The eNOC featuring DownCollector

Guillaume Cessieux
(CNRS/IN2P3-CC, EGEE networking support)

IN2P3-CC / UREC meeting, 2010-04-07

www.eu-egee.org

Agenda

- **What is DownCollector**
  - History
  - How it works
  - Site vs host availabilty
  - Network checkpoint
- **Overview**
- **Background machinery**
  - Component
  - Implementation details
- **Possible improvements**
- **Conclusion**
- **More information**

What it is?

- **A network monitoring tool**
  - Not interested in host availability but site availability
- **Testing TCP port state of a list of service**
  - Service = Hostname + TCP port
  - Currently testing ~2200 Grid nodes each 2 minutes
    - Tool validated with 3500 nodes
  - Test done with nmap (TCP syn scan)
    - TCP Connection not established
      - *As light as possible*
      - *Tests not flooding logfiles*
- **Showing results and history through a webinterface**

History

- **Started within EGEE-II, 2007-04**
  - Initial release 2007-07, now in version 2.44
  - Most used tools on ENOC's portal
- **Completely home made tool by ENOC team (CNRS/IN2P3-CC) as part of EGEE-SA2**
- **Few key improvements since its release**
  1. Switched from ping to nmap
  2. Alarming system
  3. Adding network checkpoint
- **Ran without any problem since its release**
  - Mainly cosmetics changes lead to version 2.44
  - Quite mature: No service interruption since 2 years
    - Compulsory for meaningful statistics

How it works?

Once a day:

- **Full list of EGE nodes and scheduled downtime is gathered from GOCDB and stored in a db**

Every two minutes through the crontab:

- **List of node and port to scan is gathered from the db**
- **Around 600 fork in charge of 4 tests are created**
- **Results are stored in a dabatabase**
- **When all terminated results are aggregated into another table**
  - then analysed and alarms raised if necessary

Always:

- **The webinterface display results of the latest available test**
- **XML interface provided for on site monitoring (Nagios...) and dashboards (CIC portal)**

Site vs host availability

- **The tool measures host availability**
  - Results are per host
    - OK, UNKNOWN HOST, FILTERED, CLOSED, TEST ERROR, TIMEOUT
  - But we wants only site availability

- **We assume site is unreached when ALL its nodes are unreached**
  - Node unreached = UNKNOWN HOST or TIMEOUT
  - Results per sites: REACHED or UNREACHED
  - A site in scheduled downtime will be prevented to be reported in trouble

Network checkpoint (1/2)

- **Network checkpoint = border router**
  - Demarcation point for ENOC's responsibility
  - Checked only during trouble

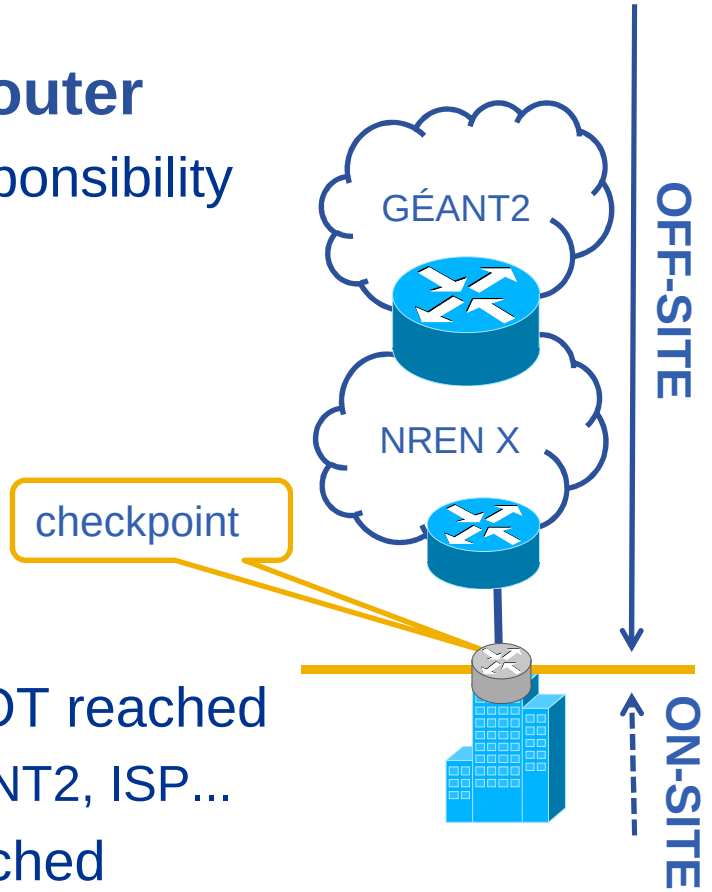

- **Three kinds of troubles**
  1. **OFF-SITE**: Network checkpoint NOT reached
     - Fault in: WAN, MAN, NREN, GÉANT2, ISP...
  2. **ON-SITE**: Network checkpoint reached
     - LAN, power, software ...
  3. **UNKNOWN**: No clear and reliable checkpoint, but site in trouble

Network checkpoint (2/2)

- **List of checkpoints manually computed per site**
  - Find the border router with traceroutes
  - https://ccenoc.in2p3.fr/DownCollector/?v=list_checkpoints
  - Only pingable checkpoints are of interest (= « testable »)

SITE	Grid status	Network checkpoint		Testable
BEgrid-ULB-VUB	Certified	gridce.iihe.ac.be vub-ulb-1.customer.brussels.belnet.net	193.191.4.94	N
BEgrid-UniversiteitAntwerpen	Uncertified	ua.ar1.antwil.belnet.net	193.191.18.10	N
ALBERTA-LCG2	Certified	gsb175-c6509-1-canet.backbone.ualberta.ca	129.128.153.201	Y
Australia-ATLAS	Certified	gw1.er1.unimelb.cpe.aarnet.net.au	202.158.200.250	Y
AUVERGRID	Certified		195.221.123.254	Y

- **This is the only things required to be manually maintained**
  - New sites = new checkpoints
  - Change in site's topology = change in checkpoint

Quick Overview

- **Homepage: Trouble list per site**
  - https://ccenoc.in2p3.fr/DownCollector/
- **History of trouble per site**
  - https://ccenoc.in2p3.fr/DownCollector/?st=INFN-ROMA1
- **Latest node result per site**
  - https://ccenoc.in2p3.fr/DownCollector/?sn=INFN-ROMA1
- **History of results for a particular node**
  - https://ccenoc.in2p3.fr/DownCollector/?nd=CE,atlas-ce-02.roma1.infn.it

Components (1/2)

1. **GOCDB module**
   - To gather list of Grid services and list of scheduled downtimes
2. **CRON module**
   - Launch tests in parallels each two minutes and store results into a db
3. **Rendering**
   - Webinterface:
     - Users: Sites, ROC
   - XML interfaces
     - Users: CIC portal, Nagios
     - Give external viewpoint to sites or ROC
       - *Node filtered by the border router…*

Components (2/2)

4. Alarming system: Appollo

- Trouble computation
- Subscription to be warned through e-mails

5. External: Statistic portal

- Into painlessly and regularly extract some statistics from the database

Current implementation

- **Single server at IN2P3-CC: ccenoc.in2p3.fr**
  - 2 GB of memory
  - Quadricore Intel Xeon CPU 3 GHz
- **Linux redhat RHEL4**
- **PHP 5.1.2**
- **mySQL 5.0.18**
- **Nmap v 3.70**
  - Required to have the setuid bit on nmap
  - Forging packets require root privilege
- **Apache 2.0.55**
- **Closest to GÉANT2 to avoid local/regional networks to interfere**

More details around database

- **Raw data: Node status database**
  - Table containing all tests per host (~2200 entries each 2 minutes)

HOSTNAME	IPv4 Resolved IP used for the test	DATE Expressed in UTC	STATUS Result of the test	TESTEDPORT TCP port tested
ipngri[illegible].fr	1[illegible]78.13	2010-01-25 00:04:06	OK	2119
grid1.c[illegible]s.gr	1[illegible]6.230	2010-01-25 23:44:05	OK	2119
ce-egee.b[illegible].es	1[illegible]39.49	2010-01-25 23:44:05	OK	2119
fal-p[illegible]c.uk	1[illegible]5.24	2010-01-25 23:44:05	TIMEOUT	2119
grid002[illegible]n.it	1[illegible].112	2010-01-25 23:44:04	OK	2119
lfc-pp[illegible].fr	1[illegible]5.163	2010-01-25 23:52:02	OK	5010

- **Aggregated data: Trouble database**
  - Computed from raw data
  - Really smaller (~ 40 entries per day)

ID_TROUBLE	ELEMENT	KIND	STATUS Current status	Counter	Date_In	Date_Out	Date_Updated	LOCATION
100065	UPorto	SITE	REACHED	6	2010-01-26 00:18:01	2010-01-26 00:29:24	2010-01-26 00:29:24	ON-SITE
100066	IEETA	SITE	REACHED	6	2010-01-26 00:18:12	2010-01-26 00:29:24	2010-01-26 00:29:24	ON-SITE
100067	CFP-IST	SITE	REACHED	6	2010-01-26 00:18:26	2010-01-26 00:29:25	2010-01-26 00:29:25	OFF-SITE
100068	UMinho-CP	SITE	REACHED	6	2010-01-26 00:18:38	2010-01-26 00:29:25	2010-01-26 00:29:25	OFF-SITE
100069	LIP-Coimbra	SITE	REACHED	6	2010-01-26 00:18:49	2010-01-26 00:29:25	2010-01-26 00:29:25	ON-SITE
100070	PPS-LIP	SITE	REACHED	6	2010-01-26 00:19:02	2010-01-26 00:29:25	2010-01-26 00:29:25	ON-SITE

AOB

- **Data access policy**
  - Granted to anyone having an IGTF certificate
    - List of Grid hosts is a good target for any attack

- **If site hosting DownCollector is poorly connected or disconneted**
  - Basic test of GÉANT2 pop reachability performed

Possible improvements (1/2)

- **External call to nmap may be rewriten to be more efficient**
- **Improve rendering and data usage**
  - Portal with graphs, stats, per sites, per countries, monthly reports etc.
- **Big improvement: Store only changes in host status**
  - UpUDownDDDDDDDDDDUpUp -> Up Down Up
  - But this is stressing database: For each update a lookup is necessary
- **Improve network checkpoint handling**
  - Automated?
  - What to do with several path to a site?
  - Intermediate checkpoint?
- **Clarify testing of failover situation for hosts and headnode**
- **Study IPv6 support**

Possible improvements (2/2)

- **Threshold on a per subscription basis**
- **More data mining: Handle flapping nodes as a single issue etc.**
- **Adding test latency details**
  - Appeared not suitable: Bunch of test launched simultaneously are interfering on source host
- **Notifications through e-mails are not enough**
  - E-mails often did not reach instantly disconnected entities
- **Correlation with stored traceroute**
  - When site is down compare current traceroute with previously computed to guess faulty element
- **PHP is good, but maybe a complete C++ rewrite could enable less requirement in hardware**
  - Threads instead of PHP fork!
- **Two instances of DownCollector in different countries**
  - But then hard to correlate data

Conclusion

- **Tool very useful and very used by EGEE**
- **Quite mature**
- **Careful Implementation required**
  - Good network connectivity & Reasonably powerful HW
- **Few regular effort necessary**
  - Only maintaining headnodes
- **Webinterface is fine but futher integration could be great**
  - We investigated dashboard etc. but maybe also within process could be good
- **Some adaptation in the EGI era might be necessary**
  - Regionalisation per country?
  - Devolving headnodes management?
  - Replaced by a central Nagios?

More information

- **Few documents were written around**
  - https://edms.cern.ch/document/885463/
  - https://edms.cern.ch/document/970586/
  - https://edms.cern.ch/document/979989/
- **Source code is entirely available under Apache 2 licensing**
  - https://cvs.in2p3.fr/egee-sa2/_2-LICENSE.txt?view=markup
  - https://cvs.in2p3.fr/egee-sa2/1-DownCollector/
- **Feel free to contact us by e-mail**
  - Guillaume . Cessieux @ cc.in2p3.fr
  - or enoc . support @ cc.in2p3.fr