

SAMSUNG

Lewis Rhodes Labs Case Study

A Giant Leap Forward in Data Search is about to Change Everything from Cybersecurity to eCommerce

The Human Brain Reveals How to Search Data Faster



Businesses and organizations today are accumulating more data than ever, and this sea of big data presents both tremendous opportunity and great risk. Until very recently, the cost of searching through vast amounts of raw data has been prohibitively expensive and time-consuming. This has made effective monetization of data a challenge for most businesses, and it has meant that recovery from malicious data breaches can cost many millions of dollars and take excessive amounts of time. But that is about to change.

The High Cost of Data Breaches

In recent years, there has been no shortage of examples of high-profile data breaches. In 2017, for example, a breach at Equifax exposed the data of 143 million customers – over 40 percent of the U.S. population. The data accessed included names, Social Security numbers, addresses, and even driver's licenses. The resulting cleanup and investigations cost nearly \$1.4 billion and took over two years. While this example is extreme, costly and damaging breaches are far too common.

According to IBM's 2020 Cost of a Data Breach Report, it takes an average of 315 days just to identify and contain a malicious breach, and in the U.S., breaches cost businesses an average of \$8.64 million. But what makes breach recovery so expensive, and why does the process take so long? The core problem is that, due to the limitations of existing technology, search on data doesn't scale with the volume of data. Or to put it another way, the more data you have, the longer it takes to search through it to track down malicious code and assess the damage. There simply has been no effective way to speed up this costly process – until now.

\$8.64^M

United States has the
highest country average cost

Introducing NPUs^{search}™ by Lewis Rhodes Labs

The introduction of Lewis Rhodes Labs' revolutionary new NPUs^{search} promises to dramatically speed up search on data, greatly reducing the time and expense of breach recovery. NPUs^{search} is built with Samsung's SmartSSD Computational Storage Drives (CSDs), which are powered by Xilinx FPGAs. The result is an appliance that searches vast amounts of data without the compromise and expense of indexing.

NPUs^{search} leverages concepts derived from research into the functioning of the human brain. The core of this technology is the Neuromorphic Processing Unit (NPU), which was modeled using select features of the human brain. Dr. Pamela Follett, co-Founder and COO of Lewis Rhodes Labs, has worked as a pediatric neurologist, neuroscientist, and biomedical engineer specializing in disorders of early brain development. Through her research, Dr. Follett gained insight on how the brain makes high-speed computations utilizing low levels of energy. Seeing an opportunity to apply this understanding to data search, she developed NPUs^{search} with the team at Lewis Rhodes Labs. The novel pattern matching capability of NPUs^{search} is an entirely new methodology for rapid search of unindexed data at low power levels.

Benefits of NPUs^{search}™

- Requires no indexing
- Rapid, predictable search times
- Decreased network, CPU, and I/O demands
- Fully scalable performance



NPUs^{search}™ Appliances
Search-In-Storage

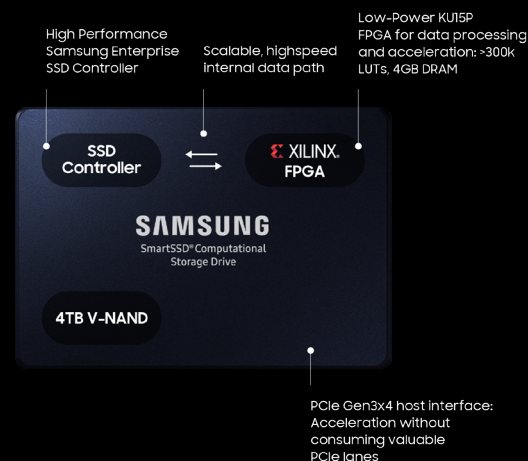
Samsung SmartSSD® Computational Storage Drive

Lewis Rhodes Labs selected the Samsung SmartSSD CSD, powered by the Xilinx Adaptive Platform as the programmable hardware for their NPUs search appliance. The SmartSSD CSD accelerates data-intensive applications by bringing FPGA compute resources directly inside the storage drive.

Benefits of SmartSSD® CSD

- Technology leadership from Samsung and Xilinx
- Leading NAND and proven SSD controller
- Fast and flexible acceleration engine
- Scalable performance
- PCIe lane and space saving


The SmartSSD CSD allows Lewis Rhodes Labs' NPU to scan data at a constant rate and to fully analyze data at the full read speed of the drive. This means that, when investigating a data breach, users can get answers to search questions in minutes rather than days – orders of magnitude faster than existing technologies, a true game-changing development.



Fully Scalable, Deterministic, High Resolution Search

In addition to its much faster data search speeds, the NPU supports fully scalable search, and, unlike traditional methods, completes a deterministic analysis of data. With previous technologies, all data is transferred to the CPU, where it is then searched. This is an incredibly slow, inefficient, and unpredictable process that can take days or more. With NPUs search, instead of transferring the full content of each drive, only search results are returned to the CPU.

Dr. Follett gives this analogy: "When you are searching data, it's like looking for lost keys in an 18-room house." With previous technologies, the contents of each "room" had to be gathered and searched consecutively. However, NPUs search effectively "puts a member of a search party in every room of the house." The NPU sits in storage, where the data is, on each SmartSSD CSD. Because searches are completed locally on the SSD, each SmartSSD CSD can be searched simultaneously, with every search taking the same predetermined amount of time. Whether searching 96 terabytes or 10 petabytes, the NPUs search can complete its search in about 25 minutes.



The advantages of *NPUSearch* extend beyond the savings in both the time and expense of responding to a data breach. For years, attackers have enjoyed an asymmetric advantage in the cyber world because they can move much more quickly than defenders can respond. Slower search times for some types of scans or data sets allow attackers to stay hidden for longer. *NPUSearch* levels the playing field.

Because *NPUSearch* times are deterministic – all scans complete in the same predictable amount of time – a potential attacker’s ability to exploit the weakness of prolonged search times is eliminated. In this way, the effective, consistently rapid, performance offered by *NPUSearch* can reduce the effectiveness of some attack tactics.

In terms of implementation, *NPUSearch* is designed to be user-friendly. The *NPUSearch* lies behind a neuromorphic driver and compiler that fully translates and integrates the technology into a familiar interface, meaning that the novel technology is invisible to the user. There is no complex software to learn, and users can conduct regular regular expression searches through a standard Python API.



Development Partners

NPUSearch was developed by Lewis Rhodes Labs in conjunction with Samsung and Xilinx. According to Dave Follett, co-Founder and CEO of Lewis Rhodes Labs, “Support from Xilinx and Samsung was critical in developing and building *NPUSearch*™.” Follett added, “LRL needed to move fast, and support from Xilinx and Samsung was essential to our success. Xilinx tools reduced the time LRL needed to deliver this product, and we received incredible support throughout the process.”

Value Beyond Breaches

As much as *NPUSearch* is a game-changer in terms of the cleanup and investigation of data breaches, this use case is just the tip of the iceberg. *NPUSearch*’s search speed on raw data also enables companies and organizations to more fully realize the potential of their ever-growing pools of data. For instance, data can be analyzed and insights sold to third parties. And data-based optimization can reveal novel insights that can improve your own business’s performance. The opportunities to exploit and monetize untapped data resources are endless for businesses, and now, thanks to *NPUSearch*, these opportunities are easier than ever to take advantage of.



SAMSUNG