DNV·GL

NETWORK STORM TESTING

VERIFYING THE ROBUSTNESS OF CONTROL SYSTEM COMMUNICATIONS

In this white paper, we discuss a stress situation that can impact communication and redundancy of networked control systems. How can this situation arise, and why should we test for it? Furthermore, we look into what we recommend, how control systems should be tested to avoid that their communication is impacted by network overload.

INTRODUCTION

Let us start with a citation from a senior engineer with the major E&P company of Norway:

"Be a demanding customer, prior to FAT, apply traffic generator packets on network segments to full bandwidth. Peer to Peer, multicast and broadcast packets. Graceful reconnect after storm or need for restart?"

– Sr. IT Security Engineer, Statoil ASA

But what is an unexpected network overload, or network storm situation actually? A network storm can be compared to a room full of people talking loudly and making a conversation between 2 individuals impossible. Network storms or stress situations are caused by excessive amounts of traffic, i.e. a flood of packets in packet switched communication, such as Ethernet-based Local Area Networks (LANs). The traffic that causes a network storm is often broadcast or multicast messages, meaning that all or a group of hosts on the same network receive the traffic. When all available bandwidth is consumed by the network storm, the network is rendered useless and all applications and devices stop functioning properly.

Unless appropriate protocols are used to block redundant network paths, communication loops may be formed. If undetected and unhandled, these loops may eventually lead to communication faults. In Ethernet networks, an unmanaged loop is dangerous because broadcast and multicast messages are continuously passed until the network gets overloaded, a situation called a broadcast storm.

A normal Ethernet switch forwards broadcast and multicast traffic on all its ports. Other (up- or down-link) switches receiving these broadcast or multicast messages, will again forward them to all their ports and so on. In an Ethernet network, any looped packet might remain on the network forever. A network storm, i.e. a network stress situation can arise in various ways and can cause a Denial of Service (DoS) in the worst case.

Probably, the most common reason for a network storm is cabling problems, in particular if a cable loop is present. Other factors contributing to a network stress situation are:

- Poor network management and monitoring;
- Improperly maintained network configuration – often due to inexperienced network engineers;
- Inadequate documentation or undocumented network design – leading to bad network management and complicated control system trouble shooting.

When a network overload occurs it is important to understand how to identify and remediate the issue quickly and understanding why it happened. Without a proper root cause analysis the symptoms will return. A large amount of uncontrolled network traffic negatively impacts business systems and processes. It can be caused by cyber-attacks, but can also be due to simple mistakes or over-reactions to normal conditions. When a network storm occurs, two important questions have to be answered: why did the network storm occur in the first place? And how can it be stopped and any future incidents be prevented?

A WELL-KNOWN INCIDENT

Brown's Ferry Nuclear Plant incident, Athens, AL, USA, Aug. 2006.
Operators manually scrammed Brown's Ferry, Unit 3, following a loss of both the 3A and 3B reactor recirculation pumps. The root cause was the malfunction of the VFD controller due to excessive traffic on the plant Ethernet based integrated computer system network [1].

Unit 3 was manually shutdown after the failure of both reactor recirculation pumps and the condensate demineralizer controller. The condensate demineralizer used a PLC; the recirculation pumps depend on VFDs to modulate motor speed. Both PLCs and VFDs have embedded microprocessors that can communicate data over the Ethernet LAN.

Both devices, however are prone to failure in high traffic environments. A device using Ethernet broadcasts data packets to any other device connected to the network. Receiving devices must examine each packet to determine which ones are addressed to them and to ignore those that are not. It appears as the Brown's Ferry control network produced more traffic than the PLC and VFD controllers could handle; it is also possible that the PLC malfunctioned and flooded the Ethernet with spurious traffic, disabling the VFD controllers; tests conducted after the incident were inconclusive. This demonstrates the effect that one component can have on an entire process control network and every device on that network.

OVERLOAD

What kind of failures are actually introduced during a network storm test? First we need to consider the ISO open system interconnection (OSI) layered protocol stack [2], illustrated in Figure 1. Network storm on OSI Layer 2 could typically be hardware failure (also known as the babbling idiot failure), whereas a storm on a multicast group, i.e. Layer 4, could be software fault or misconfiguration, as these groups are often used by user traffic in typical industrial control applications.

Normally, the intention of a network storm test is to simulate failure scenarios that might happen due to:

- hardware malfunction;
- software bug or misconfiguration;
- additional equipment connected to under-dimensioned network segment;
- broadcast storm (e.g. ARP storm due to cabling failure);
- firmware upgrade (e.g. switch re-configuration).

7. Application	HTTP, SMTP, FTP, DNS, ...
6. Presentation	HTML, JPEG, MPEG
5. Session	SSL, RPC, SQL, ...
4. Transport	TCP, UDP
3. Network	IPv4
2. Data link	Ethernet MAC
1. Physical	Ethernet PHY

Figure 1. The ISO/OSI layers with examples

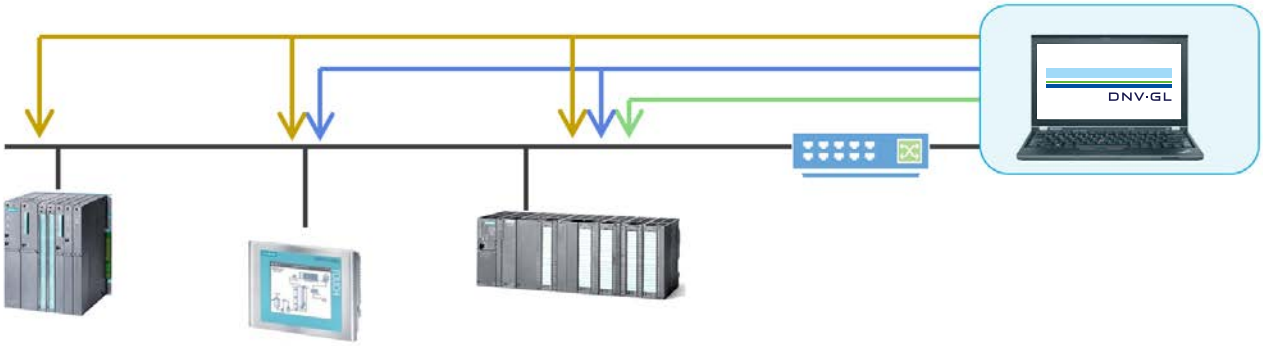

Figure 2 Network Storm test setup example

OBJECTIVE OF NETWORK STRESS TESTING

How much testing is done is always a tradeoff, e.g. the identified risk compared to available time and/or money.

Considering the OSI layers 2 and 3, there is a difference if network storm testing is conducted on a broadcast address, on a multicast address, or simply as unicast (point-to-point) traffic, in terms of which nodes will be affected (receive traffic), see Figure 2. The purpose of using broadcast is usually to test how redundancy and communication between multiple nodes (e.g. PLCs) behaves under network stress, whereas with unicast traffic behaviour of a single node can be observed.

Referring to networks storm tests on layers of the ISO/OSI stack in Figure 1:

- On Layer 2 all nodes will be targeted that are physically connected to the nearest switch, whether or not there is VLAN configured, and whether or not they are in the same Layer 3 segment.
- On Layer 3, all nodes in the same logical segment are targeted, but not others, even if they are physically connected to the same switch. However, flooding could propagate through several switching devices.
- Targeting one multicast address (on Layer 3) will affect the nodes that are subscribed to that multicast group, and will inhibit their communication.

In addition to raw data storms, "smarter" ways of testing may be performed utilising the structure of the protocols and packages, for example ARP cache poisoning, flooding with IP fragments or ICMP messages, or opening a large amount of unused TCP connections, to name a few.

A further objective is to test (the trigger of) any traffic rate limiting function (if such function is implemented). This is the reason for conducting a storm simulation where the traffic rate is ramped up step-wise from 0% to a 100%.

All in all, a network storm test aims to evaluate a control systems' performance when subjected to high network loads, to answer questions, such as:

- What are the effects of loss of communication?
- What are the effects of loss of redundancy?
- What are the effects of loss of HMI connections?
- What are the effects of a freezing controller?
- Are there any capacity problems?

Possible test targets

The outcome of a network stress test might be different depending on what type of equipment is targeted. The most common types of targeted equipment are Programmable Logic Controllers (PLCs); operator stations, or multifunctional displays, i.e. HMIs; various servers and historians; or network switching gear, such as routers, switches, and firewalls.

Acceptance criteria

When is the outcome of a network storm test satisfactory? Acceptance criteria are usually defined as a combination of verifying that:

- the control system stays functional and can be operated as expected by an operator;
- warnings or alarms are correctly issued for the component that is subjected to high network loads;
- any unexpected behaviour of the control system is detected;
- the device under test preserves a secure state if any fatal errors occur;
- the device under test is able to handle DoS attacks, and in case of resulting failure of communication, is able to recover.

TOOLS AND SOFTWARE

DNV GL Marine Cybernetics services rely on a combination of in-house developed and commercial or public software. In most projects the in-house developed MC Network Storm simulator is used, however commercial tools from Wurldtech and Codenomicon are also available, in particular to implement the CRT Test Requirements for Protocols in the EDSA Certification. These test requirements are introduced in the next section.

RELATED STANDARDS AND RECOMMENDATIONS FOR REFERENCE

In addition to the test cases developed by DNV GL Marine Cybernetics services, relevant standards and requirements can also be employed, such as the industry-wide well recognized ISA/IEC 62443-4-2 (Embedded Device Security Assurance Certification) standard [3]. In particular, the part Communication Robustness Testing (CRT) of this standard is relevant for network storm testing. The sub-specifications discussing robustness testing are EDSA-310 CRT Common, and EDSA-401 through EDSA-406 for specific protocols.

From a DNV GL perspective network storm testing is in line with the requirements from DNV GL Offshore Standard, DNV-OS-D202, Ch. 2., Sec. 3, §3.3 Network analysis and §3.4 Network test and verification [4]; and DNV GL Rules for Classification, Ships, Pt. 4., Ch. 9., Sec. 4., §3.3 Network test and verification [5].

EXAMPLE FINDINGS

DNV GL Marine Cybernetics services have conducted a number of network stress tests, often in combination with a Hardware-In-the-Loop (HIL) test program, targeting various elements of a variety of control systems [6]. In case of Dynamic Positioning (DP) systems, the following findings were results of a network storm test:

- In a DP3, triple-redundant DP system, GPS signals, Hydro-acoustic Position Reference (HPR), Gyros and Vertical Reference System (VRS) signals lost, leading to a loss of position;
- All remote I/O signals lost resulting in loss of all Gyros and VRS signals. Loss of all position reference systems, leading to loss of position;
- Overload and hence cut-off DP controller;
- DP controllers with correct sensor data being voted out during network overload;

In Power Management (PMS) systems, network stress resulted in the following findings:

- 2 out of 4 generators stopped after a successful and ended network storm, resulting in loosing half of the power plant, and a possible (partial) black-out;
- PMS controllers had to be re-started and synchronized after a ended network storm;
- A PMS controller requesting a reset for synchronization during a network storm, when it is reset a bus-tie is opened leading to a critical system state or (partial) black-out;

In some of the most complex control systems, drilling control system networks can be impacted by network storm as in the following scenarios:

- Communication lost between drillers chair and top drive controller, while the drill bit is getting stuck. Uncontrollable top-drive increases tension beyond limits, which leads to damage of the drill-string, the drilling equipment, or the well itself;
- Communication lost between the drillers chair and the drawworks controller during lowering or hoisting results in uncontrollable movement of the travelling block, which can lead to collisions with other equipment, the drill-string, or even with the drill-floor or the frame of the derrick;
- Communication of the anti-collision system for two or more drilling controllers is disturbed while machinery is moved via automated sequences, leading to a collision, damage to equipment, and a halt in production;
- Under extreme circumstances, while communication between controllers in an intelligent drilling control system is disturbed, none of the built-in interlocks is able prevent dropping a pipe during a hand-over between two machines (e.g. a piperacker and an iron-roughneck), which leads to damage to the equipment;

MITIGATION IN SWITCHING EQUIPMENT

Storm control is a feature increasingly present in managed switches. It enables a switch to monitor traffic levels and to drop broadcast, multicast, or unknown unicast packets when a specified traffic level is exceeded, thus preventing packets from degrading the network performance. As an alternative to having a switch drop packets, it is, in some implementations, possible to configure the affected interfaces to shut down temporarily when the storm control level is exceeded.

Broadcast and multicast storm control allows suppression of excessive inbound unicast, multicast, or broadcast traffic on Layer 2 interfaces. It is considered important to protect against broadcast storms resulting from misconfiguration, or even unicast storms created by malfunctioning network cards. A maximum threshold can be configured in bits or packets per second (bandwidth-based control), or as a percentage of the interface bandwidth (level-based control). If incoming traffic of the specified type exceeds its threshold during a polling interval (typically one second), traffic is blocked until the incoming rate drops below the configured falling interval. Level-based storm control is applied on the combined traffic streams, and typically has a factory default value of 80%.

Rate limiting is another function that can limit packets with an invalid source MAC address on a secure port. If the rate is exceeded, rate limiting for the port is triggered. Intelligent managed switches can also be configured with limits on the amount of traffic a port is allowed to handle, to some extent preventing network overload situations escalate from one port to the rest of the ports.

By dividing the Ethernet broadcast domain into smaller logical pieces by physical or logical separation, the effect of a network storm can be limited to the affected network segment. By using routers to connect the segments together, data exchange can be realized.

The Rapid Spanning Tree Protocol, RSTP (IEEE 802.1w) prevents loops from being formed in the network when devices are interconnected via multiple paths by logically closing a connection or path until a failure is detected and the connection or path is re-enabled.

CONCLUSION

Regardless of what mitigating actions are put in place to avoid network stress situations impacting production, they must be tested and their functionality and configuration should be verified appropriately.

Currently, DNV GL Marine Cybernetics services can provide network storm testing for the following application areas:

- as a sub-test of a Hardware-in-the-loop (HIL) test program;
- as part of an FMEA trial or FAT;
- or as a standalone test, for example to verify mitigating actions.

REFERENCES

[1] NRC Information Notice: 2007-15: Effects of Ethernet-based, non-safety related controls on the safe and continued operation of nuclear power stations, Apr. 17. 2007.
[2] The Open Systems Interconnection (OSI) model, see https://en.wikipedia.org/wiki/OSI_model
[3] www.isasecure.org/en-US/Certification/IEC-62443-4-2-EDSA-Certification
[4] DNV GL Offshore standard, DNVGL-OS-D202, Automation, safety and telecommunication systems, July 2015
[5] DNV GL Rules for classification: Ships, DNVGL-RU-SHIP, Part 4 Systems and components, Chapter 9 Control and monitoring systems, October 2015
[6] www.dnvgl.com/maritime/marine-cybernetics.html

SAFER, SMARTER, GREENER

DNV GL AS
Veritasveien 1
1322 Høvik, Norway

DNV GL – Maritime
Mate J. Csorba
Principal Specialist,
Marine Cybernetics services,
Offshore Class
mate.csorba@dnvgl.com

www.dnvgl.com/maritime

About DNV GL
Driven by its purpose of safeguarding life, property and the environment, DNV GL enables organizations to advance the safety and sustainability of their business. Operating in more than 100 countries, our 15,000 professionals are dedicated to helping our customers in the maritime, oil & gas, energy and other industries to make the world safer, smarter and greener.

DNV GL is the world's leading classification society and a recognized advisor for the maritime industry. We enhance safety, quality, energy efficiency and environmental performance of the global shipping industry – across all vessel types and offshore structures. We invest heavily in research and development to find solutions, together with the industry, that address strategic, operational or regulatory challenges.

 Design: Maritime Communications