Information Security Now – 47

Intelligent Security?

John Mitchell

When IBM's *Watson* program won America's *Jeopardy!* quiz in 2011, the world woke-up to the fact that fast computers processing natural language could beat human's in a general knowledge quiz. Did this mean that Artificial Intelligence (AI) had finally arrived and if so, could it be used for other things? AI was first taken seriously in the mid-1980s when it was realised that human knowledge could be captured as a set of rules in a decision tree which could then be processed incredibly fast by a computer. Many companies were founded and died during the next decade as it was learnt that human intuition was much more fuzzy logic than hard rules. This did not dissuade firm's such as IBM, but it did make clear the huge gap between what I define as 'Specific AI' and 'General AI'. The main difference between these is that the former answers a specific question, such as 'am I under attack', while the latter formulates the question, 'what does an attack look like'. So, specific AI is great for dealing with known security threats. It is primarily rule based and can identify known threats and vulnerabilities and can suggest ways of dealing with them based around its ability to process vast amounts of data very, very quickly. Now, so called 'machine learning' has been likened to General AI, but it is no such thing. It may enable the detection of a new pattern which is not immediately obvious to a human, but any subsequent action is predicated from within the algorithm itself and not generated from outside of it. This may be a draw-back, but it is still a useful tool because of its ability to both identify new patterns and then react consistently and quickly to the given circumstances.

However, we need to be cautious in totally removing any human intervention. On 26 September 1983 Stanislav Petrov was the duty officer at the command centre for the USSR's nuclear early-warning system. The system reported that a missile had been launched from the United States, followed by five more. Petrov judged the reports to be a false alarm and his decision to disobey standing orders, is credited with having prevented an erroneous retaliatory nuclear attack on the United States and its NATO allies. Subsequent investigation confirmed that the Soviet satellite warning system had indeed malfunctioned. So here is the dilemma for fully automated IT security services. They may be fast, but what are the consequences if they are wrong? We know that autonomous high frequency trading systems can bring financial markets to their knees. We may watch the systems, but the computers usually move far too quickly for us to intercede.

This is a straight risk analysis scenario. Just how much autonomy are we going to allow our security systems? Well, it's a straight-forward cost-benefit analysis. What is the consequence to our company/nation if our security algorithm does not act quickly enough to contain the threat, as compared with it acting very fast and being very wrong? To quote H.L. Mencken, "for every complex problem, there Is an answer that Is clear, simple, and wrong". We must be careful with

feedback-loop decision processes, because they can quickly spiral out of control, as was seen with the high frequency trading debacle. On the other hand, fast processing with an in-built safety factor is just what we need for our AI managed computer security systems. To some extent we already have this with our firewall software, which recognises when it is under attack and fails-safe by denying any further entry. I accept that this is a self-imposed denial of service, but the smart ones only stop the incoming traffic and still let outward bound transactions go through. The shutting, or slowing down of the inbound traffic allows time for the humans to intercede. However, imagine what would happen if the AI firewall automatically conducted a retaliatory attack against what it identifies as the rogue IP address? We know that cyber criminals/nations seldom attack from their home base, but tend to either piggy-back, or spoof, another identity. There is every likelihood of our AI bot attacking an innocent victim, which will itself then automatically react to our attack. Just as in the Stanislav Petrov situation we have every likelihood of a Cyber World War breaking out unless we build-in some human intervention. The banks have run some stress testing of simulated cyber-attacks against our financial systems, but the game is far bigger than that. Most of our critical national infrastructure is in private hands and the last thing those companies want to do is to throw money at a problem which is basically invisible. How do we know this? Review the recent cyber-attacks against Equifax, Talk-Talk, the central bank of Bangladesh, Tesco Bank and Yahoo to name but a few, together with the WannaCry ransomware attack which hit the NHS along with many others. So how much should we be spending on security and where should that spend be targeted?

Recent research by the Sans Institute shows that no industry sector spends more than 12% of its IT budget on security, with the majority at less than 9%. You will notice that this is from the IT budget, yet the impact is felt company wide. The consequences to UK plc if our national infrastructure (power, water, sewerage, internet, rail, air and road systems) on which we all depend, is far greater than a mere 12% of our combined IT budgets. We know that Russia cyber-attacked Estonia as far back as 2007 in a dispute over a statue. They used distributed denial of service attacks against websites of Estonian organizations, including its parliament, banks, ministries, newspapers and broadcasters. Imagine what their capability must be now? We really do need competent AI security systems if we are to survive such attacks. Fortunately, our government has recognised this and has proposed allocating more funds from the defence budget for cyber warfare (bad luck you kinetic chaps), but if most of our infrastructure is in private hands how is it going to persuade them to stump-up the necessary money for their own defence? After all, they first need to consider their shareholders' dividends against what they may perceive as a vague future threat. But it is not a future threat it is a current clear and present danger to our society. Perhaps the regulators have a role to play here?

John is a previous member of Council and the Risk, Audit and Finance Committee. He is currently Treasurer for the Information Risk Management and Assurance (IRMA) specialist group. He can be contacted at: john@lhscontrol.com, www.lhscontrol.com, or on +44 (0)7774 145638

.