E-BOOK

Combatting Ransomware by Exposing the Myths

By: Brien M. Posey

NASUNI®

Table of Contents

Ransomware is now one of the leading causes of data loss in the enterprise. Organizations go to great lengths to prevent ransomware attacks and to ensure data is recoverable if an attack should occur. Yet ransomware remains an ever-present threat. Check Point listed ransomware among its biggest cyber security threat. Similarly, a recent report by Splunk found that 84% of organizations have suffered a significant security incident in the last two years, with 31% of respondents reporting a ransomware attack.

Electronics manufacturer Canon suffered a ransomware attack that compromised 24 of its domains. The notorious Colonial Pipeline attack resulted in a major fuel shortage across much of the United States and a $4.4 million ransom payment to the attackers. Other high-profile victims of ransomware attacks include Acer, Quanta, The National Basketball Association (NBA), and Kia Motors.

Ransomware continues to pose a credible and costly threat to the enterprise, and misinformation only compounds the problem. Today, several ransomware-related myths remain prevalent. These could prevent companies from developing solid defense and recovery plans.

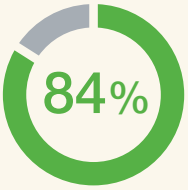

Organizations have suffered a **significant security incident** in the last two years

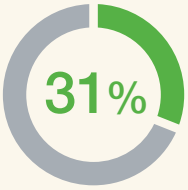

Respondents reported a ransomware attack

Myth 1: It Can't Happen in the Cloud

The first of these myths is that a ransomware infection cannot occur in the cloud. This myth likely stems from the fact that early on, ransomware was designed to attack the system upon which the infection occurred but lacked the sophistication to attack connected systems. Today, however, ransomware strains can attack external resources such as mapped network drives and even data within cloud-based file shares.

While ransomware has become more capable and increasingly sophisticated in recent years, this vulnerability is also a side effect of the way enterprise computing has changed.

A key trend of the 2020s is location-agnostic computing. Users are no longer working from domain-joined desktops in the corporate office that are tied to backend systems in the organization's datacenter. Instead, the pandemic has forced nearly all of the world's knowledge workers to work remotely, usually from home, where they are likely to be using untrusted personal devices. Similarly, the systems users connect to in order to do their jobs are widely distributed. While some data may remain in an organization's datacenter, an organization may also have data stored in SaaS applications like Microsoft 365 and in cloud-based file shares such as Azure Files. Never mind all the structured data an enterprise-class organization relies on.

Since users are now remote and decentralized, so too is most of the data those users rely upon. A user's physical location no longer matters, nor does the location of the data. Modern computing systems have been designed to work seamlessly, regardless of where the various resources are physically located. Just as a user who is working from home on a personal device can open a file stored in Azure Files, a ransomware infection triggered on that user's device may be able to access and encrypt that same data.

Myth 2: All You Have to Do is Restore a Backup

A second prevalent ransomware myth is that recovering from a ransomware attack is simply a matter of restoring a backup. Like all good myths, it includes a kernel of truth. Backups are one way of restoring data from ransomware attacks, but they can result in extended recovery times. Often, a company's backup solution simply won't get the business up and running quickly enough. This can lead to major disruptions in productivity.

To furher address this myth, here are four reasons why a backup does not completely resolve a ransomware attack.

Why Backup Fails #1: Time Bombs

Early on, ransomware behaved in a very specific way. As soon as a user triggered the infection, the ransomware would immediately begin encrypting files. The ransom demand would only be displayed once the encryption process was finished.That way, the organization would not realize a ransomware attack had occurred until the damage had already been done.

Although this has long been standard operating procedure for ransomware, some ransomware authors have begun taking a completely different approach. Time-bomb-enabled ransomware initially infects files without encrypting them. This infection is generally tied to a clock and causes infected files to become encrypted weeks or even months after the initial infection.

This technique can neutralize an organization's backup strategy. Most organizations retain data backups for a matter of days, or perhaps a few weeks at the very most. Older backups would contain outdated data that is of little use to the organization, so it makes little sense for an organization to hold on to backups whose age exceeds the useful life of the data. Herein lies the problem.

Imagine for a moment that an organization's file data becomes infected with time-bomb-enabled ransomware in January. Let's also suppose this particular organization has a backup retention period of three months, but the ransomware has an activation period of six months. This would mean the data encryption would occur in June, six months after the initial infection.

The organization's natural response would be to restore a backup. The problem is that the files will have been infected for six months. Because the organization's backup retention period is only three months, all of the backups will therefore contain infected files. If the organization attempts to restore one of these backups, the restored data will immediately become encrypted.

In some cases, it may be possible to overcome a time-bomb-enabled ransomware attack by leveraging various clock manipulation techniques. Even so, it can be extremely difficult for an organization to get its data back following this type of attack. The organization probably will not know exactly when the infection happened or how long the ransomware was designed to remain dormant.

Why Backup Fails #2: Ransomware Targets Backups

Another reason why backups are not an ideal solution is that some types of ransomware are specifically designed to seek out and attack backups. Ransomware authors have long known that once they have encrypted a victim's data, the victim only has two options for getting their data back - pay the ransom or restore a backup. In other words, an organization's backups are the number one thing standing in the way of a ransom demand being paid. As such, ransomware authors are always looking for new and creative ways to eliminate a victim's backups as a viable option for data recovery.

"Performing a full restoration of a large data volume takes time, and **this data will not be accessible to end users until the recovery process completes.**"

Why Backup Fails #3: Files are Encrypted at Different Times

One of the other major problems associated with using a backup to recover from a ransomware attack is that the encryption process does not happen instantaneously. Depending on the volume of data being encrypted, it can take hours or even days for the ransomware to encrypt everything. The time required to complete the encryption process matters because restoring a backup can lead to data loss.

Imagine for a moment that a ransomware attack begins encrypting data at 9 AM on a Monday morning. Now suppose a user creates an important file at 10 AM. Let's also assume that because the file was created after the ransomware attack was underway, the ransomware is unaware of the file's existence and therefore does not encrypt the file. Because the attack happened during the workday, there are likely many files being created as the attack is ongoing, but let's focus on this one file.

The IT department will almost certainly be able to tell that the ransomware infection began at 9 AM. Consequently, they may choose to restore a backup from 8:45 AM. Assuming this backup overwrites the volume, any data that has been created or legitimately modified since 8:45 AM will be lost. This includes the important file that was created at 10 AM.

The encryption process can take hours or days to complete depending upon how much data needs to be encrypted. So, there is a good chance the IT department will begin the restoration process before the attack even finishes encrypting all of the organization's data. While this will put a stop to the attack, it also means users will lose all work they have done since 8:45 AM, even if the file they were working on was never encrypted.

Why Backup Fails # 4: It Takes Too Long to Recover

Finally, it takes time to restore a backup. When organizations design their backup strategy, they frequently look at two important metrics – the Recovery Point Objective (RPO) and the Recovery Time Objective (RTO). The RPO essentially corresponds to the frequency with which data is protected, while the RTO has to do with how long a recovery could potentially take. Although organizations would ideally like to have an RTO of zero minutes, performing a full restoration of a large data volume takes time, and this data will not be accessible to end users until the recovery process completes. While the operation may result in the recovery of most of the organization's data, users will be unable to do their jobs until after the process has completed. These extended downtimes can be extremely damaging to businesses.

80% of organizations that paid a ransom were **hit by a second attack**, and half of those were hit by the same ransomware group

Myth 3: Paying the Ransom Guarantees Data Recovery

Given the complexities of restoring a backup in response to a ransomware attack, as well as the potential for data loss that is directly tied to the restoration process, an organization might understandably be tempted to pay the ransom. Unfortunately, paying the ransom does not guarantee data recovery.

According to an article in Forbes, a staggering 92% of those who pay the ransom do not get their data back. This is especially troubling when you consider that the average ransom demand in 2021 was $570,000 with some demands being far larger. The perpetrators of the Colonial Pipeline attack, for example, asked for over $4 million.

If an organization does decide to pay a ransom, it may find that the ransomware demands an additional payment. The simple fact that the organization has paid a sum of money to get its data back means that the data is valuable, and the victim has few, if any, options for recovering that data. As such, the ransomware author might exploit this vulnerability and ask for more money.

There have been reports of data being unlocked following a ransomware payment, then being re-encrypted shortly thereafter. In fact, 80% of organizations that paid a ransom were hit by a second attack, and half of those were hit by the same ransomware group. Additional malware might even be left behind after the data has been decrypted, leading to future infections.

Now that some of the myths have been addressed, it is important to have the right plan in place to handle a ransomware attack should one occur.

Recovering from an Attack

One of the most effective steps organizations can take to ensure their ability to recover file data from a ransomware attack is to migrate file data to Nasuni and Microsoft Azure Object storage. Nasuni continuously snapshots file data to Azure. It is important to note, these snapshots differ from backups (and even from other types of snapshots) in some very significant ways that collectively provide a solid defense against ransomware.

Here are a few highlights of how Nasuni recovers file shares from a ransomware attack:

Immutable Snapshots:
First, Nasuni's snapshots are immutable. This means that even if ransomware tried to attack the snapshots, it would be unable to alter or delete them. As such, snapshots will always be available for use following a ransomware attack, and organizations do not have to worry that snapshots may have been compromised by the attack.

Surgical Recovery:
The Nasuni software leverages a versioning file system that is designed to retain all previous versions of a file. Any time a file is modified, the Nasuni software saves those deltas and versions them to the cloud, but without overwriting the original file. This makes it possible to easily roll the file back to a previous state. While there are any number of tools that can make it possible to revert a file to an earlier version, Nasuni allows an organization to restore only the files that have been damaged.

There are two reasons why identifying and then recovering only the damaged files is so important. First, it accelerates the recovery because no unnecessary operations are being performed. Second, it helps an organization avoid the data loss that could potentially occur if it were to revert an undamaged file to a previous version as a part of its recovery efforts. In other words, Nasuni allows file recovery operations to be completed with surgical precision.

Rapid Restore:
Even if an organization finds that a large amount of data has been encrypted by a ransomware attack, the recovery operation can be completed far more quickly than would be the case if the organization were to restore a traditional backup. The recovery process involves moving snapshot pointers, not restoring terabytes of data from a backup. Hence, organizations often find that they can restore data extremely rapidly. Of Nasuni customers who have been impacted by a ransomware attack, *81% recovered in less than 24 hours, and 36% recovered within an hour.*

Infinite Snapshots:
Another important differentiator between Nasuni's snapshots and a traditional backup is that Nasuni allows an infinite number of snapshots to be created and retained. If an organization were to be attacked by time-bombed ransomware, it would still be able to recover from the attack because snapshots pre-dating the infection would continue to exist in an unaltered state. Hence, the organization could simply roll back files to their pre-infection state.

Of Nasuni customers who have been impacted by a ransomware attack,
81% recovered in less than 24 hours, and 36% recovered within an hour.

3-2-1 Protection

3 Copies

2 Locations

1 Offsite

Conclusion

Nasuni and Microsoft collectively provide organizations with industry standard 3-2-1 protection (3 copies, 2 locations, 1 offsite) for Azure data by leveraging an immutable global file system that cannot be corrupted by ransomware. If an organization is attacked by ransomware, it will be able to roll the infected files back to an unencrypted state quickly and easily.

In many organizations, as much as 80% of company data is stored on file shares. This joint solution from Nasuni and Azure should be part of an organization's overall ransomware protection investment as it ensures business downtime is minimized, and productivity can be restored swiftly when it comes to managing, accessing, and sharing files.

Learn more about Nasuni and Azure ransomware solutions:

Nasuni Disaster Recovery and Ransomware Mitigation

Azure Defenses for Ransomware Attack

NASUNI.

ABOUT NASUNI, CORPORATION

Nasuni provides modern cloud file storage, powered by the world's only cloud-native global file system. Nasuni is a cloud replacement for traditional network attached storage (NAS) and file server silos, consolidating file data in instantly expandable cloud object storage at a fraction of the cost. Nasuni also eliminates the need for complex legacy backup and disaster recovery infrastructure, dramatically simplifying IT administration. Nasuni is headquartered in Boston, Massachusetts, USA. For more information, visit www.nasuni.com.

NORTH AMERICA HEADQUARTERS
Boston, MA | +1 (857)444.8500

UNITED KINGDOM
London, England | +44 20 3695 7895

AUSTRALIA
Sydney, Australia | +61 2 4062 9333