PROTECTING YOUR DISTRIBUTED DATA WITH THE CLOUD

Best Practices for Meeting Your Global Data Protection Needs

START

druva

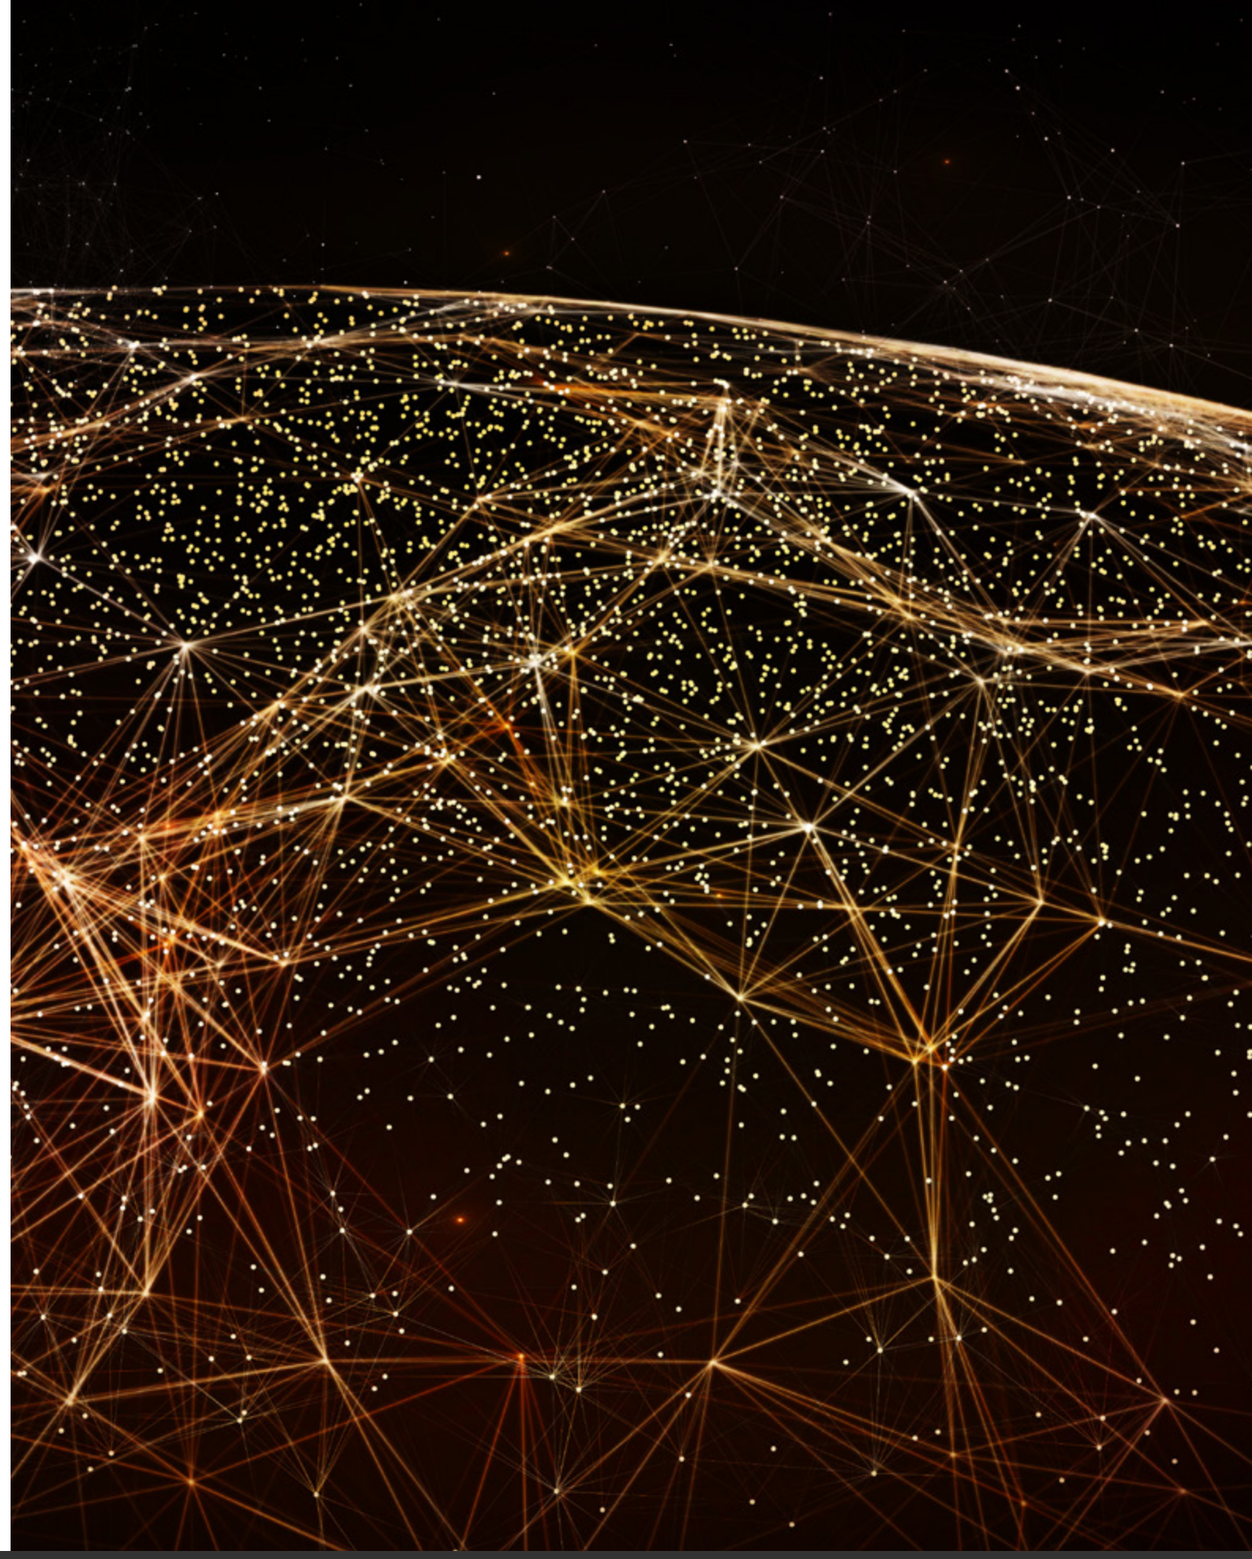

Introduction

We Are Using More Systems, Creating More Data, More Often.

As the Internet of Things (IoT) expands and evolves, many of our everyday activities will involve connected devices. By 2025, the average person with Internet service will interact with connected devices nearly 4,800 times per day—basically one interaction every 18 seconds.[1]

All this interactivity is generating a lot more data. According to IDC, the global datasphere will grow to 163 zettabytes, ten times the 16.1 ZB of data generated in 2016.[2] And much of this data will be owned by corporations. Some of it will reside in centralized data repositories, but most of it will be distributed across countless devices and local data stores. And **almost all of it will need to be protected, managed, and secured.**

Protecting distributed data is a big, complex task. And, **if remote data expands faster than IT can manage it, the consequences can be severe.** Loss of data and intellectual property, downtime, non-compliance with data privacy regulations, and increased vulnerability to malware are just some of the dangers that businesses without a solid framework for distributed data protection can face.

This ebook outlines proven strategies for protecting your distributed data and reducing your risk.

[1]https://www.seagate.com/files/www-content/our-story/trends/files/Seagate-WP-DataAge2025-March-2017.pdf
[2]IBID

Are You at Risk?

Most businesses will have to deal with distributed data at some point in their development. However, some situations present more risk than others. The following table highlights three business cases in which distributed data risk is especially high.

Business case	Characteristics	Potential risks
Newly formed company offices 	• Minimal or “TBD” IT presence • Makeshift or “TBD” backup and recovery policies • Lack of meaningful internal and external SLAs • Focus on rapid growth over compliance • A tendency to put off data protection until the next budget cycle	• Permanent data loss • Downtime resulting from incomplete or nonexistent disaster recovery policies • Failure to comply with local or regional regulations • “Data sprawl” and “data creep”: growing stores of decentralized data that become increasingly difficult to handle over time
Remote and branch offices 	• Little to no IT presence or “generalist” IT expertise that doesn’t include data protection • Lack of clear and frequent communication between local IT and data centers • Established backup and recovery policies that may be frequently disregarded • Internal or external SLAs that are not frequently monitored • Reliance on tape and other outdated technologies	• Permanent data loss • Downtime resulting from slow or cumbersome disaster recovery policies • Failure to comply with local or regional regulations • Data sprawl or creep • High costs when IT decides it’s time to upgrade outdated technology
Newly merged companies 	• Complex and piecemeal systems, including multiple data centers and data protection systems • Multiple, potentially conflicting IT organizations • Lack of consistent backup and recovery policies and SLAs • Reliance on tape and other outdated technologies • Limited funds for IT	• Permanent data loss and intellectual property risk • Downtime resulting from incomplete, conflicting, or unworkable disaster recovery policies • Failure to comply with local or regional regulations, or legal data requirements • Data sprawl, creep, and duplication • High costs when IT decides it’s time to standardize the merged company on a single data protection platform or upgrade outdated technology

It’s Harder Than You Think

5 Reasons Why Protecting Distributed Data in Decentralized Environments Isn’t Easy

Protecting decentralized data is typically time-consuming, expensive, and incomplete. Why? According to our conversations with customers, the top five potential issues are:

1. **Lack of visibility.** Corporate IT organizations are responsible for ensuring that local offices comply with data protection policies and tracking SLAs, but they often do not have visibility into what local teams are—and are not—doing. This knowledge gap presents a significant risk.
2. **Limited expertise.** IT staff at remote offices tend to be generalists and may lack the skills to quickly restore data in the case of a failure. They may also be overworked and face pressure to prioritize customer-facing issues over routine data protection.
3. **Manual processes.** Manual processes that require non-IT staff to make tape backups and ship them offsite are extremely time-consuming, costly, and often neglected. This is especially true when teams are managing change or growth.
4. **Budget constraints.** In remote offices, IT funds may be spread thin or tightly controlled by a centralized procurement organization, which can make obtaining funds for data protection very difficult.
5. **Regional data privacy rules.** Many countries and regions have rules about copying and storing data that was collected locally. In many cases, special expertise is necessary to comply with these rules. Offices that lack this expertise may take the path of least resistance and comply by not backing up data at all.

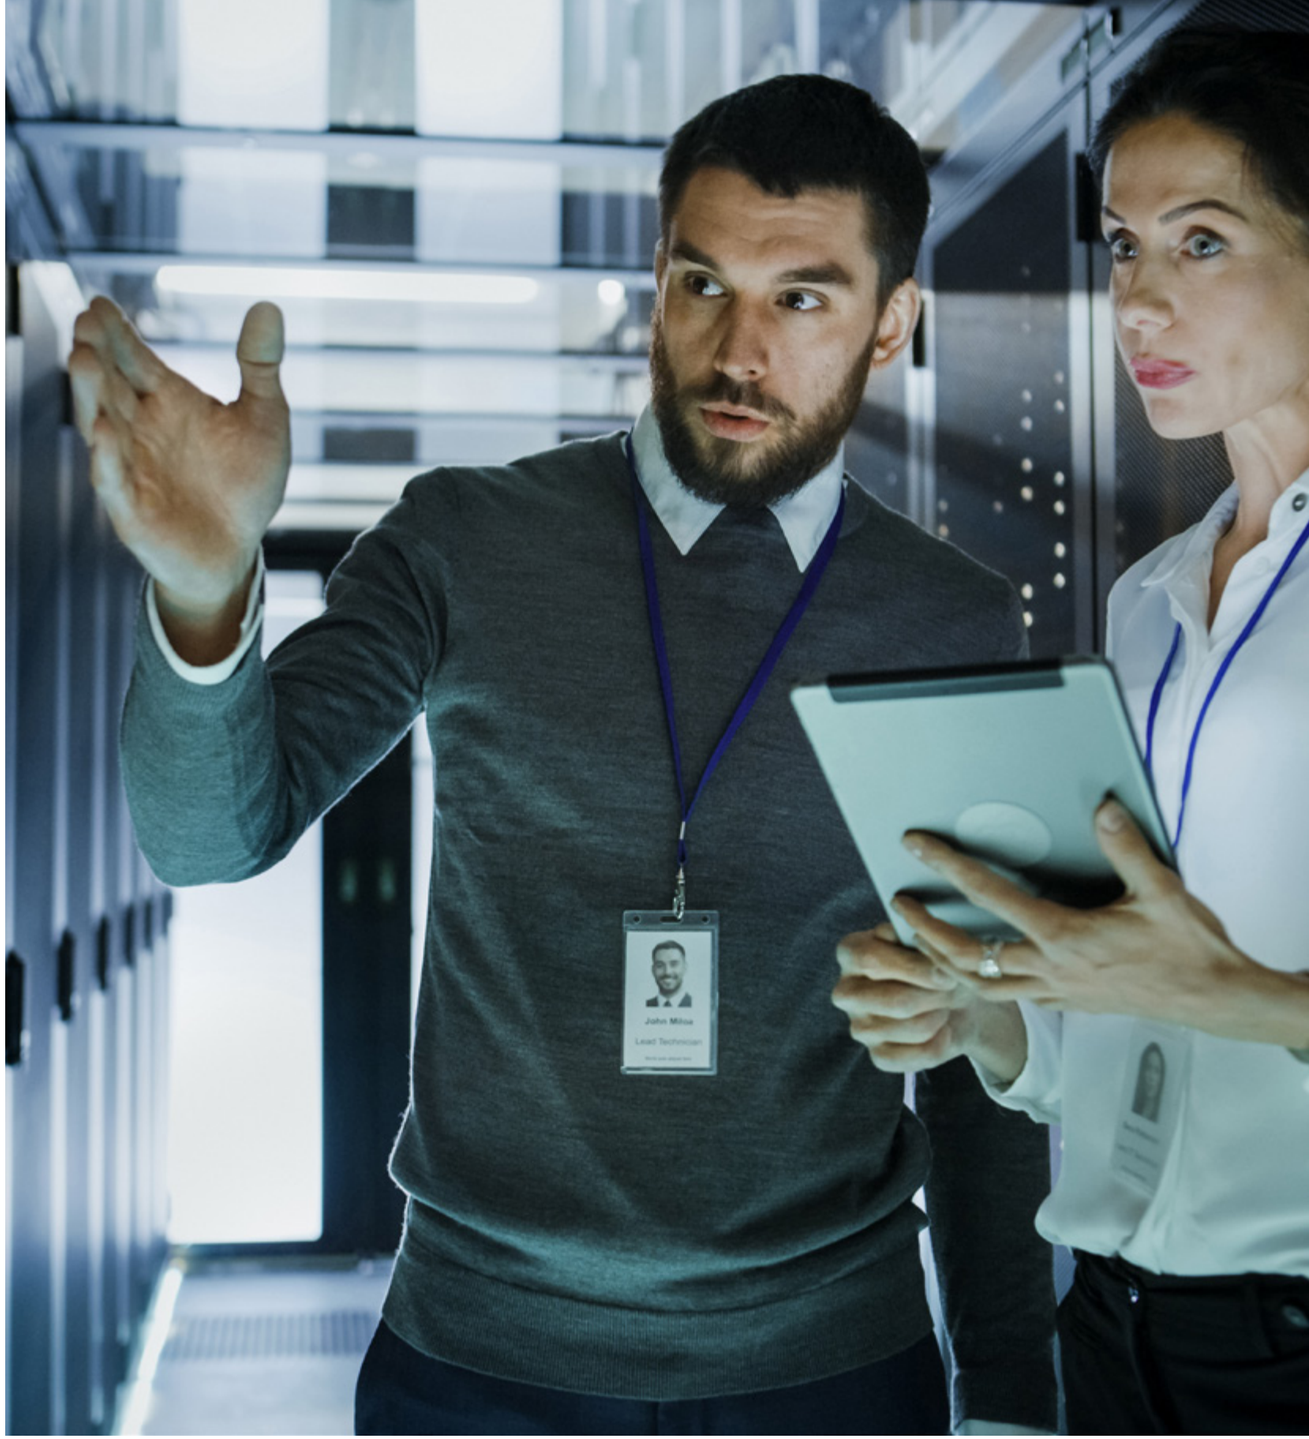

It’s in the Cloud

Why Cloud-Based Data Protection Is Ideal for Distributed Data

Remote data protection has traditionally relied on two main solutions, both with significant drawbacks:

- **Tape backups** require busy staff members to manually initiate backups and ship tapes to approved storage locations. At some organizations, tape backups are considered a low priority and only happen sporadically, which presents the potential for major data losses in the event of a failure. At other organizations, timely and frequently tape backups can absorb the lion’s share of IT resources and generate escalating costs.

- **Remote backups to on-premises servers** can be difficult for non-IT staff to manage, bandwidth-intensive, and very expensive. Adding capacity requires buying hardware and software up front, and, if you ever consolidate operations, you’ll be stuck with unused equipment. Best case, you’ll end up storing redundant data on redundant equipment. Worst case, you’ll pass on new capacity to save money and end up with permanent data loss.

What is Cloud-Based Data Protection

Both traditional forms of data protection rely on a centralized computing model that tends to treat distributed data as an afterthought. At its best, cloud-based data protection can cost-effectively protect distributed data while offering convenient, centralized management.

What to Look For in a Cloud-Based Backup and Recovery Solution

What does an effective cloud-based data protection solution look like? Some characteristics include:

- **Cloud-native architecture:** Cloud-native architecture is optimized for performance and security over the public cloud. It offers centralized management of backup and recovery processes, consistent performance even with petabytes of data, and lower TCO compared to hybrid and cloud-enabled solutions.

- **Bandwidth-friendly features:** Global, source-side deduplication and incremental storage models, which only transmits changes to data, can help businesses meet recovery point objectives (RPOs) and get the most out of limited bandwidth.

- **Disaster recovery on demand:** Cloud enables data to be immediately accessible from any backup point, down to individual files, and your VMs can be configured to failover and quickly spin-up in the cloud, eliminating system downtime.

- **Auto-tiered storage:** Auto-tiering moves data to different storage tiers, reducing overall data protection costs, based on usage policies without creating duplicate copies. Once IT defines their policies, the system carries them out automatically, no manual intervention required.

- **Data-protection-as-a-service:** Ideally, a cloud-based data protection solution should be delivered as a service, so businesses can focus on data protection policies and requirements rather than underlying infrastructure.

Case Study: Pall Corporation

Slashing Costs with Cloud-Based Data Protection

A Global Business Facing a Costly Upgrade

Port Washington, NY–based Pall Corporation has business operations that span the globe. Pall provides a rich portfolio of advanced filtration, separation, and purification equipment and services, but with such widespread operations and services, they faced a daunting challenge to efficiently manage all of their data. Their siloed legacy systems, due for an upgrade with a staggering $1.2M pricetag, still required IT managers to manually manage a collection of tape backups and disaster recovery systems.

Growing Pains and Data Loss

The time and cost required for tape backups was getting out of control. IT staff handled thousands of tapes every day and costs for replacing, storing, and transporting tapes were becoming untenable. Their initial attempt to replace this system with a hosted solution, however, was not a success. It ended with 1TB of data being completely lost, and a frustrating lack of clarity on who was to blame as the provider and AWS pointed fingers at each other.

A Second Attempt at the Cloud

To address these challenges, Pall wanted a solution that worked well with AWS infrastructure. In addition, they did not want to manage multiple vendors, given the challenges they had faced previously with data loss. Instead, they aimed to consolidate management of these backup systems into a single cloud-based software-as-a-service (SaaS) solution. After testing several solutions, they chose Druva Phoenix, a native-cloud-based data protection solution.

Benefits of Druva Phoenix

Druva Phoenix enables Pall to centrally manage remote office data protection in the cloud though one clear SaaS provider. Benefits include:

- Over 50 offices around the world are now covered by a single central solution for offsite data protection and recovery that requires only two remote admins
- 78% overall budget savings
- 95% reduction in labor related costs
- Users can now restore data on their own in seconds—a task that previously took an entire day

“We were really looking for a cloud solution all along.... For disaster recovery we cannot have something simply be onsite.”

Jessica Fletcher,
IT Analyst Supervisor,
Pall Corporation

To read the complete case study, **go here**.

Case Study: TRC

Orchestrating a Distributed, Post-Merger Environment

120 Global Offices After Multiple M&As

TRC Companies, Inc. is a national engineering, environmental consulting, and construction management firm that provides integrated services to the power, environmental, infrastructure, oil, and gas markets. With more than 120 offices in the United States and the UK, along with steady growth through mergers and acquisitions, TRC has inherited an enormous volume of legacy data and systems.

Integrating Disparate Offices and Data

Until recently, TRC struggled to integrate data from its proliferation of offices, many of which did not have dedicated IT staff on-site to ensure that mission-critical server data was being backed up in a timely or effective fashion. Geographically distributed sites used a variety of different, and often incompatible, tools—and as a result, data was being lost.

Without a single pane of glass to provide insight into the backup process, the corporate IT organization faced myriad questions, such as: “Where are we with our data? Are we backing it up, and is it secure? Who owns what, and how much of it?” And then, more importantly, “What is active data versus some of the data that’s not being used?”

A New Approach to Distributed Data

TAs TRC began the search for a comprehensive answer to their data management challenges, the company made the crucial decision that a winning solution would need to use the cost efficiencies and flexibility of the public cloud. The company’s rationale for this was simple: the system would not only need to handle their current operational needs, but also be capable of scaling to meet future requirements.

They ultimately chose Druva Phoenix, a cloud-based data protection solution, to centrally manage backup and archiving for remote office data in the cloud.

Benefits of Druva Phoenix

Druva Phoenix enables TRC to centrally manage remote office data protection in the cloud despite the fact that different offices run different technologies. It has allowed them to:

- Reduce the total cost of ownership (TCO) of data protection technology by 64.5%
- Achieve central visibility and management across 120 remote sites
- Dramatically improve backup workflows and controls for M&A
- Completely eliminate data loss

“What we were interested in was something that was truly born in the cloud and was optimized to handle the efficiencies of cloud...”

JP Saini,
CIO,
TRC Companies Inc.

TRC
Results you can rely on

To read the complete case study, **go here**.

Case Study: Andritz

Saving Time and Reducing Risk with Cloud-Based Data Protection

A Global Business with 200+ Local Offices

Andritz is a global supplier of systems and services for hydropower stations, the pulp and paper industry, and the metalworking and steel industries, as well as for solid/liquid separation in the municipal and industrial sectors. Their business, employing over 25,000 employees scattered across 200 worldwide offices, consists of many different organizations as best-in-class companies are acquired to complement existing services.

The Andritz IT team is responsible for this very diverse, decentralized environment. The company's offices range in size from less than ten to more than 200 people, and a single office may only have one file server or dozens of any mixture of file, SQL, and application servers. The IT team is geographically dispersed across multiple locations, but there are also some offices without any IT team members on site.

Problems with Tape Backup

Andritz relied on a tape-based server backup solution, which meant that in some locations non-IT staff are the ones to make sure that tapes are changed, etc. Even routine tasks, such as verifying that backups had occurred, were time consuming; the company estimates that their IT team spent about 70 hours each month just for this activity. And without any onsite IT, related tasks, such as restoring data from tape, were even more complicated; it frequently took multiple tries to locate the correct tape and then talk the onsite employee through loading it, finding the necessary files, and restoring them.

Problem Solved: Druva Phoenix

Druva Phoenix, a cloud-based data protection solution, has provided Andritz with a way to remove their legacy tape system and centrally manage remote office data in the cloud.

Benefits of Druva Phoenix

By adopting cloud-based backup through Druva Phoenix, Andritz is able to:

- Centrally manage data from 200 offices around the world through a single dashboard, with full visibility into physical and virtual server data backups and archives
- Rapidly access distributed data and recover from failures within seconds
- Automatically manage the full data lifecycle, with no more tape
- Comply with regional data residency policies

"A file restore that previously may have taken a couple of hours, or even days, can now be completed in seconds."

Brian Bagwell,
Director of IT - North America,
Andritz Inc.

ANDRITZ

To read the complete case study, **go here**.

About Druva

Druva is the global leader in Cloud Data Protection and Management, delivering the industry's first data-management-as-a-service solution that aggregates data from endpoints, servers, and cloud applications and leverages the public cloud to offer a single pane of glass to enable data protection, governance, and intelligence—dramatically increasing the availability and visibility of business-critical information while reducing the risk, cost, and complexity of managing and protecting it.

Druva's award-winning solutions intelligently collect data and unify backup, disaster recovery, archival, and governance capabilities onto a single, optimized data set. As the industry's fastest-growing data protection provider, Druva is trusted by over 4,000 global organizations and protects over 40 petabytes of data. Join the conversation at **twitter.com/druvainc**

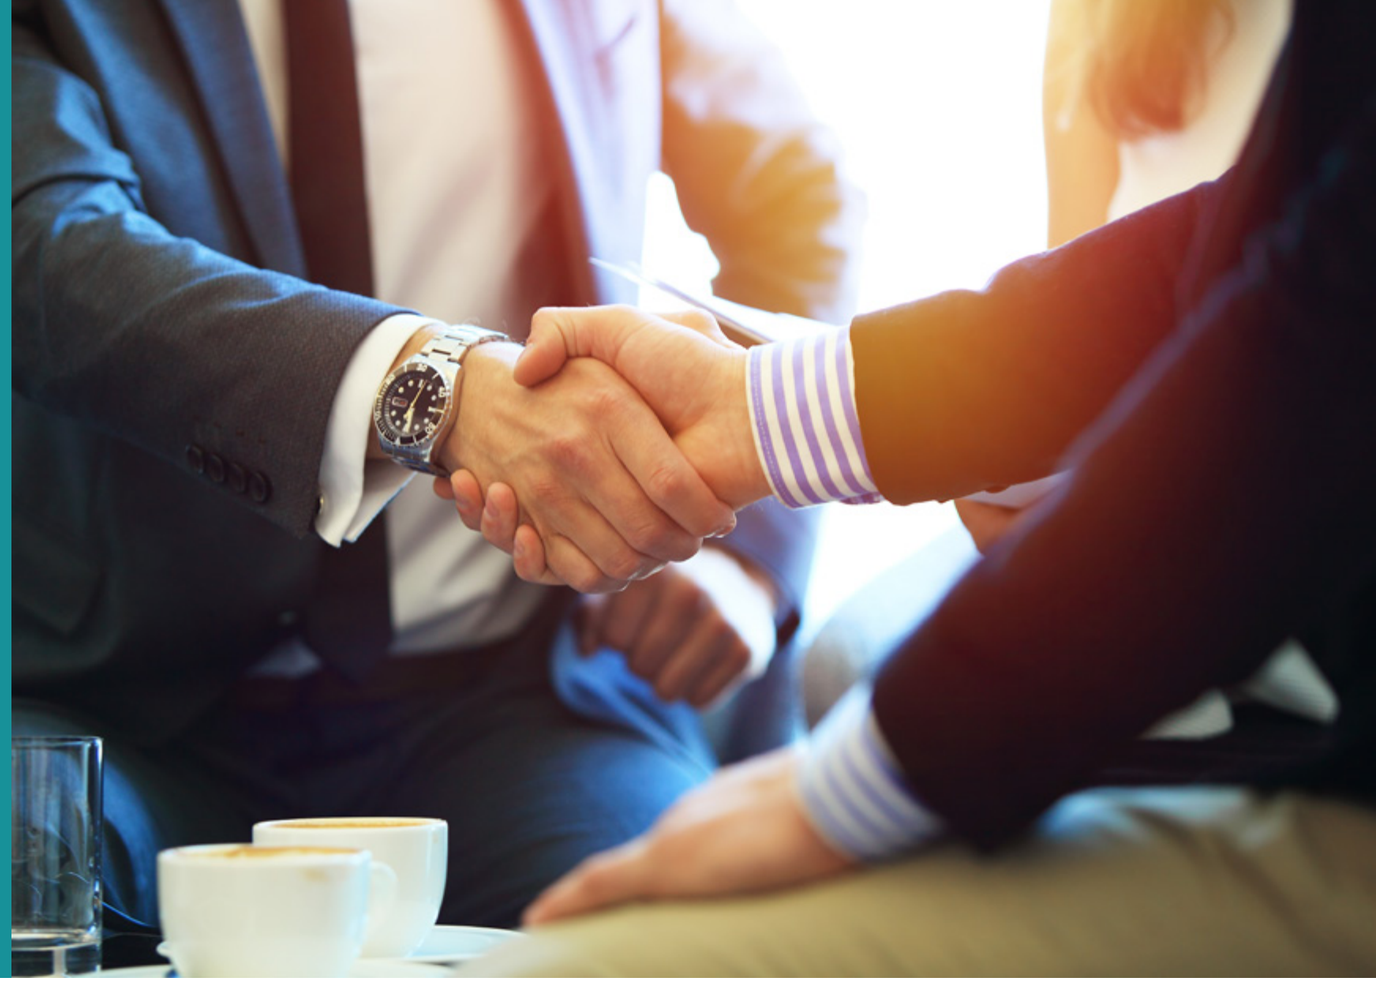

Druva, Inc.
Americas: +1 888 248 4976
Europe: +44 (0) 203-7509440
APJ: +919886120215
sales@druva.com
www.druva.com

© Copyright 2017 Druva, Inc. All rights reserved