

Why Cloud Backup/Recovery (Private, Public, or Hybrid) WILL be Your Data Protection.

**Marc Staimer, President and CDS of Dragon
Slayer Consulting**

Marc's consulting practice of 11+ years provides consulting in the areas of strategic planning, marketing, business development, as well as product and market development. He has engaged more than 100 vendors and over 400 end users. With over 28 years of marketing, sales and business experience in systems he's considered one of the leading experts in the data protection, storage, networking, server, & infrastructure markets.

marcstaimer@comcast.net

503-579-3763

Contents

Abstract	2
The Urgent Market Problems with Traditional Data Protection Technologies	2
Asigra Hybrid Cloud Backup	7
Conclusion	9

Why Cloud Backup/Recovery (Private, Public, or Hybrid) will be Your Data Protection.

Abstract

Data protection is a crucial operation required for both regulatory and corporate compliance. The ongoing paradigm shift to virtualized servers combined with escalating data protection requirements has made it exceedingly difficult for many IT organizations to keep up. From capturing the data being protected to testing and validating recoverability of that data, meeting recovery point (RPO) as well as recovery time objectives (RTO) have made data protection frequently a “crapshoot”.

Cloud backup is quickly becoming a very appealing data protection option for many IT organizations. Many have already made the move or are carefully considering moving data protection to a private, public, or hybrid cloud. A successful move to cloud backup requires ensuring data protection requirements are met while overcoming or mitigating the issues associated with traditional data protection methodologies. This paper examines these painful problems and how cloud backup technology overcomes them.

The Urgent Market Problems with Traditional Data Protection Technologies

Traditional data protection has always been problematic regardless of the technology.

Traditional Backup/Restore

The most common form of data protection is backup/restore software. It is notorious for failed backups, complicated management, frustrating agent implementations, upgrades, application disruptions, inability to meet backup windows, difficult to impossible restores, inflexibility, and high total costs of ownership. These problems can be directly tied to its anachronistic (designed for different times and technology) architecture.

IT starts with an admin or root privileged agent. The agent is installed on the operating system for backup. Another agent is installed on each structured data application (database, email, ERP, etc.) A separate agent is typically installed for additional functionality such as archiving and CDP (continuous data protection). Because of their admin/root privileges, each system must be rebooted upon agent installation or upgrade. This makes agents completely application disruptive forcing the data protection admin to schedule installs and upgrades for late nights, weekends, or holidays. That’s only part of the problems tied to these agents.




Figure 1: Agent Based Backup/Restore

Agents are also difficult to manage since each upgrade has to be pushed out by the admin and remember each upgrade also requires a reboot. Upgrades occur every time there is a major backup/restore software upgrade, application upgrade, and often when there are OS upgrades. This means there are several application disruptive agent upgrades every year.

The backup/restore agents also frustrate the security admin because they seriously compromise security. They require an open firewall port for each agent in a separate LAN segment or location. And the agents are designed to take instructions from their backup server software so they are always listening for instructions. This makes them a juicy tempting target for hackers and an ongoing headache for the security admin.

Combine these problems with the fact that each agent will require approximately 2 to 20% of each server's resources and a not very attractive picture materializes. But it gets worse. The emergence of server virtualization aggravates all of these backup/restore agent issues. Now instead of one or more agents per physical server, there are one or more agents per each virtual machine on the physical server. The net effect is a rapid burgeoning of these aggravating problems, while adding to these problems a serious IO bottleneck during backup because of excess oversubscription.

Those problems just scratch the surface is just the beginning of the problems. As data growth continues unabated traditional backup/restore is commonly incapable of completing backup jobs within the scheduled time. Recoveries on the other hand are usually a multi-layered exercise in frustration. This is true if what needs to be recovered is a file, mailbox, message, transaction, table, database, application, or volume. First the correct data on the right tape or tapes must be found. Then all of the relevant data must be restored in the correct order. Each step must be verified before the next step is completed. If all this sounds complicated, the description is understated. Rarely does traditional backup/restore address whether or not the desired data is recoverable at all.

To ensure traditional backup/restore recoverability requires testing on a periodic basis. Based on how painful it is to recover traditional backup/restore data, testing tends to be avoided or minimized increasing risk of a failed recovery when it's actually needed.

Then there's the ongoing problem of storing the data. Because traditional backup/restore tends to backup the same data over and over again, it stores a lot more data than needs to be stored. This is why expensive target storage deduplication has become so popular. Eventually that data has to be archived offline to tape or other media. When it is, the data must be rehydrated or undeduped making the amount of data stored in archive highly duplicated and far too expensive for data unlikely to be recovered.

But if those offline tapes are not encrypted, there are other aggravating and ultimately costly issues to be resolved. If a tape with personal or corporate data is lost or misplaced, then the IT organization will have to by law, inform the public and provide some form of identity protection for a period of time.

Server-to-Server and Server-to-Server via Proxy

Other forms of data protection such as server-to-server replication and server-to-server replication via proxy have similar (albeit slightly different) issues as traditional backup/restore. Traditional server-to-server data protection is essentially mirroring between servers. It doesn't protect against the rolling disaster, malware, human error, and data corruption. The newer server-to-server via proxy is functionally very similar to backup/restore with most of the same problems and recovery complexity. They have others as well including limited scalability.

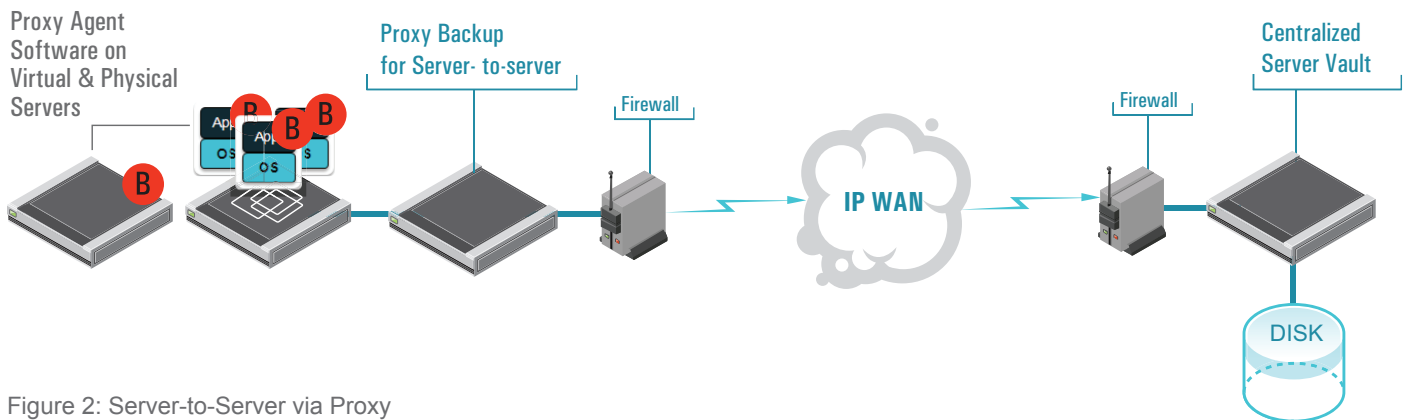


Figure 2: Server-to-Server via Proxy

Storage System Snapshot

Storage based snapshots lack structured data crash consistency (e.g. they are not application aware). In other words, it is very likely that a snapshot of a structured data application (databases, email server, ERP, etc.) will not be recoverable. Unless the administrator is testing the snapped data to check on recoverability on an ongoing basis, the admin will suffer from the consequences of Murphy's Law and discover that the data is not recoverable when they need it most.

Their workload immediately expands exponentially. The admin will attempt recovery on earlier snapshots until they find one that by luck is recoverable. For some structured data applications, the admin will be able to journal forward (a painstaking time consuming function) to bring the application up to date. For others, there will be unacceptable lost data. Either way, depending on luck for data recovery is a dicey proposition.

To solve this problem requires a work-around. The widespread answer is to put software (an agent) that works either with the storage system controller or a backup/restore server that works with the storage system controller. The agent quiesce the structure data application, flushes the cache, completes all writes to the storage, the index, and the metadata, then tells the storage system directly or through the backup/restore server to take the snap.

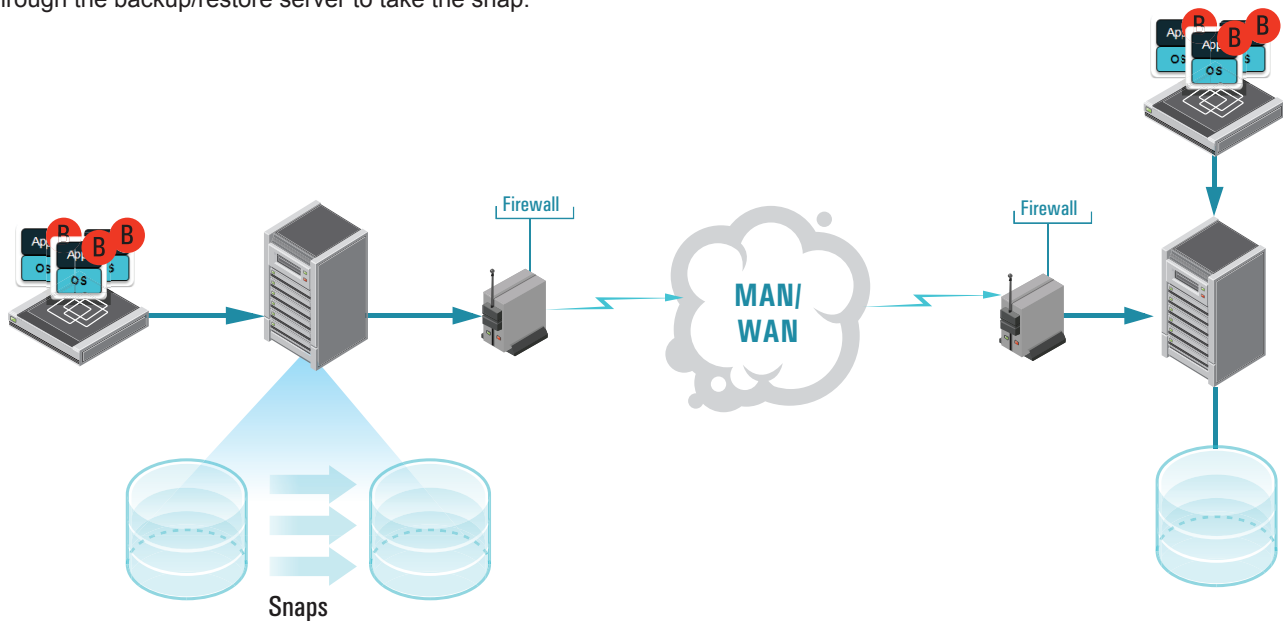


Figure 3: Snapshot plus Backup Agent

The agents, storage system snapshots, and backup/restore server all must be managed separately. Missteps or admin errors are widespread.

Hypervisor Snapshot

Hypervisor (server virtualization) snapshots also lack structured data crash consistency (e.g. they too are not application aware). There are exceptions. Most hypervisor snapshots are integrated with Microsoft VSS (Volume Shadow). For Microsoft Windows server structured data applications that are VSS enabled (SQL Server, Exchange, Oracle on Windows, etc.), VSS will quiesce the applications, flush the cache, complete the writes of the data, metadata, and index. But all other structured applications are not crash consistent with hypervisor snapshots and have the same recoverability issues as storage based snapshots.

There are other annoying issues with hypervisor snapshots. With the exception of Windows server, recovery is very coarse grain. Recovery is limited to the entire virtual machine image making partial recoveries of files, applications, transactions, mailboxes, etc. a major pain in the neck.

In the unique case of VMware, there is VMware Consolidated Backup (VCB). VCB is based on VMware snapshot and has the same severe limitations and issues. Unfortunately, VCB also requires a separate physical Windows Server to act as a proxy. This proxy is then backed up using traditional backup/restore technologies including an agent. The Windows proxy server is limited to a maximum of 32 concurrent data streams with 5 being VMware's recommended best practice. These additional limits makes either the backups very slow or significantly increases the number of proxy servers to implement, operate, and manage.

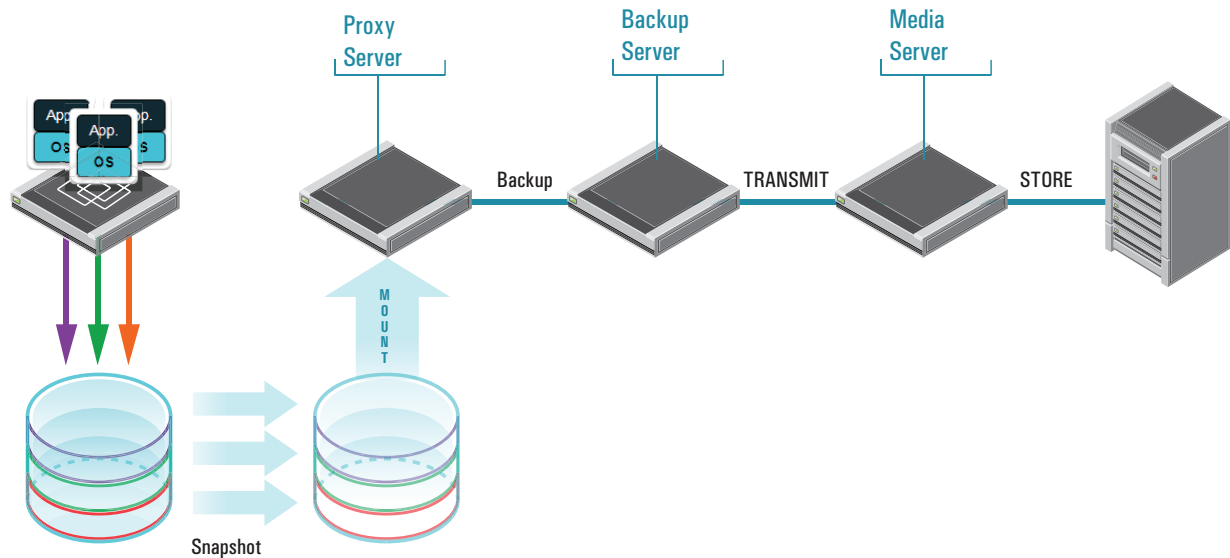


Figure 4: VMware Consolidated Backup (VCB)

There has to be a better way. Fortunately there is.

Private, Public, and Hybrid Cloud Backup/Recovery Requirements

For cloud backup/recovery to finally eliminate these painful data protection problems requires a host of new capabilities. It must provide:

- Effective agentless backup/recovery.
- Extensive support for virtualized servers, OS, structured data applications (databases, ERP, and email application).
- Data reduction technologies including incremental backup, local deduplication, global deduplication, and compression.
- Restore validation with autonomic data healing of “bad” data.
- User enabled simple online one-pass restores.
- Encrypted transmission and storage of protected data with simplified key management.
- Multi-tenancy.
- Flexible backup and recovery RPO & RTO granularity options.
- Automatically move protected data based on its age/value to correlated storage tier.
- Scalable backend and front-end architectures.
- Private and public cloud integratio.

Effective agentless backup/recovery

Effective agentless backup/recovery requires being able to provide the same or more data collection as agent-based backup/recovery, but without agents. This includes both visible and hidden files. And effective agentless backup/recovery cannot be application disruptive when implemented, operated, managed, or upgraded meaning everything is done online.

Extensive support for virtualized servers, OS, structured data applications (databases, ERP, and email)

Extensive support simply means that it supports whatever virtual server infrastructure, OS, and structured applications currently in use or with plans for future use.

By providing incremental backup, local deduplication, global deduplication, and compression, the least possible amount of data is stored. Less data stored equals less storage requirements along with reductions in both capital and operating expenses. It also

reduces the bandwidth requirements at remote offices branch offices as well as the bandwidth between the backup data vault and a second vault either within the organization or a managed service provider backup/restore cloud.

Restore validation with autonomic data healing of “bad” data

Provides the peace of mind that should there be a recovery event, that the data will be recovered. And it eliminates the massive headaches of actually having to run manual tests to see if the data is recoverable.

User enabled simple online one-pass restores

No backup/restore admin would be required. It would also make restores relatively simple “point-and-click” affairs. If a Microsoft Windows Exchange server goes down, there would be no need to first restore the Exchange server and then spend hours, days, weeks restoring the individual mailboxes and messages. One pass restores would accomplish this all at the same time.

Encrypted transmission and storage of protected data with simplified key management

Whether the data is local or remote it could not be sniffed nor intercepted. Once stored, it cannot be hacked or accessed except by authorized users.

Multi-tenancy

Supports multiple customers or clients in the same system but with hard firewalls between them. This ensures that each client can access their own data and only their own data. Multi-tenancy also requires some form of billing for what is actually being utilized.

Flexible backup and recovery RPO & RTO granularity options

Different application data have different values. As data value increases, so does its mission criticalness to the organization. High value data usually requires a finer grain RPO & RTO. Lower data value can have coarser grain RPO & RTO, which in turn has, lower costs. One size should not be forced to fit all.

Automatically move protected data based on its age/value to correlated storage tier

As data ages so does its value. The same is true for backup data. The older the backup, the less likely it is to be restored or accessed. There is no justification to leave lower valued data on high valued storage. Having a way to move the aged data to lower cost lower performance storage makes good sense.

Scalable backend and front-end architectures

A backup cloud has to be able to scale for both users and storage repositories. This is essential and critical in keeping complexity down, manageability up, capital expenditures down, operating expenditures down, and admin productivity up. Being able to increase performance or capacity “on-demand” is fundamental to cloud architecture.

Private and public cloud integration (hybrid cloud)

There should be two levels of private public cloud integration. The first allows up the last backup to reside at each office on local storage. Both the latest and all other incremental backups continue to reside in the centralized backup/recovery vault. The second allows for the organization to provide complete cloud backup/restore, and allow them to replicate their vault to a public cloud backup supplier for offsite backup/restore and archive.

These are the requirements for private, public, and hybrid clouds. The question is whether there is a software product or service that meets these requirements. The answer is yes. The company and software product that not only meets but also exceeds these requirements is Asigra’s Hybrid Cloud Backup™/Restore.

Asigra Hybrid Cloud Backup

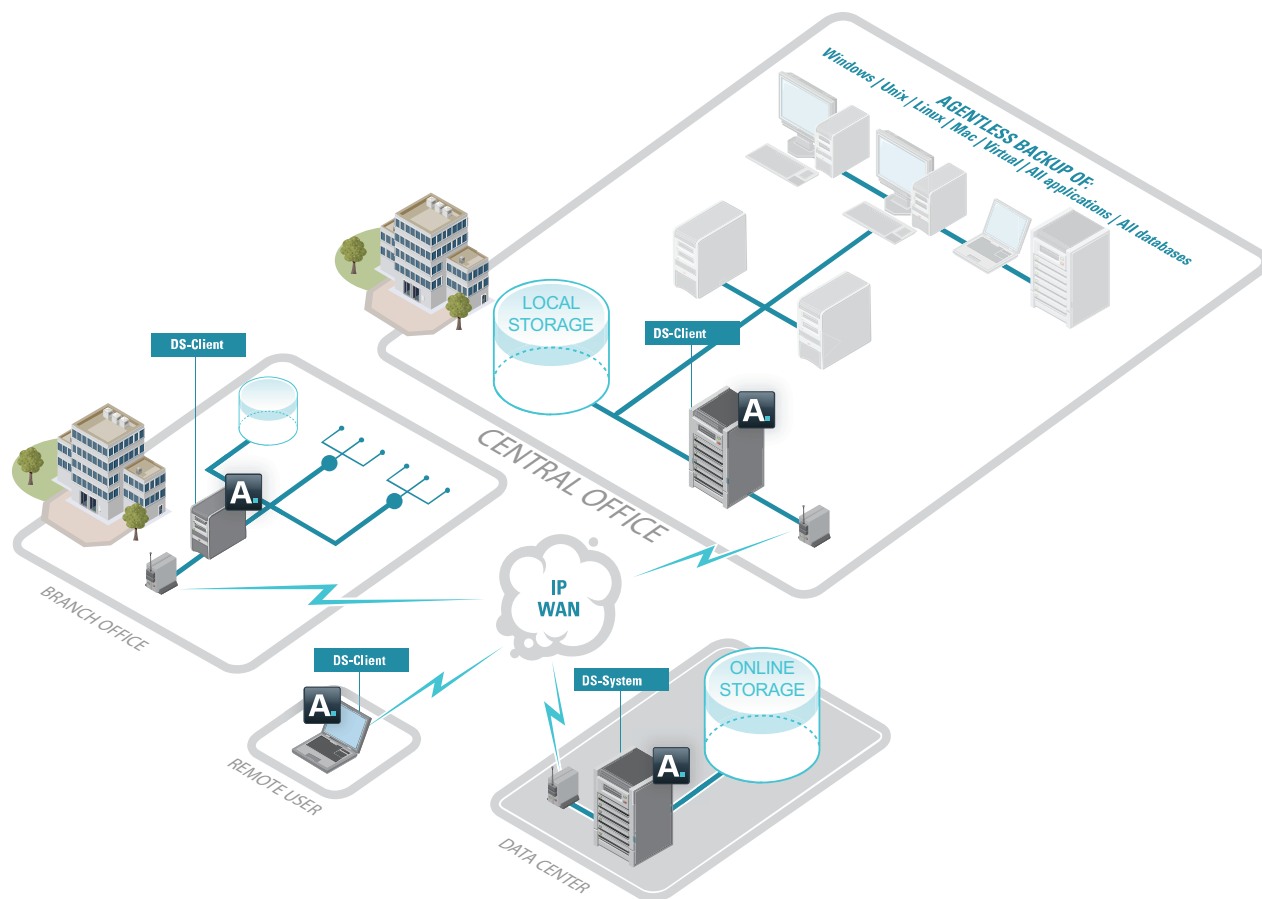


Figure 5: Asigra Hybrid Cloud™ Backup/Recovery Agentless Software

Asigra has been a leading provider of distributed backup software for 23 years delivering the first multi-tenant public cloud software in 1986. Asigra has continually produced an impressive string of cloud backup/restore firsts. Some of which include: the first to provide public cloud incremental forever, local and global deduplication with additional compression in 1993; the first to provide a public cloud based service oriented architecture, encryption in flight, and encryption at rest in 1995; the first to provide autonomic healing of backed up data ensuring recoverability in 2003; the first to provide public and private agentless continuous data protection (CDP), agentless VMware backup/restore/CDP, and cradle to grave backup life cycle management and archive in 2006; the first to provide public and private 64bit agentless backup/restore/CDP in 2007; and the first to provide public/private cloud hybrid integration, on-demand backup load distribution; FIPS 140-2 certification, and automated storage migration plus retirement.

Asigra's Hybrid Cloud Backup/Restore meets and exceeds all of the market requirements.

Effective agentless backup/recovery

Asigra invented agentless backup/recovery and have been perfecting it for over 23 years. There is nothing that an agent can do that Asigra agentless cannot do.

Extensive support for virtualized servers, OS, structured data applications (databases, ERP, and email application)

Asigra's Hybrid Cloud Backup/Restore OS and structured data application support is one of the most extensive in the industry. It includes: Windows (2000, 2003, 2008), Linux, AIX, Solaris (x86 and SPARC), HP UX, Tru64 UNIX, BSD, Mac OSX, System I, NetWare, VMware ESX 3/3.5/3i, VMware vSphere4, Virtual Iron, XenServer, Parallels, Hyper-V, SQL Server, Oracle, MySQL, PostgreSQL, Exchange, Notes, GroupWise, SAP, and more.

Data reduction technologies including incremental backup, local deduplication, global deduplication, and compression

Asigra invented data reduction for backup/recovery. Asigra's Hybrid Cloud Backup/Restore captures the least amount of data at the source with its incremental forever and time based versioning; then dedupes locally, compresses, transmitting the least amount of data possible to the backup cloud vault; finally dedupes globally from all sources before it stores the least amount of data possible on disk.

Restore validation with autonomic data healing of "bad" data

Asigra's automated autonomic healing tests the backed up data for recoverability when first stored. If it is not recoverable the corrupted data is again backed up and tested. Asigra has currently the only backup/recovery software providing it. The restore validation capability gives the backup vault admin the ability to simply and easily test any or all data set recoverability without application disruption or hassle. It is a belt and suspenders approach because in the end, it's all about recoverability.

User enabled simple online one-pass restores

Asigra enables users to recover their own data. This is essential in a multi-tenant environment.

Encrypted transmission and storage of protected data with simplified key management

After Asigra collects, dedupes, and compresses the data, it encrypts data up to AES 256 and is US Federal Government certified FIPS 140-2. It also manages the user encryption keys.

Multi-tenancy

Asigra's Hybrid Cloud Backup/Restore software has been multi-tenant for over 23 years. It was designed for multi-tenants (whether the tenants are internal departments or external clients). No tenant can ever see or recover another tenant's data. It is just not possible. Whereas the licensing, bill back, or charge back are built into the software management.

Flexible backup and recovery RPO & RTO granularity options

Asigra's Hybrid Cloud Backup/Restore software offers a complete range of granularity from CDP level fine grain (continuous backups) all the way up to coarse grain and everything in between. It is application and server configurable. Its CDP function is designed for both the data center and ROBO (remote office branch office). It is unique in that the CDP requires no additional licensing and for remote offices will provide the level of granularity that the bandwidth can support.

Recovery is one pass not multi-step. Granularity is entirely up to the user. It ranges from file, table, message, and mailbox, to complete applications, databases, or volumes. In the event of a disaster, many Asigra based public backup/restore clouds will actually ship a disk subsystem to the client.

Automatically move protected data based on its age/value to correlated storage tier

The Hybrid Cloud Backup automatically moves data based on user determined time value (age) from online backup/recovery, to lower cost nearline (2-step recovery), to even lower cost offline (2-step recovery) archive. Each step of the process is automated and reduces the software licensing as well.

Scalable backend and front-end architectures

One of the basic concepts of cloud computing and cloud storage is that it can linearly scale both up front and in the backend seamlessly. Asigra's Hybrid Cloud Backup/Restore maintains that lineage. Whether it is the front-end data collection where Asigra has on-demand data collection where multiple backup servers (called DS-Clients) share the load based on policies; or the backend where the backup vault itself (called DS-System) utilizes grid technology and extensible storage to provide a peerless scalable vault.

The backend DS-System Vault also automatically migrates data as storage systems are replaced.

Why Cloud Backup/Recovery WILL be Your Data Protection

Private and public cloud integration (hybrid cloud)

Asigra's Hybrid Cloud Backup/Restore software can be deployed as a private cloud, public cloud, or as the name implies, uniquely as a hybrid. This comes from Asigra's ability to replicate backup vaults to another site ongoing incremental, deduped, compressed, encrypted.

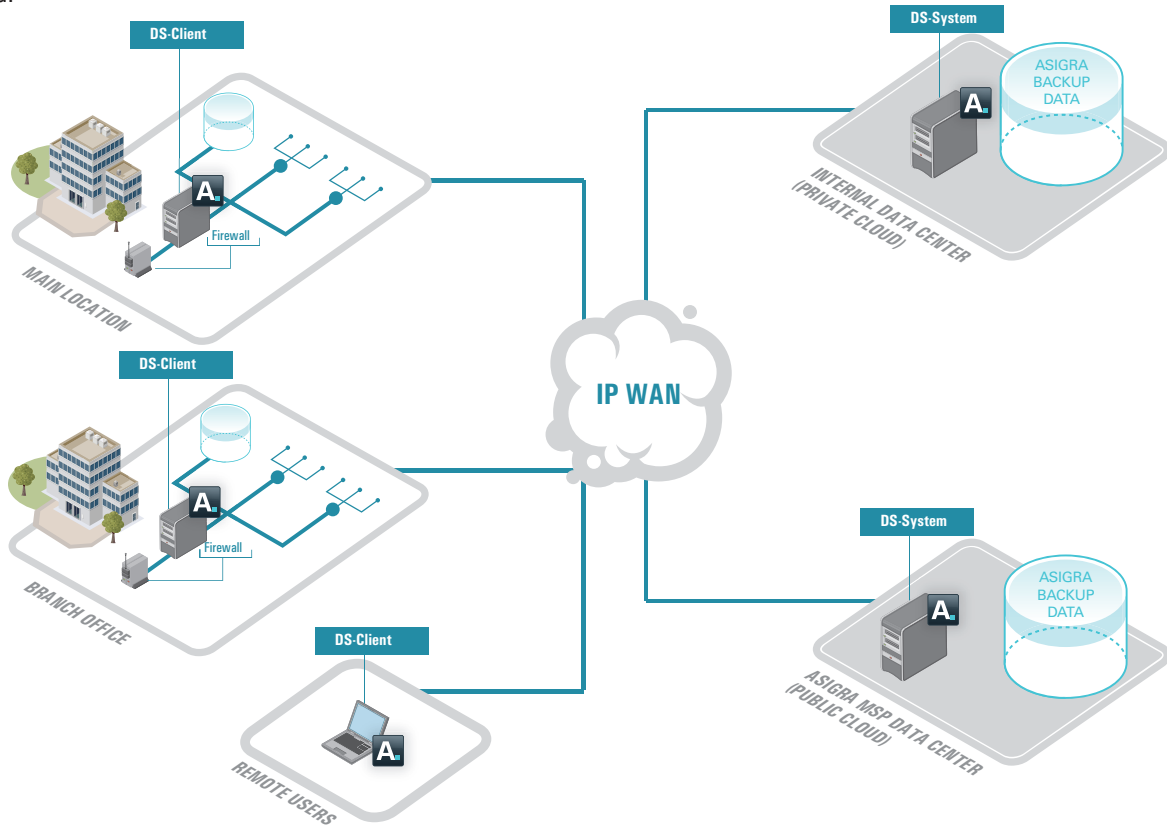


Figure 5: Asigra Hybrid Cloud™ Backup/Recovery in Private-Public Hybrid Mode

The public cloud architecture also has another optional hybrid capability by providing the last backup on local storage. Figure 6: Asigra Hybrid Cloud™ Backup/Recovery in Private-Public Hybrid Mode.

Conclusion

In an era of increasing regulatory and corporate compliance traditional data protection strategies and products have proven to be painfully frustrating. They're complicated, time-consuming, and error prone; leading many IT organizations to a series of mixed overlapping solutions. The hope is that this complex soup of data protection will allow the organization to protect and if necessary recover their data. Regrettably, that hope is far too often unrealized.

The solution is private, public, and hybrid cloud based backup/recovery. It eliminates the complexity and frustration while actually meeting or exceeding IT organization data protection requirements.

The one solution that meets the requirements of private, public, and hybrid cloud based backup/recovery and goes above and beyond, is Asigra's Hybrid Cloud Backup™/Restore.

Get on Cloud 9 with Asigra v9.

About Asigra.

For more than 20 years Asigra has stayed ahead of the market with a secure, agentless, scalable and automated backup and recovery solution that aligns the value of data with its storage costs. IT Leaders evolve their environment, without being constrained by their recurring backup challenges, with innovative solutions from Asigra that currently protect over 100,000 sites worldwide.

Tel: 416.736.8111 Fax: 416.736.7120 Email: info@asigra.com

RecoverYourCool.com

Asigra.