Quantum

WHITE PAPER

WHAT IS GDPR

and Does It Affect Your Backup and Archive?

CONTENTS

WHAT IS GDPR?

The General Data Protection Regulation (GDPR) is a new privacy regulation across the European Union (EU) that takes effect on May 25, 2018. It provides EU individuals ("data subjects") with more control over their personal data, ensures transparency about the use of data, and requires security and controls to protect data.

The GDPR specifies the roles, processes, and technologies organizations (whether EU based or not) must have in place to collect (with consent), protect, and appropriately use EU residents' personal data. IT teams shouldn't fall into the trap of considering GDPR to be solely a legislative exercise and therefore assume the effort to implement the changes should be run by the legal department (if you have one).

WHY GDPR?

Today, gathering, processing, and exchanging personal data are daily activities for many organizations; a direct outcome of this situation is the growing economic value for the data availability and protection.

Recent IT failures, such as the one suffered by British Airways in May 2017 that affected 75,000 passengers or data leakage like the one suffered by Deloitte in March 2017[1] illustrate how difficult it is to deliver 100% data protection and availability. The GDPR is about building trust in the digital economy—it does not forbid anything, but regulates everything.

GDPR benefits and opportunities:

- **Simplification:** With only one regulation for all EU countries, your organization will need to interact only with the data protection authority in the Member State you selected. You will no longer have to administer regulations in multiple EU countries.
- **Harmonization:** The GDPR provides a level of harmonization across 28 Member States (ignoring Brexit for now), meaning only one set of rules no matter where you are based.
- **Cleaning opportunity:** The GDPR requires a top-down approach with board-level sponsorship. It offers a good opportunity to review your data flows not only for compliance, but also for business efficiency. Does all your data deserve to be backed up? It is the right time to start a new archiving process—or to introduce an automatic storage tiering mechanism?
- **Security enhancement:** The GDPR clearly recommends different techniques, such as encryption, anonymization, and pseudonymization, that will help IT teams improve security.

DOES THE GDPR APPLY TO MY ORGANIZATION?

If your company offers goods or services (even for free) to EU residents, monitors EU residents' behaviors (including via cookies), or has any type of physical presence in the EU, then your organization is probably subject to GDPR compliance. Basically, the GDPR applies to any company worldwide that targets EU residents. If you decide what happens to the data, you're a "data controller," and your obligations are numerous. But if you're told by someone else what to do with the data, then you're a "data processor," and you'll have fewer obligations, but you can't sit back and wait.

In both cases, you'll need to know where the EU residents' personal data is stored, understand how data is processed, and start putting privacy and data protection higher up in your list. If you don't do so, your organization will be exposed to heavy fines (up to 4% of the worldwide annual turnover or €20 million—whichever is greater). The GDPR enforces a broad definition of personal data—including IP addresses, Internet of Things (IoT) data, or even your favorite color. This may explain why 87% of surveyed CIOs believe their current policies and procedures leave them exposed to risk under the GDPR[2].

KEY GDPR CONCEPTS AND WHAT THEY MEAN TO IT

- **Data protection by design:** Data protection should be part of any process that requires processing personal data. Organizations have a general obligation to implement technical and organizational measures to show that they have considered and integrated data protection into their processing activities.
- **Right for information, data access, rectification, object to, and to be forgotten:** Individuals will have the right to access their personal data, to have it rectified if it is inaccurate, to object and/or "block" the processing, and to request the deletion or removal where there is no compelling reason for its continued processing.
- **Data transfer and portability:** The GDPR allows individuals to obtain and reuse their personal data for their own purposes across different services.
- **Stricter definition of consent:** Consent under the GDPR must be a freely given, specific, informed, and unambiguous indication of the individual's wishes. Clear affirmative action is required (for example, no pre-checked boxes are allowed).
- **72-hour breach notification:** Organizations have to report within 72 hours certain types of data breaches to the relevant supervisory authority, and in some case to the data subjects affected.

GDPR IMPACT ON BACKUP

> ***Article 4 Definitions***
> *(1) 'personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person;*
> *(2) 'processing' means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organization, structuring,* ***storage,*** *adaptation or alteration,* ***retrieval,*** *consultation, use,* ***disclosure by transmission,*** *dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction; […]*

As explained earlier, the GDPR is not solely a legislative exercise, it has direct impact on backup, archive, and disaster recovery (DR).

Article 17 Right of erasure ('right to be forgotten')

The data subject shall have the right to obtain from the controller the erasure of personal data concerning him or her without undue delay and the controller shall have the obligation to erase personal data without undue delay where one of the following grounds applies: a) the personal data are no longer necessary in relation to the purposes for which they were collected or otherwise processed; b) the data subject withdraws consent [...] d) the personal data have been unlawfully processed; [...]

The EU is calling for organizations to optimize their data handling practices in a number of key areas, including data removal and deletion. EU residents have the right to be removed from the records of companies they have previously authorized to collect and store their data.

It is called the "right to be forgotten" (which means purging the data, including from backups and DR copies), and it is perhaps the most written-about obligation of the GDPR—and probably the most impactful one for IT administrators. It gives an individual the right to order a business to erase his or her personal data. Data controllers will have to erase all copies or links to personal data where the data subject withdraws consent, and there is no legal ground for processing it.

Chapter 3 Rights of data subjects

Article 12 Transparent information, communication and modalities for the exercise of the rights of the data subject

(1) The controller shall take appropriate measures to provide any information [...] to the data subject in a concise, transparent, intelligible and easily accessible form, using clear and plain language [...]

(3) ***The controller shall provide information on action taken on a request*** *under Articles 15 to 22 to the data subject* ***without undue delay and in any event within one month of receipt of the request.*** *That period may be extended by two further months where necessary, taking into account the complexity and number of the requests. The controller shall inform the data subject of any such extension within one month of receipt of the request, together with the reasons for the delay. [...]*

Organizations have one month to answer a data removal request and ensure that all traces of personal information are wiped from their systems. But removing an individual from a historical backup can be challenging! What happen to long retention backups? How do you deal with a removal request when your backup files include thousands of database entries?

The GDPR is open to interpretation, so we asked an EU Member State supervisory authority (CNIL in France) for clarification. CNIL confirmed that you'll have one month to answer to a removal request, and that you don't need to delete a backup set in order to remove an individual from it. Organizations will have to clearly explain to the data subject (remember, "using clear and plain language") that his or her personal data has been removed from production systems, but a backup copy may remain, but will expire after a certain amount of time (indicate the retention time in your communication with the data subject). Backups should be used only for restoring a technical environment, and data subject personal data should not be proceeded again after restore (and deleted again). While this adds some complexity, it allows organizations to have some time to re-engineer their data protection processes.

Based on our understanding, here is our recommendation on how to complete a removal request with a limited impact on your existing backup processes and infrastructure.

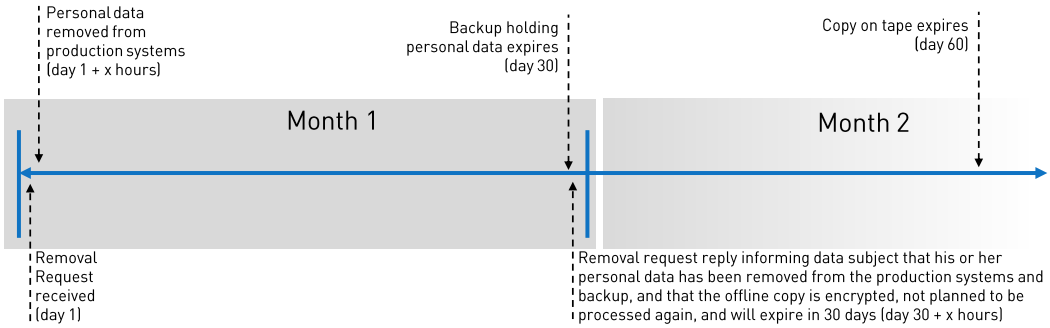

Article 17 Right of erasure ('right to be forgotten')

(1) The data subject shall have the right to obtain from the controller the ***erasure of personal data*** *[...]*

(3) Paragraphs 1 and 2 ***shall not apply to the extent that processing is necessary:*** *[...]*

b) ***for compliance with a legal obligation which requires processing by Union or Member State law*** *to which the controller is subject [...]*

In most of the cases, data will automatically expire with backups (most companies have backup retention schedules of three to five weeks—specific regulations are not taken into account here). But backups can be kept for longer, and the right to be forgotten won't apply if the data controller needs to apply a longer retention for other compliance reasons.

But if an organization is using backups for longer than a month or two, not for compliance, ask yourself this question, "Do I really need to keep backups for months/years or should I start archiving instead?"

RANSOMWARE AND GDPR

Viruses and ransomware attacks in particular can be seen as a major breach of the GDPR (some viruses are known to upload personal data to hackers). Employees are one of the main causes of ransomware breaches, often fooled by sophisticated social engineering attacks combined with ransomware. GDPR requires organizations to properly train their staff about security.

Article 39 Tasks of the data protection officer

(1)The data protection officer shall have at least the following tasks:

a) to inform and ***advise the controller or the processor and the employees who carry out processing of their obligations*** *pursuant to this Regulation and to other Union or Member State data protection provisions;[...]*

The GDPR offers a broad definition of data breaches and ransomware attacks, as they can potentially encrypt servers or workstations that process personal data, clearly "a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorized disclosure of, or access to, personal data transmitted, stored or otherwise processed." So, a ransomware attack is likely to trigger the 72-hour notification obligation, unless an organization can demonstrate that the data is unreadable (read encrypted).

Some consider data backups to be the best defense against these kind of viruses—but more and more ransomware hackers are hitting backups before they start encrypting production servers or workstations. And the only way to be fully protected is to use offline media such as tape.

As we saw earlier, it can be challenging to delete individual entries to meet the right-to-be-forgotten requirements, and this is even more true for tape. A pragmatic approach can be to keep one copy of data on tape for ransomware protection, with a one- or two-month retention policy. This way organizations can be protected against ransomware and deal more easily with the right to be forgotten. Tape libraries can also be used, in parallel with data protection, as an archive target. Archiving applications are by default easier to handle when it comes to data deletion, and can help keep storage costs low and data access simple for archiving data.

GDPR IMPACT ON DISASTER RECOVERY (DR)

> ***Article 32 Security of processing***
> *(1) Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the controller and the processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including inter alia as appropriate:*
> *[...]* ***b) the ability to ensure the [...] availability and resilience of processing systems and services;*** *c) the ability to restore the availability and access to personal data in a timely manner* ***in the event of a physical or technical incident;*** *d) a process* ***for regularly testing, assessing and evaluating the effectiveness of technical and organisational measures for ensuring the security of the processing.***

Simply keeping a backup of the data will not be good enough. You'll need to have a DR plan and test it regularly to be able to "restore the availability and access to personal data in a timely manner in the event of a physical or technical incident." Disasters happen, whether they are caused by a human error, as happened to Amazon in March 2017[3], or by a natural disaster, like the Vodafone UK data center outage due to floods in 2016[4]—but at the end of the day, disasters happen when people least expect them.

DR is a must-have to protect your data, but also to comply with the GDPR, as outlined in article 32(1). Your DR facility will also need to meet GDPR compliance, and if you're outsourcing all or part of your DR infrastructure, then your DR provider may be considered a "data processor," subject to GDPR compliance as well. This will also require that you use encryption during the replication between your main data center and the DR facility (data protection by design).

A couple of things to keep in mind:

- The right to be forgotten will apply to DR as well, adding extra steps and time to the removal process.
- The 72-hour breach notification also applies, requiring even more reactivity if the breach happens at the DR facility.
- If you're replicating (transferring) data outside the EU for DR purposes, you'll need to meet the conditions of chapter 5 of the GDPR.

Chapter 5 Transfers of personal data to third countries or international organisations

Article 44 General principle for transfers

Any transfer of personal data which are undergoing processing or are intended for processing after transfer to a third country or to an international organisation shall take place only if, subject to the other provisions of this Regulation, the conditions laid down in this Chapter are complied with by the controller and processor, including for onward transfers of personal data from the third country or an international organization to another third country or to another international organization. All provisions in this Chapter shall be applied in order to ensure that the level of protection of natural persons guaranteed by this Regulation is not undermined.

Article 45 Transfers on the basis of an adequacy decision

(1) A transfer of personal data to a third country or an international organization may take place where the Commission has decided that the third country, a territory or one or more specified sectors within that third country, or the international organization in question ensures an adequate level of protection. Such a transfer shall not require any specific authorization. […]

HOW QUANTUM ENABLES GDPR COMPLIANCE

There isn't a one-size-fits-all solution, but Quantum can help you in your journey to GDPR compliance.

As mentioned earlier, technology plays a part in GDPR compliance, but be skeptical of vendors saying their solution is "GDPR compliant" when there is no certification to support this claim yet. A certification is planned, but no details have been published at the time of writing. No one piece of software or hardware can ensure all aspects of the GDPR are met, but Quantum solutions can help you with very specific aspects of your GDPR compliancy regarding backup, DR, and archive—limiting the impacts on your day-to-day operations.

Quantum's portfolio includes purpose-built solutions uniquely designed for data protection and archiving, tightly integrated with many regulation compliance solutions, e-discovery tools, and robust backup applications. But at the same time, the solutions are flexible, agnostic, and multi-protocol to accommodate any changes in your infrastructure, which is happening more and more due to the high rate of acquisition and mergers. Quantum's portfolio can help your organization demonstrate that you have addressed "data protection by design and default."

The following diagram presents Quantum's portfolio for a rock-solid data protection and compliance solution.

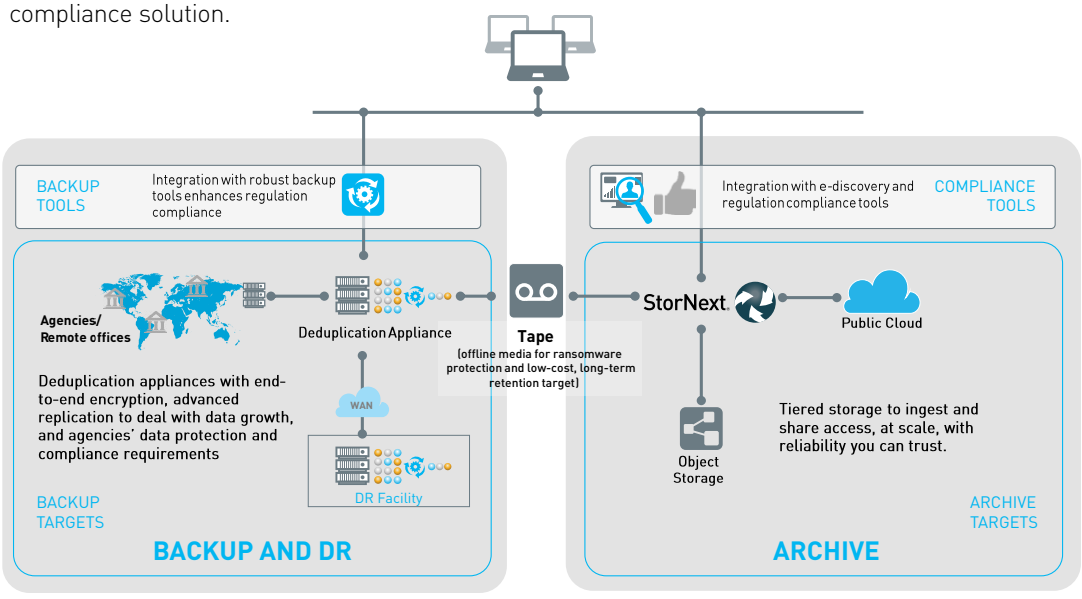

SECURE BACKUPS, DR COPIES, AND DATA TRANSFERS

The GDPR requires you to *"ensure the [...] availability and resilience of processing systems and services; [...] and the ability to restore the availability and access to personal data in a timely manner in the event of a physical or technical incident"* and to *"implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk."* In other words, you need to protect your data with secure backups, DR, and encrypted data transfers. Securing your data throughout its lifecycle and protecting against the threat of data breaches are critical elements to take into account. That's why Quantum's data protection solutions can encrypt data from the moment it's backed up until expiration— all using hardware-based, military-grade encryption—so your data is secure without any impact to performance.

Not all backup targets have been created equal in regards to encryption. The GDPR requires you to back up personal data and encrypt it where possible (e.g., encrypt backups). But if you're using a deduplication appliance to reduce the size of the storage required to keep your backups, this may create an issue—as most vendors use software-based encryption. While the level of encryption may be good enough, you can expect a performance drop of 15% to 20% when activating software-based encryption; this can lead to non-compliance if you can't meet your backup window, hence, personal data will not be backed up.

At Quantum, our DXi® deduplication appliances use self-encrypting drives (SED) to encrypt all file data and metadata, configuration files, and the DXi software and operating system. DXi Federal Information Processing Standards (FIPS)–certified SEDs are paired with their respective controllers, and data can only be read back from the disk by the same controller that wrote it. SEDs have integrated encryption hardware, so the result is zero performance impact. You can comply with your backup window and the GDPR at the same time.

Since the security of SEDs is independent of the operating system (OS), software attacks on the OS, basic input/output system (BIOS), and so on are not effective against SEDs. They are not vulnerable to alternative boot approaches or memory attacks (such as evil maid and cold boot attacks, or side channel attacks like acoustic cryptanalysis and more).

Quantum DXi appliances and Scalar® tape libraries provide encryption of data-at-rest as well as data-in-flight (using 256-bit AES encryption) to protect data during replication across sites or in the cloud. But data security is about more than just encryption. That's why Quantum data protection products include role-based access and support for Active Directory (AD) and Lightweight Directory Access Protocol (LDAP), so you can easily control access across your environment.

In addition to encryption, DXi systems support other data security features (some examples below):

- **Secure File Shred:** This DXi utility can securely and permanently erase sensitive data. During secure shred, all residual data associated with deleted files or virtual cartridges is securely erased from the disk drives by performing a single-pass overwrite with zeros.
- **Built-in system integrity checks:** These routines protect the DXi solutions by continually checking the state of the system's hardware and software conditions using a built-in, automated test process during normal operations. If anomalies are detected, full data and index verification/correction processes are performed in the background.
- **T10-PI standard:** T10-PI ensures that data is validated as it moves through the data path, from the application, to the host bus adapter (HBA), to storage—enabling seamless end-to-end integrity.

RANSOMWARE PROTECTION

Deduplication disk appliance replication can help protect against ransomware by creating "air gaps." The use of non-standard NAS protocols can also help "hide" your backup targets from ransomware. DXi, for example, supports multiple protocols and presentations, so it is tightly integrated with your backup software, such as Symantec OpenStorage Technology (OST), Veeam Data Mover Service, virtual tape library (VTL), and many more. But the last line of defense against ransomware attacks is still tape.

Tape allows you to create a ransomware-free zone—not hiding anything—just being physically disconnected from your network. Even a compromised backup administrator account can't access and use ransomware to encrypt your backups on tape. Tape is unplugged, hence offline, and once removed from the system, is no longer accessible electronically.

DXi deduplication appliances make tape integration easier by supporting direct path-to-tape (PTT). DXi appliances can create physical media outside the backup window via PTT without sending data over the backup SAN and without using the backup servers for data movement. Bar-code identity is tracked across virtual and physical media, and the tape creation operation can be initiated and managed by the direct PTT feature in backup applications, which supports independent retention policies for the different locations. Moreover, Quantum's Vision® management interface manages both disk backup and tape systems from a common console—saving time and increasing data availability by giving users a unified view of all of their backup devices.

Quantum Scalar tape libraries can not only be used to protect one copy of your backup against ransomware, but also to help you to archive data.

ARCHIVING DATA

Backup should be focused on restoring a service and helping with GDPR availability obligations, not on storing historical records. Having an archiving strategy can help organizations to be more easily compliant with the GDPR (specifically, the rights to access/rectify and remove) and in particular with article 15 that states EU residents have the right to inquire if and how their data is being processed. To be able to do this, they must first understand where the personal data of that individual is stored. The right to access or be forgotten becomes the right to be found.

How do you store more data in the most cost-effective way and compliant with the GDPR?

Quantum's multi-tier, policy-driven design solutions, such as Xcellis® workflow storage, allow higher-priority files to be stored on high-performance disk while lower priority files can be moved off to less expensive storage platforms, such as file-based tape or cloud storage. Cost and performance are dynamically adapted according to the value of your data.

Archiving data also reduces the pain and cost associated with data growth. Rather than spending more on expensive primary and backup storage to accommodate data growth, you can move data to protected archive or extended online storage. Your data can then be accessed and protected on less expensive storage without disrupting your day-to-day workflow. As data is moved to archive storage, it is no longer part of the regular backup process, so there is much less data to back up and backups take much less time.

Regarding GDPR compliance, archive is usually more GDPR friendly than backup, as it usually stores data in native format, so your data protection officer (DPO) can directly access data without requesting support from the IT team. (DPO is an enterprise security leadership role required by the GDPR; see GDPR articles 37 and 39.) This makes data rectification or removal operations easier. Quantum utilizes intelligent data movers to transparently locate data on multiple tiers of storage. This enables you to store more files at a lower cost, without having to reconfigure applications to retrieve or delete data from disparate locations. Instead, applications or authorized users continue to access files normally, and Xcellis automatically handles data access—regardless of where the file resides. As data movement occurs, Xcellis also performs a variety of data protection services to guarantee that data is safeguarded both on-site and off-site.

In addition to the GDPR, Quantum solutions are compliant with many international regulations when combined with Quantum's partner software, iTernity. iTernity software bundles the data and the corresponding metadata in special archive containers. These containers are secured against manipulation and unauthorized deletion. Files that have reached the retention date and are not on legal hold can be securely deleted. iTernity offers compliance solutions for many international regulations and laws.

FINDING DATA

Some unstructured data can be difficult to find. To help our customers find data subjects' personal information, Quantum has partnered with Veritone, developer of the world's first artificial intelligence (AI) operating system, to offer aiWARE™ for Xcellis. It can help you search for content in video files, text, audio, and more.

SOLUTION BENEFITS

- Strong integration with robust backup tools, and compliance and e-discovery software to ensure end-to-end compliance
- Deduplication appliances with hardware encryption and advanced replication for DR
- Tape integration for ransomware protection
- Tiered-storage archiving to simplify archiving

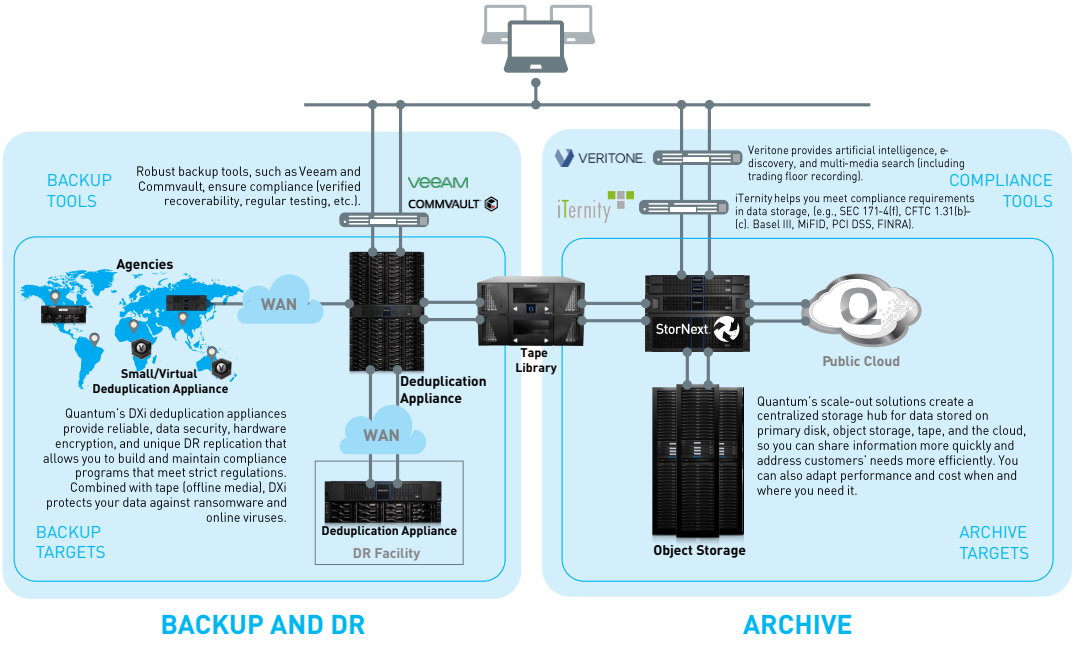

FEATURED PRODUCTS

- DXi deduplication appliances provide high-performance, scalable storage for backup and multi-site disaster recovery, with the industry's most efficient design. Learn More
- Scalar LTO tape storage provides the lowest cost long-term storage for archiving and retention, and offline storage to protect against ransomware. Learn More
- Scalar Key Manager is a scalable, secure, simple key management solution designed to work with Quantum's Scalar tape libraries. Learn More
- aiWARE for Xcellis provides on-premise AI and media data mining for StorNext® content repositories. Learn More
- Quantum and iTernity have partnered to help you comply with data retention mandates. Learn More

For more information, visit **www.quantum.com**.

(1) Hopkins, Nick. September 25, 2017. Deloitte Hit by Cyber-Attack Revealing Clients' Secret Emails. www.theguardian.com/business/2017/sep/25/deloitte-hit-by-cyber-attack-revealing-clients-secret-emails

(2) IDC&Commvault 2017. Are You Ready For GDPR?. https://www.commvault.com/solutions/by-topic/gdpr

(3) Del Ray, Jason. March 2, 2017. Amazon's Massive AWS Outage Was Caused by Human Error. https://www.recode.net/2017/3/2/14792636/amazon-aws-internet-outage-cause-human-error-incorrect-command.

(4) Smolaks, Max. January 5, 2016. Vodafone UK data center suffers outage due to floods. http://www.datacenterdynamics.com/content-tracks/security-risk/vodafone-uk-data-center-suffers-outage-due-to-floods/95438.fullarticle.

ABOUT QUANTUM

Quantum is a leading expert in scale-out tiered storage, archive, and data protection, providing solutions for capturing, sharing, and preserving digital assets over the entire data lifecycle. From small businesses to major enterprises, more than 100,000 customers have trusted Quantum to address their most demanding data workflow challenges. Quantum's end-to-end, tiered storage foundation enables customers to maximize the value of their data by making it accessible whenever and wherever needed, retaining it indefinitely and reducing total cost and complexity. See how at **www.quantum.com/customerstories**.

Quantum®

www.quantum.com
800-677-6268

WP00236A-v01 Dec 2017