Flood; Port Scan and Network Scan.

Before making operational use of the anomalies signalled by NetReflex, we ran (in November '09) a four-week trial phase where, twice per week, we analysed every signalled anomaly for the previous day and validated its authenticity via several manual checks; as during tool selection, these mainly involved access to raw NetFlow records and “whois” and/or DNS queries on the IP addresses involved in the anomalies. These checks, which varied depending on the type of anomaly being detected, resulted in a number of best practice guidelines and ensured that the false positive ratio was in the order of a few per cent, which is operationally acceptable. (The full paper will include snippets of flow records for each type of anomaly.)

The main drawback of manual validation, besides being a resource-intensive task, is that it relies on the experience, knowledge and skills of the operator to extract the evidence of an attack from NetFlow data and reach a correct conclusion. Providing an automatic way to retrieve all the flows potentially involved in an anomaly and present them back to the operator – automatic anomaly validation – removes this dependence and allows the operator to take a final decision more quickly and reliably. We are currently exploring and developing a method proposed by Brauckhoff et al. [8] for automatically extracting the flows related to a particular anomaly. Their approach is composed of three stages: anomaly detection, flow filtering and flow mining. We have added to the original algorithm the capability to automatically extend the analysis to time periods around those flagged by NetReflex as anomalous, to look for suspicious activity related to the signalled anomaly and thus gain a more complete understanding of it. Preliminary results (see figure below) show that in several cases the time periods where it's relevant to look for suspicious activity are much larger than those on which the tool initially focuses.

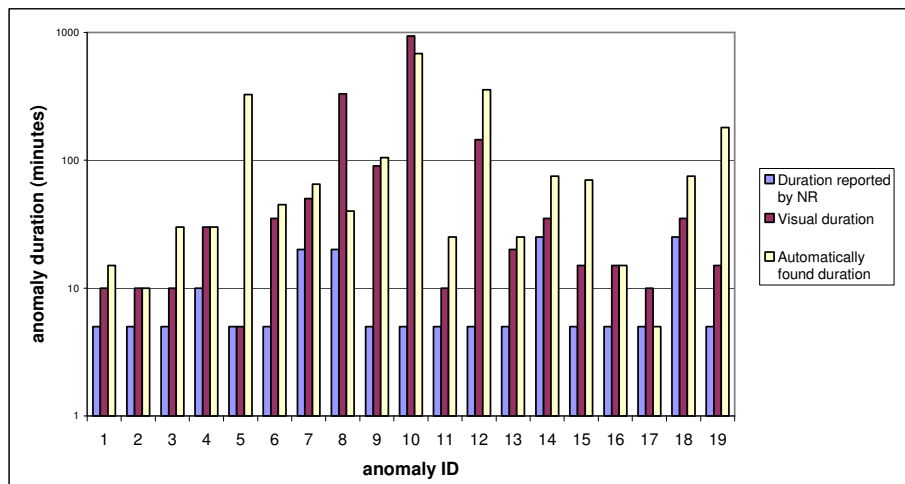


Figure 1 - Comparison among anomaly durations reported by different approaches: NetReflex (left bar), visual duration (middle bar) and automatic approach duration (right bar)

DANTE plans to further refine and validate the method for automatic extraction of flows related to anomalies, and make it part of the anomaly signalling service.

The service is currently in its first phase (final tool tuning and validation, internal testing of results). In its second phase, starting in January 2010, selected anomalies will be signalled to some of the NREN CERTs, to increase confidence in the usefulness and correctness of the information reported.

It is envisaged that at its apex the service will play a fully integrated role in the security workflow of the GÉANT community, interacting with NREN CERTs only (i.e. it will not deal directly with the institution or end-user sources or targets of security incidents). Additional sources of information (resulting from the research and service security activity of GÉANT3) will enhance the overall security service portfolio. An important aspect of the service will be the reports derived from the various systems. These reports will be circulated within the NRENS and/or other cooperating entities.

Acknowledgements

We thank Daniela Brauckhoff and Xenofontas Dimitropoulos for sharing their implementation of Apriori, and Domenico Vicinanza and Mariapaola Sorrentino for the discussion on bandwidth test tools.

References

- [1] <http://nfsen.sourceforge.net/>
- [2] M. Molina, "A Network Security Service for GÉANT2", TNC2008, Bruges
- [3] <http://www.arbornetworks.com/>
- [4] <http://www.lancope.com/>
- [5] <http://www.guavus.com/>
- [6] W. Routly, "A Quantitative Cross-Comparative Analysis of Tools for Anomaly Detection", TF-CSIRT/FIRST technical seminar, Riga, Jan 2009, <http://www.terena.org/activities/tf-csirt/meeting26/routly-anomaly-detection.pdf>
- [7] I. Cunha et al, "Do you trust what flow measurement tools tell you?", Thomson Technical Report CR-PRL-2008-10-000, <http://www.thlab.net/~fernando/papers/CR-PRL-2008-10-0001.pdf>

[8] D. Brauckhoff, X. Dimitropoulos, A. Wagner, K. Salamatian, “Anomaly Extraction in Backbone Networks using Association Rules”, ACM Sigcomm Internet Measurement Conference (IMC), Nov. 2009.

Biographies

Maurizio Molina graduated in Electronic Engineering from the Polytechnic of Turin in 1993. Since then, he has worked in the telecommunications industry, mainly in research centres, including Telecom Italia Labs (Turin, Italy) and the NEC Network Laboratories (Heidelberg, Germany). He has published several papers about IP, ATM traffic modelling and network measurements. He contributed to the ITU-T ATM standardisation process, and to working groups in the IETF (on IPFIX and PSAMP). Maurizio joined DANTE in November 2004, working on performance monitoring, security and authentication. He is a GIAC Certified Incident Handler and currently leads the GÉANT3 Multi-Domain Security task.

Wayne Routly graduated from the Port Elizabeth Technikon in 2004 with a Msc. Information Technology. After graduating he joined the Nelson Mandela Metropolitan University as their Security Engineer & Information Security Officer. After moving to London in late 2005 he joined Fabric Technologies as their Infrastructure Engineer and Team Leader before joining DANTE in 2008 as a Security Engineer and also acquired his CISSP certification. He is currently working on the Network Security Service for DANTE and the growth of DANcert.

Ignasi Paredes-Oliva obtained his Engineering Degree in Computer Science together with a Master in Computer Architecture, Networks and Systems on July 2009 at the Barcelona School of Informatics (FIB) of the Technical University of Catalonia (UPC). He joined the CCABA research centre of the Computer Architecture Department of the UPC on July 2007. He is part of the Load Shedding in Network Monitoring Applications group and he is involved on SMARTxAC project. He started his PhD on July 2009 in UPC focusing his research on sampling techniques and anomaly detection

Ashish Jain is a member of the research staff at Guavus, Inc., where he conducts research on root-cause analysis of security anomalies in backbone networks. Prior to Guavus, Ashish was a member of the technical staff at Juniper Networks and iPolicy Networks, where he worked on intrusion detection and prevention engines. Ashish holds CISSP, CISA and SANS GCIH certifications, and has broad interests in backbone network security, netflow-based spam detection and botnet forensics.