

Druva Phoenix FAQ

About Phoenix

What is Phoenix?

Druva Phoenix is a cloud based backup and archival solution aimed primarily at remote office servers. Since Phoenix is cloud-targeted backup, there is no elaborate deployment needed to get started. All that is needed is an install of the Phoenix agent on servers that need to be backed up. Phoenix agents establish a connection between the servers and Phoenix cloud. After performing some quick configuration tasks from the Phoenix Management Console, the Phoenix setup will start backing up data.

What types of servers does Phoenix back up?

Phoenix backs up servers using Windows and Linux operating systems. Here is a list of Windows servers that Phoenix supports:

- Windows Server 2003 (32-bit / 64-bit)
- Windows Server 2008 (64-bit)
- Windows Server 2012 (64-bit)

Here is a list of Linux servers that Phoenix supports:

- CentOS 6.x (64-bit)
- Red Hat Enterprise Linux (RHEL) 6.x (64-bit)

With Windows, NTFS file systems are supported, and on Linux there is support for ext3 and ext4, but in both cases a journaling file system is required.

How does Phoenix work?

The Phoenix setup includes a Phoenix Master, which is hosted in our cloud. The Phoenix Master orchestrates all backup activity, and is the place where an administrator defines policies and manages servers. The master acknowledges server requests and redirects them to assigned storage locations. The second piece is the agent - Phoenix agents are deployed on the servers that need to be backed up. At specified intervals, the Phoenix Master instructs Phoenix agents to transfer data from associated servers to the storage assigned for this purpose. When a restore is requested, Phoenix agents communicate directly with the storage assigned for this purpose to receive data.

What are the unique capabilities of Phoenix?

- Phoenix is a cloud-based, unified server backup and archiving solution. Built on the advanced cloud technologies offered by Amazon Web services, Phoenix promises best-in-class durability and infinite scalability for growing server data needs.
- The powerful global deduplication that Phoenix employs ensures reduced backup windows and faster restores. With Phoenix, one can expect faster backups, with native server connections of up to 100Mbps. When deduplication is taken into account, effective speeds of up to 1Gbps can be seen.

- The Phoenix CloudCache Server can further increase the speed and efficiency of both backups and restores. Ideal for those with extremely bandwidth challenged locations, or with very short RTO objectives, data of a determined time interval can be cached on an appliance close to the server. This allows for LAN speed backup and restore, with the ability to send this data to the cloud on a schedule that makes sense for your enterprise.
- A Phoenix deployment is entirely cloud based. Free up existing IT staff from having to manage multi tier implementations with secondary storage locations, or increase the reliability of backups in offices that may not have appropriate levels of IT support.
- Licensing is based off of the size of the current data on the server, not the size of the archived data, so costs don't increase with longer retention requirements or higher change rates.
- The flexible design of Phoenix can be leveraged to support backup and restore needs of organizations of all sizes, from a single remote office file server to a globally distributed enterprise - all from a single console.

How is Phoenix different from inSync?

For the existing inSync customer, they will feel right at home with Phoenix. For those new to Druva's backup technologies, they will be impressed with the ease of use and low administrative overhead.

- Phoenix is a backup and restore solution for server data. In comparison, inSync is backup and restore solution for all endpoint devices, such as laptops, desktops, phones and tablets.
- Phoenix is entirely in the cloud; an install and activation of a server agent is all that is needed to get started. inSync comes as two offerings - either cloud or on-premise.
- Phoenix is designed for server data needs; data can be retained indefinitely, while permitting restores of data from any snapshot. Backups can happen at a faster rate, which is ideal for the larger amount of data and higher change rates of the average server, as compared to an endpoint device.
- The administration console of Phoenix caters to server backup administrators, allowing for quick insights and the ability to manage server data. By contrast, inSync is designed for endpoint devices, the console caters to the backup and restore needs of modern users, recognizing that most employees work with corporate data across many devices.
- Phoenix has been qualified at native backup speeds approaching 100Mbps. InSync is rated at less than 20Mbps. With server workloads being larger, these speed differences help speed up backup and restore speeds from the cloud.



Pricing and Licensing

How much does Phoenix cost?

Licensing Phoenix depends on the amount of source data being backed up, and the retention period for the data. For further information, contact a Druva account executive, or e-mail sales@druva.com.

What retention options does Phoenix support?

Phoenix retains data on warm storage and archival storage, depending on how old the data is. The retention options are up to 3 year, 7 year and unlimited retention.

What is source based pricing?

A Phoenix license is based on the size of the source data. The source data is the sum of all current data from the folders that are being backed up. For example, if the total size of an existing set of folders is 500GB, the source data is just that - 500GB. Phoenix retains all previous versions of the source data for the licensed retention period no matter what the change rate is. The license cost is always based on source data.

Do I pay separately for restores?

Unlimited restores are included in the license cost, there is never an additional cost to restore data.

Sizing and setup

How do I determine how much storage to purchase?

The storage required will depend on the amount of source data, including expected short term source data growth. Additional storage can be purchased as needed.

I purchased 10TB of data. How do I know how much of that is being used?

Storage consumption is displayed on the storage page (Manage > Storages) in the admin console. The Storage used column displays the total storage that is consumed across all servers. This is the ideal place to see storage utilization against purchased storage.

I consumed all the storage that I had paid for. Will my next backup fail?

No, it will not. If all licensed storage is consumed, additional storage can be requested through the sales team. In the interim, server backups will function normally.

How much target data am I allowed?

Phoenix considers only source data - which is the amount of data currently on a server. If the size of source data increases, additional storage can be requested. Changes in source data does not affect the data already under management. If the size of source data is reduced, Phoenix continues to store all of the historical data at no extra cost.



Backup

Where is my data backed up?

Data is backed up to Amazon Web Services, and is encrypted in transit and at rest. Phoenix leverages the robustness and the security standards of Amazon Web Services to ensure the security of backup data.

What type of backups does Phoenix perform?

The first backup of the server is a full backup, with deduplication of data across every server that has been previously backed up by a specific customer. Thereafter, every backup Phoenix performs is an incremental backup.

What happens to my backups after my contract expires?

After the expiration of a contract, backups will stop. However, Phoenix retains the data on the cloud for a period. If a customer chooses to extend a contract, backup and restore activities can continue. At the end of a contract, a customer can request immediate disposal of all of their data.

What speeds can I expect for backups and restores?

The speed of backup and restore depends on available bandwidth. Transfer rates of up to 100 Mbps can be seen, if there is enough client side bandwidth available to support these speeds.

How does the Phoenix deduplication algorithm work? How does it improve the backup experience?

Phoenix uses a variable block size technique (based on file type) to determine globally duplicated data across all of a specific customer's servers. The Phoenix agent only transfers unique blocks of data to the cloud. By doing this on the client side, massive bandwidth savings can be realized - up to 90% in many cases.

What happens if I lose connectivity during a backup?

Phoenix is designed to work efficiently even from sites with intermittent connectivity. If the Phoenix agent cannot reach Phoenix cloud, backups or restores stop temporarily. When the connectivity resumes, backups or restores start again. A single backup can therefore span multiple backup windows.

Security and Certifications

What certifications does Phoenix adhere to?

Phoenix adheres to the same certifications as InSync. As an IaaS provider, Amazon regularly updates their compliance certifications for all the standard certifications. Further details are available and up to date here: <http://aws.amazon.com/compliance/>. As the SaaS provider, Druva holds a number of specific certifications related to how we control and manage our own applications. This includes a ISAE 3000 Type II audit performed yearly by KPMG, as well as certifications around Safe Harbor and a review of our security and privacy controls for handling HIPAA-compliant protected health information (PHI).

What security standards do you employ for protecting my data?

To ensure data stays protected, data is encrypted using 256-bit SSL encryption while in transit. At rest, Phoenix encrypts data using a unique 256-bit AES key for each customer. In addition to encryption, data scrambling techniques are used to further obscure access to data.

I am required to enforce a policy where data must stay completely within my own country (Safe Harbor). Does Phoenix support that?

Yes, when provisioning a server the data can be scoped to server clusters within a specific geography ("availability zone", as it is known within Amazon Web Services). This can provide complete assurance that data never leaves a specific geography.

Phoenix CloudCache Server

What is Phoenix CloudCache Server? What are the benefits of using it?

Phoenix CloudCache Server is a dedicated virtual appliance that stores up to 30 days of backups locally (to enable quick local restores), and then periodically synchronizes this data to the Phoenix cloud. Since Phoenix agents communicate directly with CloudCache on a local network, the speed of backups and restores increases significantly. Moreover, administrators can schedule synchronization operations from Phoenix CloudCache Server to Phoenix Master during off-peak hours, thus ensuring an optimized utilization of WAN bandwidth.

How is CloudCache licensed?

Licensing is based on the storage space that is required for source data. This is a function of the source data on the server, the number of days retention is required, and the change rate. For more information about sizing or licensing, contact your account team, or sales@druva.com.

What difference will deploying a Phoenix CloudCache Server make to my backup and restore experience?

With CloudCache, two main differences will be seen. The backup window will reduce significantly as Phoenix agents back up data directly to CloudCache Server on the LAN. The data saved to CloudCache can be restored at any time during the pre-defined retention period directly from this local server. Local restores are faster, and they reduce the demand on WAN bandwidth. Additionally, data can be sent to the cloud during off hours, which is ideal for sites with limited available bandwidth.

How long does the Phoenix CloudCache Server hold my data?

An administrator can configure the number of days (up to 30 days) for which the CloudCache Server holds a copy of data locally. This is independent of the scheduling to sync data to the cloud, as data can exist in both locations.

Administration

How many administrators do you recommend for managing Phoenix?

Phoenix supports a flexible server administrator model, with the concept of a delegated administrator model if needed. These are admins who only have access to specific groups of servers. This is all controlled by Cloud Administrators, which have unrestricted access to the server. A deployment can have as many administrators as needed, but each environment should always have at least two Cloud Administrators at all times. If the credentials were lost for a single admin account, there would be no way to recover data from the system. This is by design, and ensures that only the customer has access to any data.

How can I delegate administrative activities to local administrators?

Group administrators manage a subset of servers within a deployment, with controlled access to the server groups to which they are associated. They have necessary permissions for server group activities such as managing backup policies, and managing servers belonging to their server groups. A group administrator can be assigned to manage multiple server groups. Similarly, a single server group can be managed by more than one group administrator.

Can I mass deploy Phoenix agents?

A third-party tool such as Microsoft SCCM can be used for installing and activating Phoenix agents. Once activated, the server can be configured through the master management console.

How do I leverage Activation Tokens for my deployment?

Activation tokens are created for the purpose of registering an agent to the associated deployment. Tokens can be set to expire after a specific number of uses or invalidated at any time by an administrator. A best practice is to use a single token to activate servers that share something in common. For example, one might want to use a token to activate all servers that are geographically co-located. Once registered, the token is not required for further operation.

Do you store my activation tokens?

No, Druva does not store activation tokens, as this token can be used to bind any server to a specific environment. Administrators have full control over creating and invalidating tokens at any time.

Managing Servers

What are the best practices for creating server groups?

In a production environment, we recommend that a server group is created for each set of servers that are geographically co-located and belonging to the same platform. Server groups are backed up to a specified storage location, can have the same delegated administrators and all have the same backup policy associated with them. Creating a server group for each location ensures faster backups and restores to and from region-specific storage, while reducing the demand on bandwidth. One example: create a server group for all Windows file servers located in Eastern Europe, and another server group for all CentOS file servers located in California.

Where can I find the Phoenix agent installer?

The Phoenix agent installer is available from a link in the Phoenix Admin portal, as well as during the server registration wizard process.

How do I adjust bandwidth available for a server backup?

Bandwidth is defined at the server group level. Servers within a server group inherit their bandwidth and backup policies from the group to which they belong.

Can I assign more than one backup policy to a server group?

A server group can only have one backup policy at a time. However, a backup policy can be detached from a server group, and a new one attached in its place.

Can I install Phoenix agent first, and then register servers?

When an administrator goes through the workflow to register a server, Phoenix generates an activation token. This token must be used to activate the Phoenix agents on a server. However, if installing Phoenix on multiple servers, this operation can be batched as well. With an appropriately generated token, multiple servers can be registered and the order of the install and registration does not matter.

Reporting

Can I get regular reports of backups and restores?

Yes, reports are sent out to administrators for the servers that they manage on a scheduled basis that is fully configurable. These reports can contain information on any errors or unsuccessful backups and restores.

How often do the graphs on the dashboard get refreshed?

The graphs on the dashboard view are a real time look into the activities on the system. These are refreshed in real time every time the page loads.



Supported Platforms

What versions of Windows servers are supported?

Here is a list of Windows servers that Phoenix supports:

- Windows Server 2003 (32/64-bit)
- Windows Server 2008 (64-bit)
- Windows Server 2012 (64-bit)

What versions of Linux servers are supported?

Here is a list of Linux servers that Phoenix supports:

- CentOS 6.x
- Red Hat Enterprise Linux (RHEL) 6.x

What versions of MS SQL server are supported?

In a future release, Phoenix will support SQL Server 2008 and 2012.

About Druva

Druva is the pioneer and market leader in data protection and governance at the edge of the enterprise, bringing visibility and control to business information in the era of the mobile workforce and consumerization of IT. Druva's solutions for managing data outside the corporate firewall reduce the loss of corporate information assets and address organizations' compliance, governance, forensics and eDiscovery needs. Headquartered in Silicon Valley with offices worldwide, Druva is privately held and is backed by Nexus Venture Partners, Sequoia Capital and Tenaya Capital. For more information, visit www.druva.com and connect at www.druva.com/blog.



Druva, Inc.

Americas: +1 888-248-4976

Europe: +44.(0)20.3150.1722

APJ: +919886120215

sales@druva.com

www.druva.com