SELECTING ENCRYPTION FOR "DATA AT REST" IN BACK-END SYSTEMS: WHAT RISKS ARE YOU TRYING TO ADDRESS?

June, 2015

→ **Derek E. Brink,** CISSP, Vice President
and Research Fellow, IT Security and IT GRC

Report Highlights

As tempting as it may be to simply use a security technology that has worked in one area as the solution for another, the right approach is to carefully consider what risks you are trying to address, and to choose the technologies that are most appropriate to address them. Aberdeen Group's analysis shows why *full-disk encryption* has become an attractive solution for the risks to data in use at the endpoints — and why *file-level encryption* is a better fit for the risks associated with data at rest in back-end systems.

ABERDEEN GROUP

2

Over the past several years, Aberdeen's research has reflected the sharp growth in enterprise use of encryption for *data at rest* in back-end systems, *data in motion* on the network, and *data in use* at the endpoints. But the *purpose* of using encryption — i.e., the risks that organizations intend for encryption to help them address — can be markedly different for each. As always, decisions about security have to be made in a specific business context.

Traditionally, discussions about **encryption** are segmented into three distinct data protection problems:

- *Data at rest* in back-end systems (e.g., file servers, network storage, cloud-based storage, databases, backup media)
- *Data in motion* on the network
- *Data in use* at the endpoints (e.g., laptops, mobile devices, removable media)

Over the past several years, Aberdeen's research has reflected the sharp growth in enterprise use of encryption for all three of these areas. But the *purpose* for using encryption — i.e., the **risks** that organizations intend for encryption to help them address — can be markedly different for each. As always, decisions about security have to be made in a specific *business context*.

For "Data in Use" at the Endpoints, the Biggest Risk to Address is that Endpoints are Commonly Lost, Stolen or Missing

For example, consider the use case of protecting data in use at the endpoints. In this scenario, Aberdeen's research shows that **file / folder encryption** has increasingly been overtaken by **full-disk encryption (FDE)** — where FDE includes both *software-based* solutions and *hardware-based* solutions, including *self-encrypting drives* — in recent years, as shown in Figure 1.

Why is this? As far back as 2009, Aberdeen's research has found that faced with a choice between the *precision* of encrypting only specific files or folders, based on decisions about each piece of content and pre-existing policies — or the *simplicity* of just encrypting everything on the hard drive — simplicity is consistently correlated with the companies achieving the best results. In other words, why worry about whether users or

3

technologies can consistently make the right decisions about which data to encrypt, when you can just encrypt everything?

Figure 1: Aberdeen's Research Reflects the Sharp Growth in Enterprise Use of Encryption, on Both Endpoints and Servers

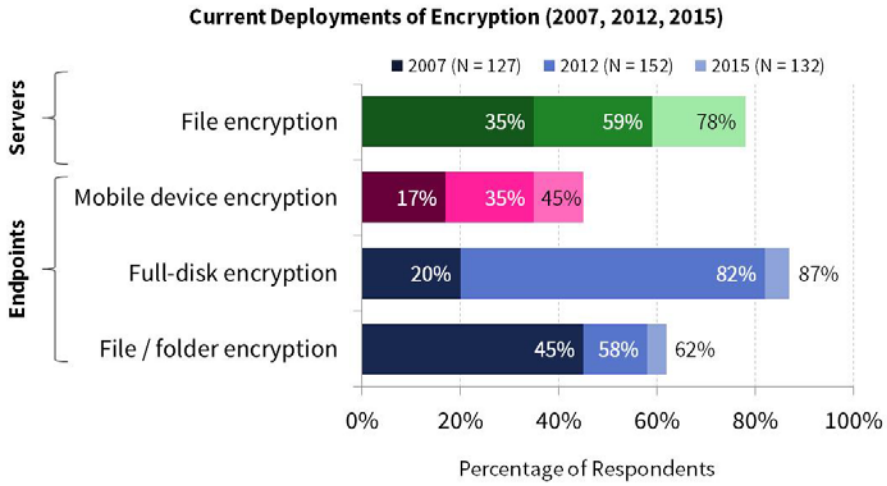

Source: Aberdeen Group, June 2015

What's important to keep in mind, however, is the primary problem that encryption for data in use at the endpoints is aiming to solve — which is that laptops and mobile devices that contain sensitive enterprise data are very commonly *lost*, *stolen* or simply *unaccounted for*. In other words:

- ➔ The risk of **loss or exposure of sensitive data** in use at the endpoints is high.
- ➔ The use of full-disk encryption provides a **higher level of assurance** than file / folder encryption that the desired data protection is actually in place, if the endpoint becomes lost, stolen or missing (as an aside, companies should be sure to confirm that the FDE solution they deploy is activated for endpoints that are in *sleep* or *hibernate* mode).

4

For data in use at the endpoints, the primary risk that full-disk encryption is meant to address is that laptops and mobile devices containing sensitive enterprise data are very commonly lost, stolen or simply unaccounted for. But the risks that need to be addressed for data at rest in back-end systems are different!

➔ FDE solutions require **no involvement or decisions** by the organization's users about whether or not to encrypt, and have evolved to have **little to no impact** on endpoint performance or on the user's operational experience.

When used in conjunction with a *hardware root of trust* (e.g., a *trusted platform module*), FDE solutions can also provide the additional benefit of assuring that the endpoint and its platform-level software boots up in a known, unaltered and trusted state — i.e., free of *rootkits* or other malware.

For "Data at Rest" in Back-End Systems, the Risks are Different!

In contrast, consider the use case of protecting data at rest in back-end systems, e.g., on *file servers*, *network servers* or in *cloud-based storage*. In this scenario, Aberdeen's research again shows steady growth in the use of encryption: nearly 4 out of 5 (78%) respondents indicated current use of **file-level encryption** in 2015, up from about 1 out of 3 (35%) back in 2007 (refer again to Figure 1).

Again, why is this? Why, for example, would the simplicity of the full-disk encryption, "let's just encrypt everything and be done with it" approach not apply equally well for data at rest on file servers and in cloud-based storage, as it seems to do for data in use at the endpoints?

Because the risks that need to be addressed are different!

To be sure, leading providers for **network-attached storage (NAS)** solutions and **storage-area network (SAN)** solutions have begun offering options for full-disk encryption — typically at a significant price premium over their non-encrypted offerings, and with potential implications for the flexibility to move between competing vendors in the future. But the risk of your NAS or SAN implementation being lost, getting stolen or

going missing is obviously not as high as it is for the laptops, tablets and smartphones that are being used throughout your extended enterprise. So on that basis, the case for the simplicity of FDE just isn't there when it comes to selecting encryption for data at rest in back-end systems. (If your concern is for storage that gets decommissioned, then like any encryption solution, FDE would help to ensure that the organization's data is not accessible.)

In addition, we can reason that whenever the file servers, network storage or cloud-based storage that contain our organization's sensitive data are online, available and accessible — which is to say, virtually all of the time — then full-disk encryption is not actively protecting us. On the contrary, the biggest risks to our sensitive data are more likely to be infiltration and **unauthorized access** by external attackers, **fraud or theft** by trusted insiders, or non-malicious **errors** made by authorized, well-intended users. It stands to reason that when selecting encryption for data at rest in back-end systems, these are the biggest risks that need to be addressed. **File-level encryption**, which is actively protecting the data in any of these scenarios, stands out as the more appropriate choice.

When it comes to selecting encryption for data at rest in back-end systems, the case for the simplicity of full-disk encryption just isn't there. To address the risks to sensitive data in file servers, network storage and cloud-based storage, file-level encryption stands out as the more appropriate choice.

What Encryption to Select to Protect Data at Rest in Back-End Systems Depends on What Risks You Are Trying to Address

When we're talking about **risk**, we have to use the proper definition. That is, for any given undesirable event that could potentially happen, we have to consider both the *likelihood* that it might occur, and the *business impact* if it actually does occur. If we're not talking about both likelihood and business impact, we're not really talking about risk! Too much of the time, technically oriented IT and Security teams focus intensely on identifying the undesirable events — i.e., by talking about the latest *threats*, *vulnerabilities* and *exploits* — but not necessarily

framing them properly in terms of risk, and in the business context that's relevant to their organization.

From the likelihood perspective, it's only logical that attackers would find enterprise data at rest in back-end systems to be an attractive target — as in the infamous quip about robbing banks "because that's where the money is" — and this is true whether they are motivated by financial gain (i.e., *cyber criminals*) or by political causes (i.e., *hacktivists*). As Aberdeen noted in *SOS! (Secure Our Servers)* (August 2014), industry sources also echo this point:

- "Attackers are increasingly going after central, strategic targets as a means to optimize their efforts and increase their return on exploit." (Source: IBM, *X-Force Threat Intelligence Quarterly*)
- "Servers have typically been on top [in a breakdown of data breaches by asset category], probably because attackers know that's where the data is stored." (Source: Verizon, *Data Breach Investigations Report*)

So the big question becomes this: Do we have any hard data to support the assertion that *external attacks*, *insider threats* and *human error* are the most likely risks for sensitive data at rest in back-end systems? Happily, the investigation and analysis of tens of thousands of actual data breach incidents over the course of a decade by Verizon and its investigation partners in the *DBIR* series provides us with invaluable evidence for our risk-based analysis.

Based on roughly 2,500 actual data breach incident investigations that have been made available for sharing in the VERIS Community Database, the high-level categories for the **actions** that were taken on server and endpoint **assets**, respectively, are summarized in Table 1.

Table 1: A2 Grid — Actions on Assets for Data Breach Incidents Provide Insights into the Biggest Risks to be Addressed

	Hacking	Misuse	Error	Malware	Unknown	Social	Physical	Environmental
Servers	**610**	**241**	**162**	53	46	40	33	4
Endpoints	27	30	**73**	**36**	9	12	**409**	0

Source: *VERIS Community Database*; adapted by Aberdeen Group, June 2015

In the VERIS framework, the high-level categories for actions include:

- **Hacking** – i.e., intentionally circumventing security controls to access or harm an asset without authorization
- **Misuse** – i.e., the inappropriate or malicious use of assets by an actor with authorized access
- **Error** – i.e., an action that causes or significantly contributes to an incident, and which deviates from the normal processes of the organization
- **Malware** – i.e., malicious software that makes an unauthorized change in the state or function of an asset
- **Social** – i.e., exploiting human behaviors as the means to install malware, gain information or access, etc.
- **Physical** – i.e., deliberate actions on assets that involve possession, proximity or force
- **Environmental** – i.e., incidents caused by natural events (e.g., earthquakes, floods) or other hazards (e.g., power,

Definitions

Vocabulary for Event Recording and Incident Sharing (VERIS) provides a common framework that describes the "4 A's" of a data breach incident:

- **Actors** – Whose actions affected the asset?
- **Actions** – What actions affected the asset?
- **Assets** – Which assets were affected?
- **Attributes** – How was the asset affected?

The **A2 Grid** is a simplified view of the VERIS data, focused on the *actions* that were taken on specific *assets*.

Source: Verizon DBIR

water, temperature) associated with the environment in which assets are located

And there we have the evidence we are looking for — based on the data in Table 1:

- For *data in use* on laptops and mobile devices, more than 80% of the actual data breach incidents were related to *physical loss or theft* and *human error* — which aligns with the strong adoption of **full-disk encryption** for this use case, as observed in Aberdeen's research. The use of full-disk encryption provides a higher level of assurance than file / folder encryption that the desired data protection is actually in place, requires no involvement or decisions by the organization's users about whether or not to encrypt, and has little to no impact on endpoint performance or user experience.

- For *data at rest* in file servers, network storage and cloud-based storage, more than 85% of the actual data breach incidents were related to external threats (*hacking*), insider threats (*misuse*) or *human error* — which aligns with the strong adoption of **file-level encryption** for this use case, as observed in Aberdeen's research. File-level encryption is actively protecting the organization's data, whenever these back-end systems are online, available and accessible — even if unauthorized access to these systems has been achieved by any of these actions.

In a time when an overabundance of options for information security technologies and controls can make it painfully difficult for the security team in any given organization to sort through all the alternatives — and to make the necessary choices for the mix of controls that represents the best fit for their specific business context — it can be tempting to seek a shortcut, and simply use a

technology that worked in one area as the solution for another. But as we have seen, this is not necessarily the right approach.

Selecting an Encryption Solution: Use the Right Tool for the Job

Organizations that are concerned about protecting their sensitive data should carefully consider what risks they are trying to address, and choose the appropriate security technologies to address them. In this analysis, we have seen why *full-disk encryption* has become an attractive solution for the risks to data in use at the endpoints — and why *file-level encryption* is a better fit for the risks associated with data at rest in back-end systems.

Summary and Key Takeaways

- Traditionally, discussions about encryption are segmented into three distinct data protection problems: *data at rest* in back-end systems (e.g., file servers, network storage, cloud-based storage, databases, backup media); *data in motion* on the network; and *data in use* at the endpoints (e.g., laptops, mobile devices, removable media). Over the past several years, Aberdeen's research has reflected the sharp growth in enterprise use of **encryption** for all three of these areas.

- But the *purpose* of using encryption — i.e., the **risks** that organizations intend for encryption to help them address — can be markedly different for each of these problems. As always, decisions about security have to be made in a specific *business context*.

- For the use case of protecting data in use at the endpoints, Aberdeen's research shows that **file / folder encryption** has increasingly been overtaken by **full-disk encryption (FDE)**, for two main reasons: first, the *simplicity* of just encrypting everything on the hard drive;

and second, the fact that in this scenario lost, stolen or missing endpoints represent the greatest *risk*.

➔ In contrast, for the use case of protecting data at rest in back-end systems — e.g., on *file servers*, *network servers*, or in *cloud-based storage* — Aberdeen's research shows steady growth in the use of **file-level encryption**, with nearly 4 out of 5 (78%) respondents indicating current use in 2015, up from about 1 out of 3 (35%) back in 2007.

➔ Why this difference? Because **the risks that need to be addressed are different**. When it comes to selecting encryption for data at rest in file servers, network storage and cloud-based storage, the case for the simplicity of full-disk encryption just isn't there.

➔ Qualitatively, we can reason that whenever the file servers, network storage or cloud-based storage that contain our organization's sensitive data are available and accessible — which is to say, virtually all of the time — then full-disk encryption is not actively protecting us. On the contrary, the biggest risks to our sensitive data are more likely to be infiltration and **unauthorized access** by external attackers, **fraud or theft** by trusted insiders, or non-malicious **errors** made by authorized, well-intended users. When selecting encryption for data at rest in back-end systems, these are the biggest risks that need to be addressed — and which *are* addressed by the selection of file-level encryption.

➔ Quantitatively, an analysis of the roughly 2,500 actual data breach incident investigations that have been made available for sharing in the VERIS Community Database provides us with hard evidence and corroboration for our risk-based analysis:

More on Cloud-Based Storage

Logically, encryption solutions for cloud-based storage that are implemented at the *volume* level are the equivalent of "FDE in the cloud" — i.e., once the encrypted volume is in use, the data is unencrypted and in the clear. In the context of cloud-based storage, file-level encryption solutions continue to provide protection through access controls and monitoring, even when another user or process is accessing the data.

- For *data in use* on laptops and mobile devices, more than 80% of the actual data breach incidents were related to *physical loss or theft* and *human error* — which aligns with the strong adoption of **full-disk encryption** for this use case, as seen in Aberdeen's research. The use of full-disk encryption provides a *higher level of assurance* than file / folder encryption that the desired data protection is actually in place, requires *no involvement or decisions* by the organization's users about whether or not to encrypt, and has *little to no impact* on endpoint performance or user experience.

- For *data at rest* in file servers, network storage and cloud-based storage, more than 85% of the actual data breach incidents were related to external threats (*hacking*), insider threats (*misuse*), or *human error* — which aligns with the strong adoption of **file-level encryption** for this use case, as observed in Aberdeen's research. File-level encryption is actively protecting the organization's data, whenever these back-end systems are online, available and accessible — even if unauthorized access to these systems has been achieved by any of these most-likely actions.

➔ As tempting as it may be to simply use a security technology that has worked in one area as the solution for another, the right approach is to carefully consider what risks you are trying to address, and to choose the technologies that are most appropriate to address them. Aberdeen's analysis shows why full-disk encryption has become an attractive solution for the risks to data in use

at the endpoints — and why *file-level encryption* is a better fit for the risks associated with data at rest in back-end systems.

For more information on this or other research topics, please visit www.aberdeen.com.

Related Research

A Renaissance in Enterprise Rights Management; May 2015
Flash Forward: Putting "Critical Security Controls" in Perspective; January 2015
Insider Threat: Three Activities to Worry About, Five Ways They're Allowed to Happen - and What Enterprises Can Do About It; September 2014
Securing the Evolving Datacenter; April 2014
Understanding Your Encryption Footprint: Your Reliance on Security and Trust; June 2013
SOS! (Secure Our Servers): Managing Servers in the Evolving Enterprise Datacenter; August 2014

Author: Derek E. Brink, CISSP, Vice President and Research Fellow, IT Security and IT GRC (Derek.Brink@aberdeen.com)

About Aberdeen Group

Since 1988, Aberdeen Group has published research that helps businesses worldwide to improve their performance. Our analysts derive fact-based, vendor-neutral insights from a proprietary analytical framework, which identifies Best-in-Class organizations from primary research conducted with industry practitioners. The resulting research content is used by hundreds of thousands of business professionals to drive smarter decision-making and improve business strategies. Aberdeen Group is headquartered in Boston, Massachusetts, USA.