No.44
November 2020

HardenStance Briefing

Trusted research, analysis & insight in IT & telecom security **PUBLIC/UN-SPONSORED**

New Standards for Securing AI

- AI algorithms deployed out in the field today are either secured by proprietary security mechanisms of varying quality or they are not secured at all. There don't appear to be any global AI security standards of note.
- ETSI is addressing this with its Secure AI (SAI) Industry Specification Group (ISG) in which 46 companies are participating. The work is very impressive, focused on key areas like AI threat ontology; AI supply chain security and AI threat mitigation.
- Industry and society would benefit from cyber security vendors participating. There may also be more value in taking part than these firms recognize. There's a lot to be gained from the Webscale giants participating too – but only if it's in good faith.

More Effort Going into AI Ethics than AI Security

There aren't too many companies that aren't committing time and money to Artificial Intelligence (AI) nowadays. IDC reckons $50 billion will be spent on AI during 2020 and forecasts this will rise to $110 billion in 2024.

The battle of AI ethics is only just beginning.

For those of us who worry about giving algorithms ever more power over our lives, it's at least partially reassuring to observe some of the collaborative work being done to constrain, limit and regulate the use cases that we allow present day AI algorithms to be applied to. One high profile example of AI ethics asserting itself is the decisions by IBM, Amazon and Microsoft – following a lot of public pressure – not to sell their facial recognition technology to law enforcement agencies for fear of exacerbating social divisions based on race.

Significant Effort to drive AI Ethics into Regulations and Standards

These battles are only just beginning. Facebook's algorithms still drive us to intensify our prejudices rather than challenge them – it has the white supremacist and other extremist groups flourishing on its platform to prove it. Un-remediated though these toxic outcomes still are, we can at least be encouraged by the significant collaboration across government, industry, academia and the media now to try and drive AI ethics deeper into government regulations, corporate responsibility and technical standards.

Why ETSI for Artificial Intelligence Security Standards?

'Why ETSI?' for creating AI security standards, you might reasonably ask. Telecom operators certainly need AI for critical functions like automating their networks – and they need those AIs to be secure. But their needs and requirements are arguably no greater than many other industrial sectors.

One key reason is that the telecom sector has a comparatively good track record of collaborating to define technology standards – unlike many other sectors of industry, the telco business model simply doesn't work without standards-based inter-operability. The other reason is that by their nature, telecom operators are also best placed to shut down volumes of connected devices that have gone rogue – whether that's due to an AI having been corrupted or for any other reason.

AI Security and AI Ethics are Complementary

But what if an AI that is perfectly optimized from an ethical perspective gets hacked? This is the field of AI security. It doesn't concern itself so much with the ethics of AI. Rather it's concerned with protecting AI against any form of unauthorized manipulation that generates decisions, hence outcomes, that are contrary to the original purpose for which it was deployed. The two fields are complementary but distinct.

Whereas industry-wide collaboration around AI ethics goes back many years, there hasn't been anywhere near as much collaboration in developing standards for AI security. Despite some progress in cyber security awareness, organizations still think first about how AI can improve their business processes or generate more revenue. Security still tends to be an afterthought.

There don't appear to be any significant international industry standards for AI security. All the AIs that are commercially deployed today are either secured by proprietary security mechanisms of varying quality or they're not secured at all – and that includes all those cyber security products out there that are positioned as being 'AI-driven'.

All the AIs that are commercially deployed today are either secured by proprietary security mechanisms of varying quality or they're not secured at all.

The ETSI SAI ISG was Formed at the End of Last Year

This is why HardenStance is delighted to report on the progress of the European Telecommunications Standards Institute's (ETSI) new Industry Specification Group (ISG) focused on AI Security. Formed at the end of last year, the link to the Securing AI (SAI) page on the ETSI portal is provided in the "More Information" section at the end of this Briefing. The scope of the SAI ISG is captured in **Figure 2** below. As also shown in **Figure 3**, the SAI group already has 46 members, 6 participants and one counsellor.

Figure 2: The Scope of ETSI SAI ISG

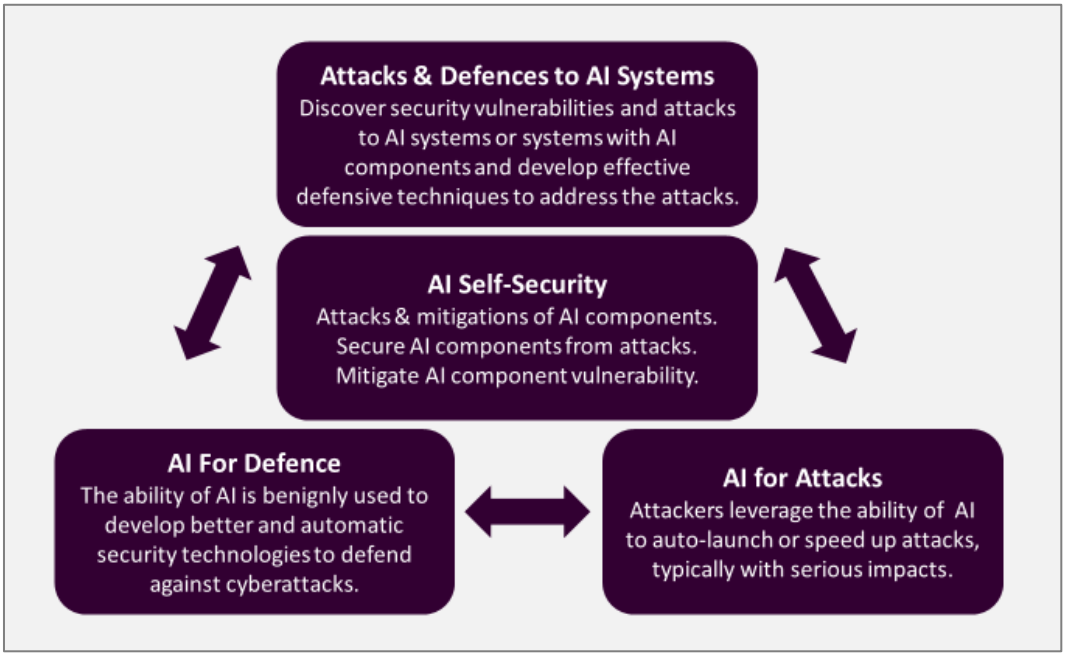

Source: ETSI SAI ISG

The volume and quality of the work is remarkable, especially when you consider that the group has only been going for a year. It's well worth delving into some of the detail. The work is organized around five Phase one Work Items as follows:

1. The Problem Statement

This group is attempting to arrive at a comprehensive definition of exactly what the challenge of securing AI is. It aims to describe some of the main challenges of securing AI-based systems and solutions, including challenges relating to data, algorithms and models in both training and implementation environments. The focus is on challenges which are specific to AI-based systems, including poisoning and evasion. In the case of poisoning for example, it is looking at the risk of poisoning of the data set and of the model as well as the algorithm itself. (Rapporteur: Queens University, Belfast)

The Threat Ontology group is developing an ontology to capture an overarching philosophy around how specific threats to an AI can be described.

2. AI Threat Ontology

This group is developing an ontology to capture an overarching philosophy around how specific threats to AI can be described. This group is constructing a language of relationships between things where an AI is one of those things. To give one very simple example, if an AI is given a sensor and you want to tell it it's a camera, what are the criteria against which an AI is to determine that something purporting to be a camera is indeed a camera? What is the language used to describe the relationships between things where an AI is one of those things? (Rapporteur: C3L).

3. Secure AI Supply Chain

This group is looking at how a trusted supply chain in AI components and systems gets established, whether that be at the level of an organization assembling its own products or at the level of a buyer investing in a finished AI-enabled product. How can systems and components be certified as meeting a given security standard and what security benchmarks will AI components, systems, and products must achieve to obtain certification? (Rapporteur: UK National Cyber Security Centre (NCSC).

4. Security Testing AI

We know how to test routers and firewalls against cyber security attacks but how do you test AIs in the same way? This group is looking at developing a set of test specifications for how AIs get tested to ensure the maximum possible robustness against a variety of different types of attack (Rapporteur: Fraunhofer Focus).

5. Mitigation Strategy

This group is developing guidelines for mitigating against threats introduced by adopting AI (Rapporteur: Huawei).

Phase Two

In Phase two, which is only now getting underway, a new Work Item with InterDigital as Rapporteur has the remit of identifying the role of hardware, both specialised and general-purpose, in the security of AI. This will address the mitigations available in hardware such as roots of trust, secure components, and trusted boot processes to prevent attacks and address the general requirements on hardware to support SAI.

Figure 3: Participation in ETSI SAI ISG

46 Members, 6 Participants, 1 Counsellor		
MEMBERS	InfoCert s.p.a	Queens University Belfast
Affnor	Intel Corp UK Ltd	Rogers Communications Canada
Approach Affinity	InterDigital Inc	SBS aisbl
Blekinge Tekniska Högskola	Irdeto BV	SGDSN
BMWi	Medtronic Bakken Research BV	Telefonica S.A
BT plc	Dutch Ministry of Economic Affairs	Thales
Cadzow Communications	Motorola Mobility UK Ltd	Traficom
Certicar SL	Motorola Solutions Danmark A/S	Uninfo
CIS	NCSC (UK)	University of Bradford
CKH IOD UK ltd	NEC Europe Ltd	**PARTICIPANTS**
DEKRA	Nextra Partners	ENISA
Deutsche Telekom AG	Nimbus Technologieberatung	Last Press Label
Euro Nova	Nokia Germany	NetMagic Associates LLC
Fabasoft AG	NXP Semiconductors Netherlands	SK Telecom
Facultad de Informatica	OPPO	Tecnalia Research & Innovation
Fondazione Links	PCCW Global B.V	Telecom Foresight Consulting
Fraunhofer Focus	Philips International B.V	
GIE ANEC	Public Safety Canada	**COUNSELLORS**
Huawei Tech (UK) Ltd	Qualcomm Technologies Int	European Commission

Source: ETSI SAI ISG

Membership and Participation in SAI ISG

Encouragingly, the membership of SAI has more than doubled from 21 at the end of last year to 46 now. This growth is in significant part down to the effects of the Coronavirus pandemic. Although standards-goers lament that they can no longer reach compromises with peers and competitors quite as easily over coffee in corridors or white-board in person any more, on the upside some companies are more willing to justify the time and cost of someone participating in standards groups remotely than traveling for in-person meetings.

None of the Webscale giants will commit their resources to multi-lateral, peer to peer standardization in securing AI at this time.

The current membership of SAI ISG is shown in **Figure 3** above. Among the factors that stand out are what seem to be a disproportionately high representation of both British and Dutch organizations. The UK's participation is notable for the number of participants from government, industry and academia; the leadership role taken by its National Cyber Security Centre (NCSC), and the use by American (Motorola and Intel), Chinese (Huawei) and Japanese (NEC) companies of UK subsidiaries to participate in the group.

The Webscale Giants are Noticeable by Their Absence

Most segments of the ICT industry are well represented in the SAI ISG. Two groups are noticeable by their absence, though. One of them is the Webscale giants, some of whom (Google and Facebook) are actually ETSI members.

These firms are at the cutting edge of developing, trialing, and implementing AI. But consistent with how most of them have so far rigorously defended their work in advanced algorithms from meaningful public scrutiny or accountability, none of them will commit their resources to multi-lateral, peer to peer standardization in securing AI at this time. Their participation in good faith would, of course, be enormously valuable. That said, their participation with a view to engaging in spoiling tactics would not.

The other noticeable absentee is the cyber security community itself. Look at **Figure 3** again and look for a leading pure-play cyber security vendor lending its expertise to create baseline security standards for protecting AI. There aren't any. Striking isn't it?

Maybe it's the 'lending a hand' that's the problem. Participating in global standards may not appear to be all that directly relevant to many of these companies' business models. But take a look at which countries are investing heavily in AI across multiple different global industry standards forums and one sticks out: China. Firms that sit back and leave AI security standardization to take whatever path it may, might come to regret this later on when other companies' IPR forms the core of those standards ■

More Information

- ETSI SAI ISG
- HardenStance White Paper: "AI In Cyber Security: Filtering Out The Noise"
- HardenStance Principal Analyst: patrick.donegan@hardenstance.com
- Register for free email notifications whenever HardenStance publishes new content.
- HardenStance received no payment - direct or in-kind - for publishing this Briefing.

About HardenStance

HardenStance provides trusted research, analysis and insight in IT and telecom security. HardenStance is a leader in custom cyber security research and leading publisher of cyber security reports. HardenStance is also a strong advocate of industry collaboration in cyber security. HardenStance openly supports the work of key industry associations, organizations and SDOs including NetSecOPEN, AMTSO, The Cyber Threat Alliance, The GSM Association, OASIS, ETSI and TM Forum. www.hardenstance.com

HardenStance Disclaimer