KASPERSKY lab

The Memory Market: Preparing for a future where cyberthreats target your past

October 2018

Contents

Introduction

There is an episode in the dystopian near-future series Black Mirror about an implanted chip that allows users to record and replay everything they see and hear. In the episode, *The entire history of you*, this capability ultimately proves devastating. Despite this, a recent YouGov survey found that 29% of viewers would be willing to use the technology if it existed.

The desire to remember things is overwhelming. From our most precious personal experiences to the academic knowledge that gets us through exams, our lives are enriched and even defined by the memories we have. It follows that the loss or decline of memory, for example through illness or injury, undermines our ability to function normally and our overall quality of life.

If the Black Mirror scenario sounds a bit too much like science fiction, it's worth noting that we're already well on the way to understanding how memories are created in the brain and how this process can be restored. Earlier this year proof of concept experiments showed that we can boost people's ability to create short-term memories. These experiments were undertaken by DARPA (the U.S. Defense Advanced Research Projects Agency) as part of its program to help military personnel whose head injuries have impaired their ability to form memories.

The hardware and software to underpin this exists too. Neurostimulators, connected implants that can target and stimulate the brain to restore its function, are being used to address the symptoms of Parkinson's disease and even depression. It's not a huge leap for these devices to become 'memory prostheses'. Five years from now we may be able to electronically record the brain signals that build memories, and then enhance or even rewrite them before putting them back in the brain.

To function effectively, any memory implant will need to communicate over a wireless network with a handheld device or medical professionals. Even at today's level of development – which is more advanced than many people realize – there is a clear tension between patient safety and patient security. For example, the implants need to have a software backdoor and override capability in case of a medical emergency. This makes them safer, but more insecure.

Unlike other implants, such as cardiac pacemakers or insulin pumps, neurological implants have more complex (although less life-critical) functions, and are vulnerable to cyberattacks with a broader range of consequences, including influencing a patient's thoughts and behaviors. This raises many moral, ethical and legal issues, but also some important cybersecurity ones.

How can we best prepare for the threats and vulnerabilities likely to target this rapidly evolving technology over the coming years and decades? The Memory Market project considers that question.

The first and most important step is to understand the current and emerging threat landscape and identify the points of greatest risk. Researchers from Kaspersky Lab and the University of Oxford Functional Neurosurgery Group have undertaken a practical and theoretical threat review of existing neurostimulators and their supporting infrastructure. This report is the outcome of that research. It looks at the current landscape in terms of hardware and software, the potential attack vectors and opportunities for compromise, and suggests what we need to do to address these and prepare for the future.

The report in a nutshell

- Researchers have found potential threat vectors in terms of exposed connected infrastructure, insecure data transfer, programming software stored on inadequately protected commercial-grade devices and insecure behavior by on-site medical staff.

- Evolution of memory implants over the coming decades could enable cyberattackers to steal, sell, spy on, manipulate, implant or alter memories.

- Collaboration between healthcare professionals, the security industry, manufacturers and developers and associated professional bodies will be essential to understand and address emerging risks.

Neurostimulators – the current landscape

Deep brain stimulation (DBS) is a neurosurgical procedure that involves implanting a medical device called a neurostimulator or implantable pulse generator (IPG) in the human body to send electrical impulses, through implanted electrodes, to specific targets in the brain for the treatment of movement and neuropsychiatric disorders. DBS has been selectively approved for use since the late 1990s, and has shown therapeutic benefits for otherwise treatment-resistant disorders such as Parkinson's disease, essential tremor, dystonia, chronic pain, major depression, and obsessive-compulsive disorder.

Fig: 1.0 - The connected architecture of implantable pulse generators. Source: University of Oxford, Functional Neurosurgery Group

Electrodes are carefully implanted in the patient's brain, precisely targeted at the region that neurosurgeons think will best treat the patient's ailment, and connected to wires that run beneath

the skin (Fig: 1.0). The wires are attached to an IPG, which is implanted underneath the skin, most often just below the clavicle. The IPG contains a battery, which may or may not be rechargeable depending on design, a microprocessor that controls the stimulation, a small amount of digital memory, and a radio antenna for sending and receiving instructions from programmer devices. All of these components are contained inside a titanium case that is designed to be compatible with the tissue of the human body.

In the most cutting-edge IPGs, the radio antenna communicates using the Bluetooth radio protocol, allowing them to connect with smartphone or tablet-type devices. Users are able to program in the stimulation parameters (such as the frequency and amplitude of the current being delivered to the brain) check the status of the battery, and switch between different device modes using a touchscreen interface similar in design to apps that even elderly patients will have some familiarity with. The IPG is also capable of having its firmware updated over the air, allowing new features to be added or bugs to be patched after first release.

Security overview – the current landscape

Brain-jacking and malicious memory alteration pose a variety of challenges to security – some quite novel or unique, some almost universal. Because neurostimulators/IPGs are just starting their long path of evolution and much of the work involving them is handled in medical research laboratories, it's not easy to practically test the technology and associated software for vulnerabilities. However, much can be learned from handling the devices and seeing them used in situ, and our research involved both.

The challenge of connected infrastructure

Before talking about the practical aspects of cybersecurity, possible vulnerabilities and different drawbacks of neurostimulators, it is worth looking at the infrastructure around them.

As can be seen from Fig. 1, the trend towards using wireless communications between patient implants and other external equipment could result in unwanted consequences if the pairing and communication process is not secure (for example, if data transferred to and from the device is not encrypted).

As connected healthcare evolves over coming years, we are likely to reach a point where manufacturers and medical institutions remotely monitor patients' telemetry and adjust the settings of implanted devices in order to provide responsive treatment to anyone, at any time.

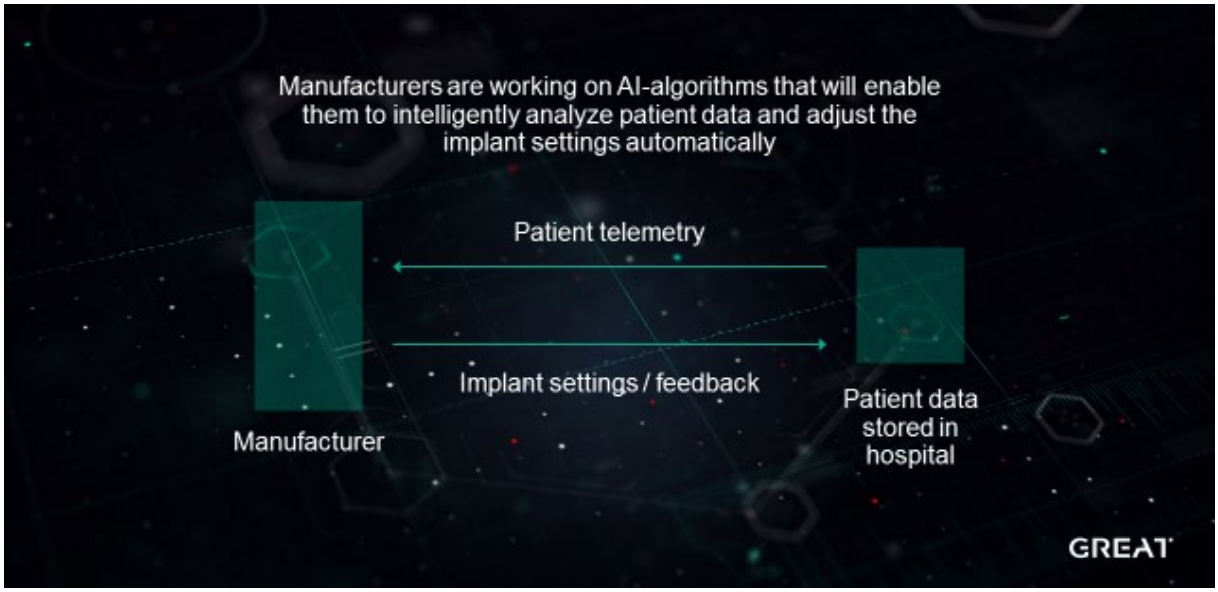

Fig: 2.0 - The supporting infrastructure for neurostimulators in the next few years. Source: Kaspersky Lab

In fact, this capability already exists. In a number of particularly challenging cases, cardiac implants use such remote controls for constant monitoring of the patient. This gives rise to a potential risk where an attacker could seize control of the network and take over the remote control of all the devices on that network.

Fig: 3.0 - The supporting infrastructure for neurostimulators in 10 years. Source: Kaspersky Lab

Within a decade, clinicians may not be directly involved at all, only called in as experts in unexpected or critical incidents.

Exposed infrastructure

Additional research into medical perimeter assets that are open to the internet uncovered a management platform that's popular in surgical theaters. Here, the researchers found a critical vulnerability and two misconfigurations that could lead an attacker to sensitive data and treatment procedures. Since fitting and managing the functionality of neurostimulators is a surgical process, such security oversights could potentially offer attackers direct access to details of patients with connected implants, or even affect the implant surgery process.

Other potential attack vectors

Currently, all roads lead to data. Our evaluation of current cyberthreats facing connected neurostimulators found that every hypothetical exploitation scenario led inevitably to unauthorized access, either to a patient's personal data (electronically protected health information, ePHI) or to personal device settings, causing issues with a patient's confidentiality or the integrity and safety of the whole system.

Here is a list of potential scenarios that are all feasible with the current level of technical development:

- Remote tampering with implant settings to cause harm or pain to a target patient or patients (if we are dealing with network-level tampering). For example, you could manipulate the impulse settings of a neurostimulator implant inside a patient with Parkinson's disease, effectively paralyzing them for a while. Just imagine the consequences if a hacker uses this to immobilize someone driving a car.

- Depleting the power supply. Neurostimulators are currently powered by a battery and if they fall below a certain critical level, they cease to operate, and need to be surgically replaced. A hacker could drain the battery, perhaps by actively connecting the device in a loop, although they would need to be in physical proximity. Even though in newer models a magnet is required to initiate the process of pairing with the programmer, this vector cannot be ruled out.

- Extortion is another potential threat – where an adversary can seize control of, lock or threaten to sabotage a device unless you pay or do something they want you to do.

The human factor

Humans represent one of the greatest vulnerabilities. Current implants are used in busy healthcare environments, and people working in hospitals are not security experts. Our on-site research shows that they don't always change hardcoded passwords on the clinical programmers, misuse the latest generation of programmers (consumer-grade hardware such as iPads or smartphones supplied by the vendor with clinician software pre-installed), for example, to browse online or install an additional application. In other words, healthcare professionals often make the same mistakes as

ordinary users, only in this case the device also carries patient-critical software. This is really important, because any system is only as secure as its weakest part.

Though manufacturers work to minimize the risk of insecure or malicious third-party software interfering with the function of the programming software and IPG, the risk of insecure code infecting these devices is still very real. And, while users may be instructed to follow best security practices, in reality it is common for default passwords to be used and more laborious steps ignored.

Hospitals have to ensure that people working with implants and other sophisticated devices are aware of and apply basic cybersecurity hygiene.

Design and functionality

Some vulnerabilities are necessary for proper functioning of the neurostimulators – their design is constrained by factors that engineers working on most electronic devices don't have to worry about. For example, a medical implant needs to be controlled by physicians in emergency situations, including when a patient is rushed into a hospital far from their home. This precludes use of any password that isn't widely known among clinicians, but such a password would be very easy for malicious actors to bypass.

Ultimately some trade-offs will need to be made – power usage, size, shape, and a host of other limitations are key to designing implanted devices to a degree unmatched by most electronic devices.

It's important to mention here that a number of organizations are aware of the potential cybersecurity issues and are introducing measures and guidelines for the security of connected medical devices. For example, in October 2018, the U.S. Food and Drug Administration (FDA) published a draft paper on pre-market security considerations for manufacturers of medical devices. Other FDA reports feature mitigation plans for known issues in older generations of implants.

The growing recognition of cybersecurity issues and the need to address them 'at source' is a good thing. The possibility to update firmware over the air is one such example. But as the technology evolves, so do the potential threats facing it.

Signposts to future threats

In 2012 an experiment showed that observing P-300 brainwaves can leak information such as pin-codes from our brains, using commercially available headsets. If even passive access to the information recorded in the brain can make it possible to learn basic thoughts, it follows that if an attacker can gain active access and control the brain's stimulation, they can potentially manipulate brainwaves in order to 'control' behavior and feelings – and, ultimately, memories.

Future risk predictions

Within five years scientists expect to be able to electronically record the brain signals that build memories and then enhance or even rewrite them before putting them back into the brain. A decade from now, the first commercial memory boosting implants could appear on the market – and, within 20 years or so, the technology could be advanced enough to allow for extensive control over memories.

The healthcare benefits of all this will be significant, and this goal is helping to fund and drive research and development. However, as with other advanced bio-connected technologies, once the technology exists it will also be vulnerable to commercialization, exploitation and abuse.

New threats could include the mass manipulation of groups through implanted or erased memories of political events or conflicts; while 'repurposed' cyberthreats could target new opportunities for cyber-espionage or the theft, deletion of or 'locking' of memories (for example, in return for a ransom).

Attackers will take advantage of the fact that addressing the complex and unpredictable cybersecurity needs of this rapidly evolving technology will not have the same priority as realizing the scientific or medical potential, and that even when taken into account it is likely to be implemented by people who lack professional cybersecurity understanding. The first and most important step, therefore, is to bring healthcare professionals and the cybersecurity industry closer together.

Conclusion

Current vulnerabilities matter because the technology that exists today is the foundation for what will exist in the future. Although no attacks targeting neurostimulators have been observed in the wild – a fact that is not altogether surprising since the numbers currently in use worldwide are low, and many are implemented in controlled research settings, several points of weakness exist that will not be hard to exploit.

When it comes to future security, there are two things to bear in mind. The first is that many of the potential vulnerabilities could be reduced or even eliminated by appropriate security education for clinical care teams and patients. The second thing is that patient needs will always take precedence, which means that compromises will inevitably need to be made. Security is one of many goals, and it is impossible and unreasonable to expect medical professionals to be cybersecurity experts.

But reducing the vulnerabilities posed by these implants can and should be a key goal for all supporting stakeholders – the hardware and software manufacturers and vendors, the cybersecurity industry, and professional bodies. Collaborating to understand and address emerging risks and vulnerabilities, and doing so while this technology is still relatively new, will pay off in the future.