

Modern data security and management topologies: A guide for IT leaders

Blueprints and best practices to reduce risk and strengthen business resilience



TABLE OF CONTENTS

Introduction	3	Basic	9
Key design factors	4	What is a cyber vault?	9
The 3:2:1 rule still applies	4	Enhanced topologies (including adoption by industry)	14
Your requirements may vary, but commonalities exist	5	Mission Critical	19
Blueprints for modern data security and management: Types and topologies	7	What is your Minimum Viable Company?	19
Blueprints: The complete list of topologies	8	Conclusions and next steps	23
		About Cohesity	25

Introduction

Three factors are driving the need for a fresh approach to data security and management. First is the imperative for digital transformation and API-driven infrastructure. IT leaders are modernizing every aspect of their IT estate to support greater automation, extensibility, cloud scale, software-defined architectures, and “shift left” security principles.

Second, the cyber threat landscape is evolving in complex, unpredictable ways. The existing data estates of many organizations are siloed, resulting in greater operational risk from a cyberattack. A recent survey indicates that 32% of organizations believe that rapid recovery times are hindered by antiquated backup and recovery systems. Similarly, 34% say that a lack of integration between IT and security teams also prolongs recovery times.

And third, the advent of AI has caused leaders to search for modern data platforms that make enterprise data accessible to generative AI technologies.

Every IT executive leading a data modernization project can benefit by “standing on the shoulders of giants.” In this white paper, we describe the most important design considerations, and how to arrive at the best data resilience approach given common enterprise requirements. Our experience across thousands of deployments and our knowledge of best practices have informed our perspectives.

The guidance that follows is meant to be vendor agnostic, but we’ll use Cohesity brand names for simplicity.

Key design factors

Every worthwhile modernization project carries risk. This certainly applies when transforming enterprise data security and management processes and tooling. But there's good news: you can build from the work others have done before you. Many organizations have successfully modernized their data estate, and we've cataloged these best practices in this paper.

IT and cybersecurity leaders are tasked with balancing agility, risk, and cost across their IT estate. This estate is dynamic, and now encompasses data and apps running in on-premises data centers, public clouds, colocation facilities, and edge locations. Modernization efforts grow in scale and scope over time as the sheer volume of apps, data, and data sources compounds.

Modernizing the enterprise data estate is further complicated by:

- Diverse infrastructure targets in different locations and multiple workloads, resulting in data fragmentation and inefficient backup and recovery processes
- Lack of sufficient IT and cybersecurity skills within most organizations
- A rapidly changing cybersecurity landscape, with hundreds of attacks every minute. The evolution and sophistication of attacks makes early identification critically important.

That said, there are some “first principles” to keep in mind.

The 3:2:1 rule still applies

The 3:2:1 rule states that you should keep at least three copies of data, that these backups should be stored on two different types of media or platforms, and that at least one of the copies should be kept offsite.

The enduring value of the 3:2:1 rule stems from three concepts that continue to drive system topology design in the cloud-native era:

- Business requirements
- Failure domains
- Acts of God

We're going to describe each of these concepts in detail. They've always been key focuses in the industry and remain so today—especially with intensifying concerns around cyberattacks. Cyberattacks weren't always front and center, but they certainly are now.

We'll define how these three concepts have informed backup and recovery designs over time and discuss how concerns around cyberattacks are forcing our customers to ask themselves if their existing deployment topologies are still sufficient (or not).

Let's take each one at a time:

Business requirements

Companies must adhere to a broad set of business, regulatory, and compliance requirements. Many of these requirements drive the need for the retention of both current and past copies of data. For example, a compliance

team might need to pull a three year old contract to respond to a request from an industry regulator. Or a tax team may need to recover files for an ongoing audit. Or, on a more relatable note, perhaps a critical file was accidentally deleted and needs to be restored.

Failure domains

It's well known in the IT industry that both software and hardware will fail. Hardware and software vendors make extraordinary efforts to design around this inevitability, yet failure persists. IT teams must plan for failure, and ensure that these failures don't adversely impact the organization or the business. Examples of failures include workload failures, such as VM or storage volume corruption, or an OS patch implementation that failed. In both cases, the IT team needs to recover from the failure, and as part of that recovery, they'll likely require data from their backup and recovery system.

Decades ago, leaders didn't have to fret about failure due to bad actors. Today, cyberattacks are not only a key driver of system failure but perhaps the most highly visible failure cause—and one that captures the attention of the board of directors.

Acts of God

This term refers to natural disasters or other events that are outside human control and can't be foreseen or prevented by reasonable means. Adverse events such as fires, earthquakes, floods, and cable cuts can all be considered "acts of God." As with failure domains, IT organizations need to consider the potential for acts of God and build design systems that are resilient to them.

The 3:2:1 rule offers practical guidance in the face of acts of God.

Systems fail, so it's sensible to have multiple copies of the data available. Because of failure domains, it's also reasonable to keep the backups on two different media types or systems. Finally, due to acts of God, it's

responsible governance to keep at least one of these copies at a remote location, perhaps as part of a disaster recovery site or remote data center.

Your requirements may vary, but commonalities exist

While the 3:2:1 rule is a solid practice, organizations may opt for deployments with fewer than three copies. Others adhere more strictly to the 3:2:1 rule. Still others keep more than three copies. (We'll explain the rationale for these design choices later in this paper.)

Regarding design, we've grouped the deployment topologies ("blueprints") into three types.


Type	Description
Basic	A deployment topology with two or fewer copies
Enhanced	A deployment topology with three copies
Mission critical	A deployment topology with four or more copies

Familiar availability architectures are still relevant today

Let's move on to a term seasoned IT leaders and practitioners will find more familiar: availability architecture.


An availability architecture describes how an IT team may organize its hardware and software systems so they're resilient to a potential outage. In our experience, most enterprise deployments consist of one of three customer availability architectures: **active-active**, **active-standby**, and **hub and spoke**.

Each of these approaches is shown on the next page.




Active-Standby

Workloads operate out of a single data center. Often there will be a standby disaster recovery site that takes over in the event of an outage.



Active-Active

Primary workloads are split across two data centers. In the event of a data center failure, the remaining data center can take over the entire load.



Hub and Spoke

Workloads are characterized by a large set of remote/branch offices that are then connected to a single data center.

Fig. 1: Customer availability architectures

In each case, the availability architecture is designed to ensure that business operations can continue in the event of a failure. When the “active” side of an **active-standby** architecture has a failure, all processing switches over to the other standby system. When either side of an **active-active** system fails, the other side takes on 100% of the workload. And when there’s a failure in a **hub and spoke** system, such as when data is destroyed by a hardware failure or other issue in a branch office, the system and its data can be restored from the hub once the cause of the failure is resolved.

Responsible IT teams have designed these availability architectures and routinely test them to ensure that the resilience mechanisms work properly.

Our best practice set of blueprints is informed by these proven availability architectures. At Cohesity, we don’t deploy backup and recovery systems in a vacuum. We design them to align closely with the underlying IT infrastructure. You’ll see this in detail when we cover various backup and recovery topologies, and how those topologies map to IT availability architectures, later in this paper.

Blueprints for modern data security and management: Types and topologies

We've already discussed basic, enhanced, and mission-critical backup types. Now we'll introduce you to the topologies associated with these types.

But first, some definitions. Data security and management systems can store data via different methods. Copies of data can be kept as **backups, replicas, or archives**.

Some rules of thumb follow:

- **Backups** are formed from a primary copy, and result in deduplicated, compressed, and encrypted data. This processing is performed once on the data, and then the processed backup data can be copied to a replica or an archive.
- **Replicas** are generally used for short-term retention, typically months but not years. These replicas often support IT availability architectures such as active-active or active-standby.
- **Archives** are generally used for longer-term retention and are often kept for years. Archives are often used for compliance and regulatory purposes, but in some cases they're used for IT availability architectures as well.
- **Restoration from a backup or replica** is a one-step process and is faster than restoration from an archive, which is a two-step process. The archive needs to be downloaded into the vendor backup and recovery system before it's restored into the IT system.
- **Restoring from a ransomware attack** is complicated by the need to restore to a clean, uninfected copy that likely won't be the most recent backup. An organization can't simply restore from a ransomware attack in the same way it restores from a domain failure or outage caused by an act of God. The impacted organization will have to analyze the environment, and ensure that the copies being recovered are free from the malware infection that caused the attack in the first place.

The list of topologies appears in the table below. Note that every topology has a single backup where the data is deduplicated, compressed, and encrypted. In addition to the backup, the various topologies can keep one or more replicas and one or more archives as well. We've given a descriptor (B1, B2, E1, E2, M1, M2, etc.) to each topology type to help keep track of them.

Type	Copies	Topology / Type of Copies
Basic	1	B1 - Backup
	2	B2 - Backup & Archive
	2	B3 - Backup & Replication
Enhanced	3	E1 - Backup, Replication & Archive
	3	E2 - Backup & Dual Replication
Mission Critical	4	M1 - Backup & Dual Replication with Archive
	4	M2 - Backup, Replication & Dual Archive
	5	M3 - Backup, Dual Replication & Dual Archive

A type / topology combination is defined by **both the number of copies and the nature of those copies**. For example, an Enhanced type, which always includes three copies, can have two distinct topologies. One topology is backup, replication, and archive, while the other is backup and dual replication. Note that not every permutation of data copies is listed above; this list simply represents the most common deployments. Some combinations simply don't make business or technical sense.

It's helpful to understand which topologies are commonly used and which aren't. The relative popularity of each pattern can offer useful "upgrade" paths as IT teams consider additional protection for their data estate.

Blueprints: The complete list of topologies

Now that we've described types, topologies, and customer availability architectures, let's lay out the full list of industry blueprints.

This chart ties together all the concepts we discussed previously. All configurations represent a popular choice in the real world, with real enterprises operating at scale with Cohesity.

Type	Topology / Type of Copies	Customer Availability Architecture		
		Active-Standby	Active-Active	Hub & Spoke
Basic	B1 Backup (1)	•	•	
	B2 Backup & Archive (2)	•	•	
	B3 Backup & Replication (2)	•	•	
Enhanced	E1 Backup, Replication & Archive (3)	•	• •	•
	E2 Backup & Dual Replication (3)	•		•
Mission Critical	M1 Backup & Dual Replication with Archive (4)	•		
	M2 Backup, Replication & Dual Archive (4)	•	•	
	M3 Backup, Dual Replication & Dual Archive (5)			•

One caveat on the mission critical topologies. These topologies have either been deployed, demonstrated, trialed, or discussed at length with enterprises or other large customers. We noted earlier that many customers seek greater resilience to their deployments. The discussions entail practical ways to protect additional copies of data. A common solution is to enhance the deployment with a cyber vault. For that reason, many of the mission critical topologies contain an archive configured as a cyber vault. (We offer Cohesity FortKnox for this scenario.) Many topologies can be made more cyber resilient with the addition of a cyber vault.

We'll now discuss the types and topologies in detail one group at a time.

Basic

Topology	Primary Data Center	Active-Active
B1 - Backup	✓	✓
B2 - Backup & Archive	✓	✓
B3 - Backup & Replication	✓	✓

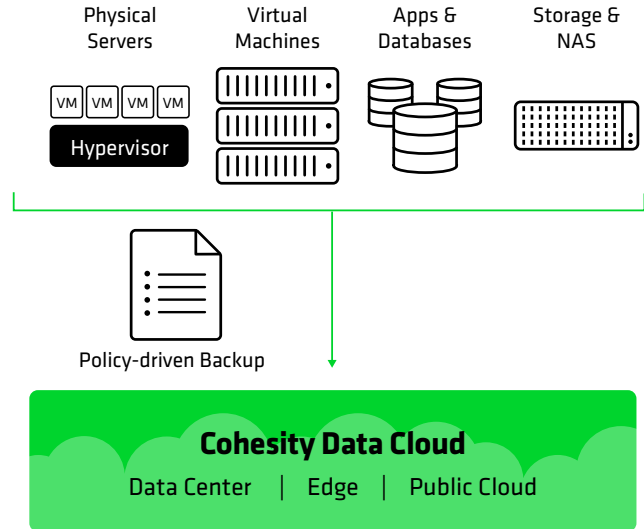
Basic is popular with low and moderate value data, or in cases where the customer already has multiple copies of their data. Basic topologies are used when there's a single, primary data center—and for active-active approaches as well.

What is a cyber vault?

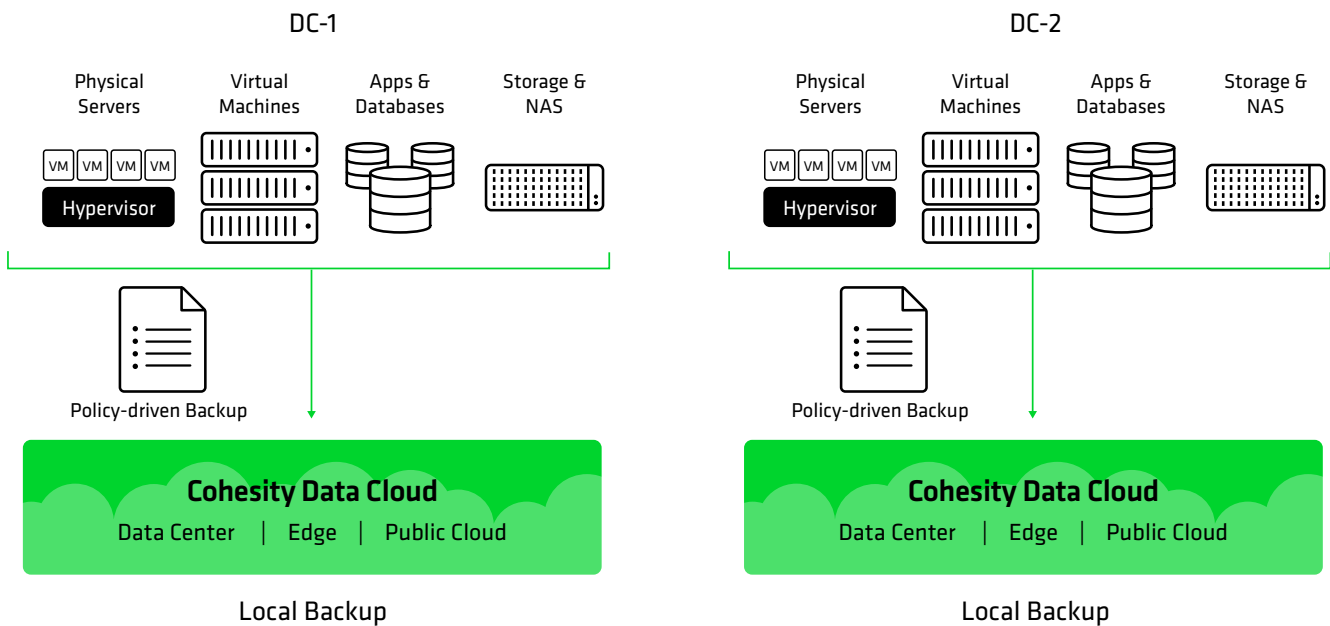
A cyber vault stores an isolated copy of production data, often offsite. With a clean, separate, and protected copy of data always on standby, organizations can rapidly recover data back to its original source, or to alternate backup locations, in case of a ransomware attack or other incident that compromises production or primary backup systems. A modern cyber vault strategy uses “virtual air gap” technology that protects backups but allows for temporary network connections to enable necessary remote access—albeit with very strong controls—while further isolating data with the cloud as needed. A well-designed cyber vault can be an effective part of a robust data isolation and cyber resilience strategy.

Basic: B1 - Local Backup

This is a bare bones approach, with only a single backup copy of the data. Many IT data centers still use this approach, though it has no provision for either disaster recovery or long-term retention of backup data. This approach is typically used for low value data.



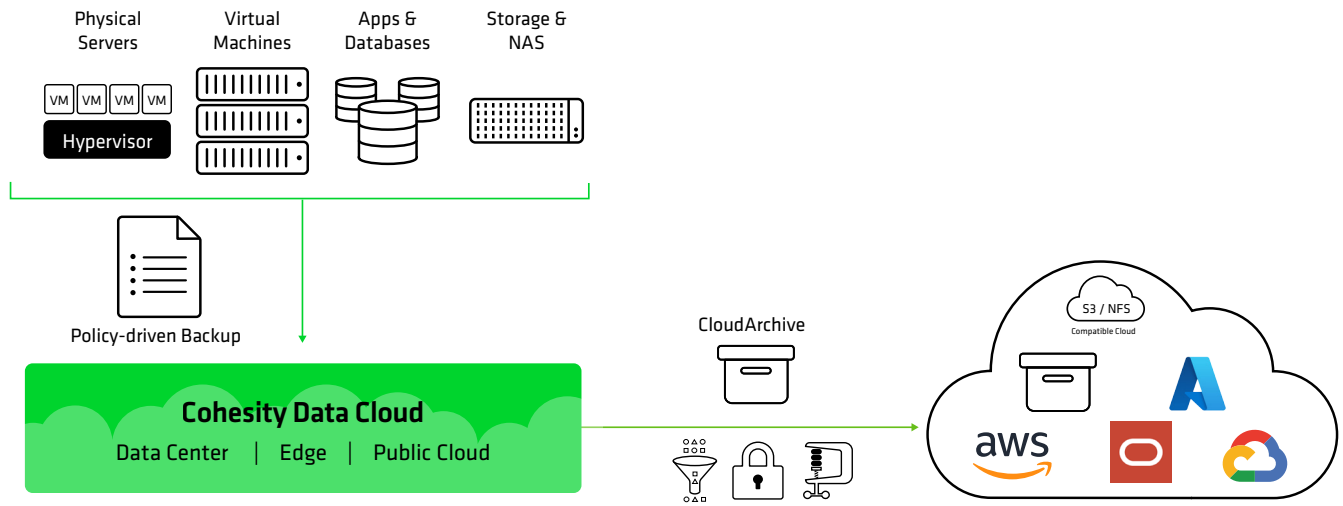
Basic: B1 - Backup (Active-Active)



Many types and topologies lend themselves to an active-active approach. An active-active approach is effectively two of a single topology type, mirrored and back-to-back. In this topology, we have a pair of active-active data centers, each with its own backup. Either data center can take over for the other in an outage. Each data center also has a

complete backup of its data and workloads. For customers with a large amount of available WAN bandwidth, even the backup can be geographically separated from the data center to provide extra disaster prevention. All replication of workloads and data occurs at the workload layer, hence this topology doesn't require replication.

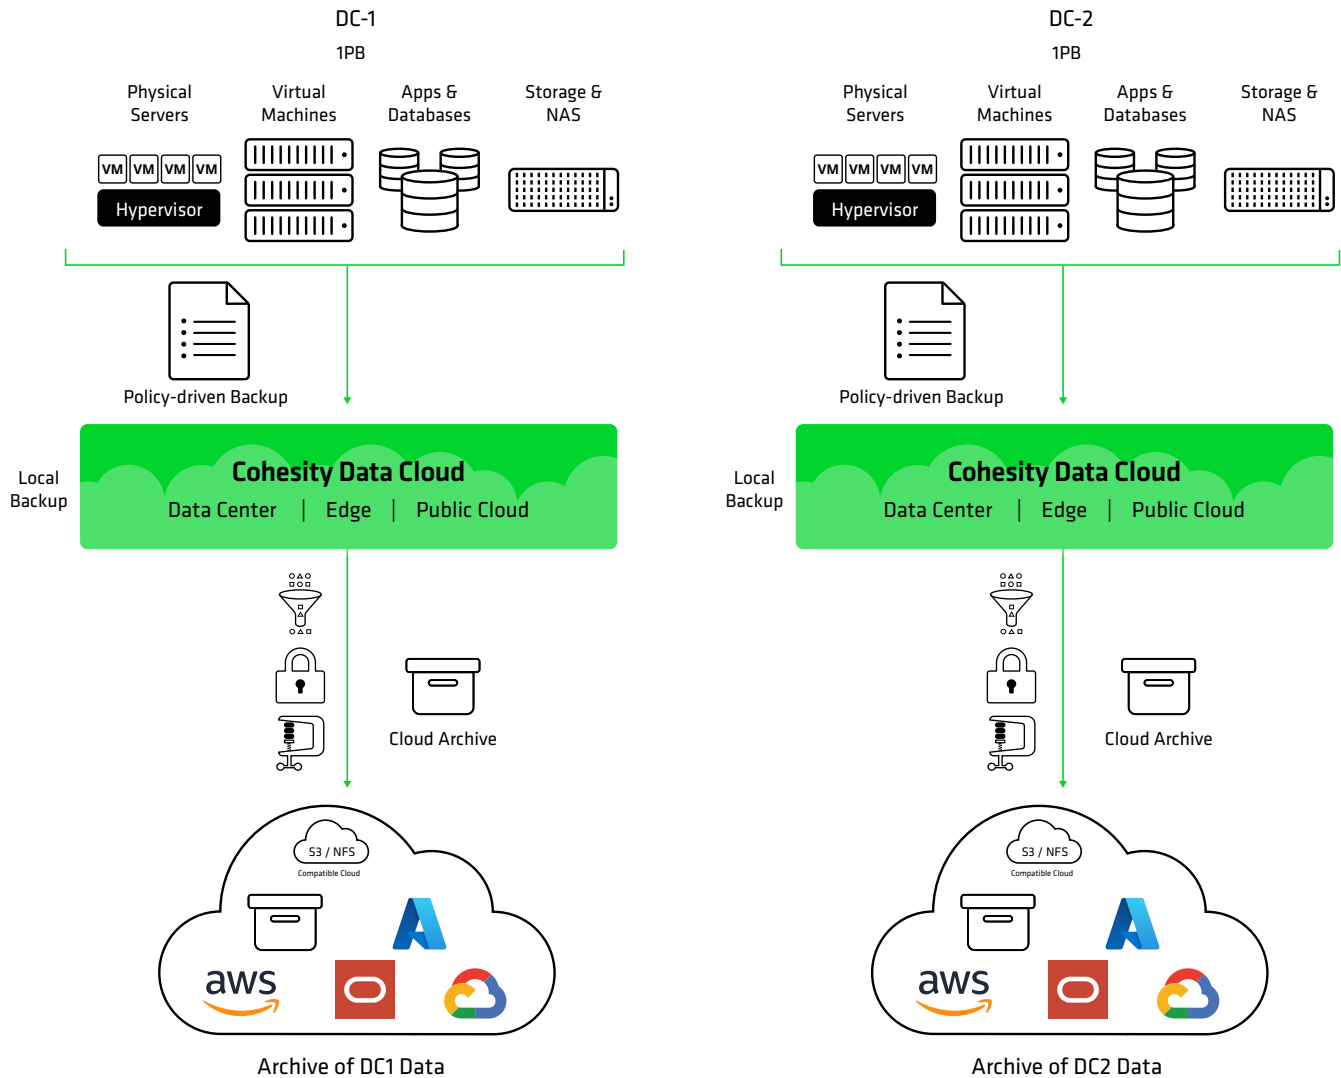
Basic: B2 - Backup & Archive



In this case, backup is coupled with an archive with a much longer retention period than the backup. Cohesity FortKnox is a popular choice here, as it provides added security for this topology. FortKnox is an isolated archive and is only connected when a write to the archive or a restore from

the archive is being performed. The archive can also be to an on/off-premises private cloud or a public cloud such as AWS, Google Cloud, Microsoft Azure, Oracle Cloud, or any S3/NFS compatible cloud service.

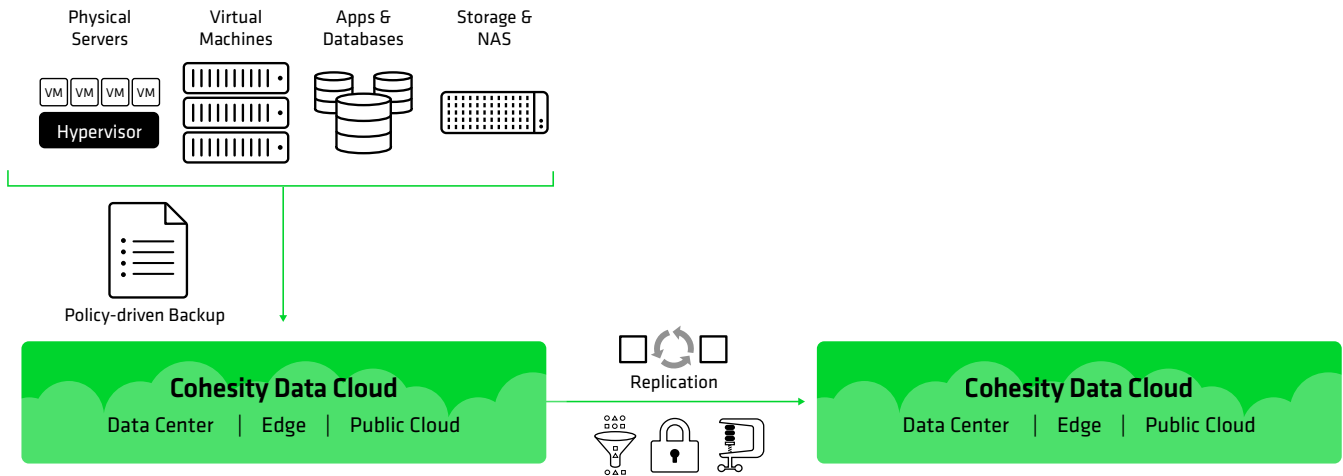
Basic: B2 - Backup & Archive (Active-Active)



In this topology, we have a pair of active-active data centers, each with its own backup and archive. Either data center can take over for the other in case of an outage, and each data center also has a complete backup of its data and workloads through the archive. All replication of

workloads and data occurs at the workload layer, which is why this topology doesn't require replication. FortKnox would be a useful choice for the archive given its isolation and added security.

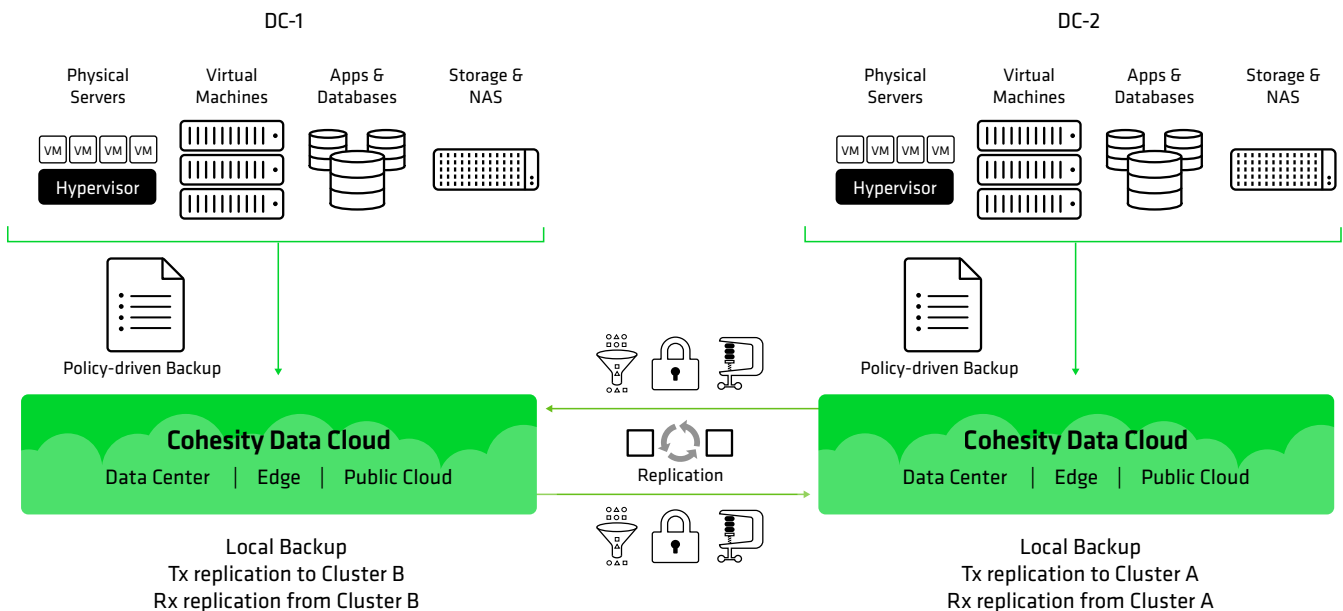
Basic: B3 - Backup & Replication (Disaster Recovery)



In this topology, the backup and replica are roughly aligned on retention period. The replica is geographically distributed and used for disaster recovery. For customers with a large amount of available WAN bandwidth, even the backup can be geographically separated from the data

center to provide extra disaster prevention. The focus of this topology is on business continuity in case the backup is unavailable. Replicas can directly restore data without requiring the two-step process that's needed when an archive is used.

Basic: B3 - Backup & Cross Replication (Active-Active)



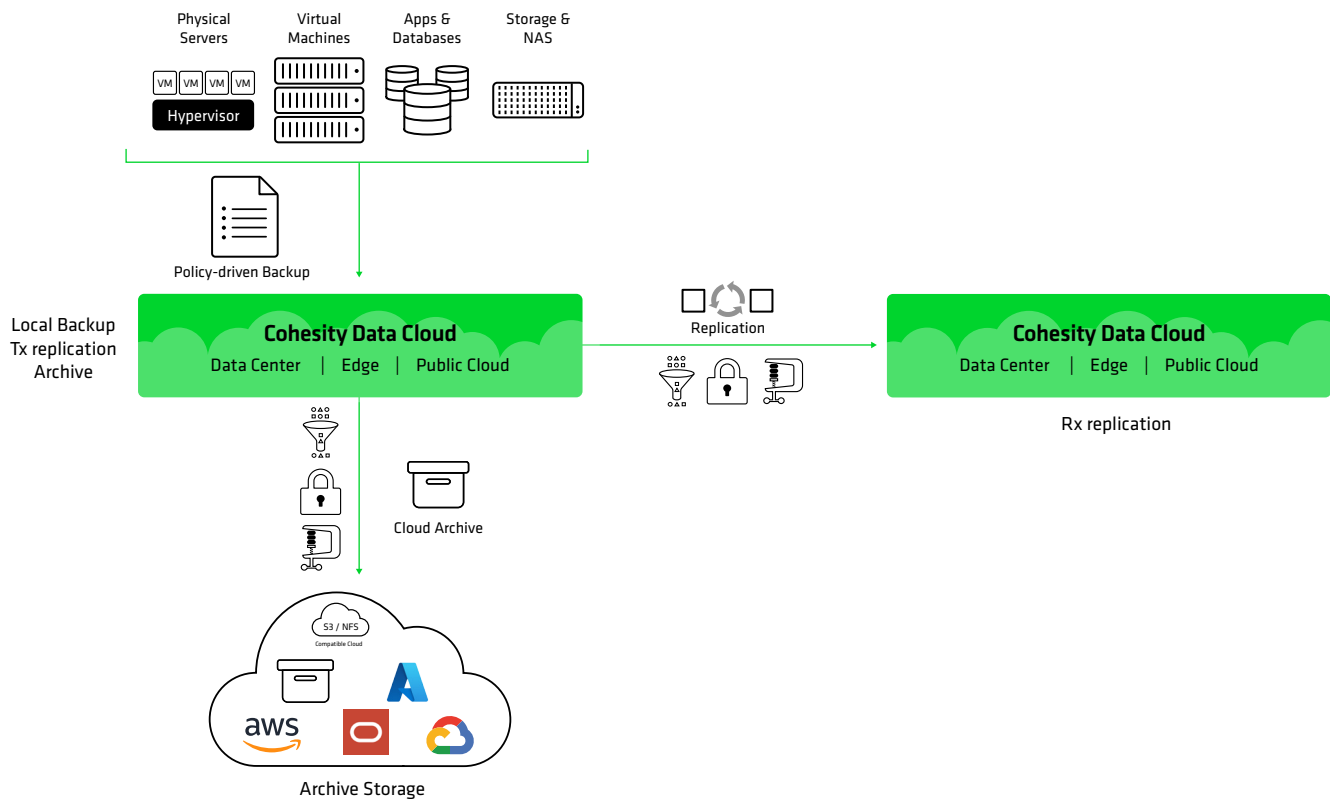
In this case, backup and replication clusters are cross replicated. The backup for the first site is the replica for the second site, and vice-versa.

Enhanced topologies (including adoption by industry)

Enhanced topologies are popular with high value data. Backup, replication, and archive (E1) is the most popular, while backup and dual replication (E2) less so. The topologies in common usage are marked below, along with noteworthy favorites by industry.

Topology	Primary Data Center	Active-Active	Hub & Spoke
E1 - Backup, Replication, and Archive	✓ All types	✓ Financial Institutions	
E2 - Backup & Dual Replication	✓ Government agencies	✓ Retail chains, some government agencies using an east-west model	

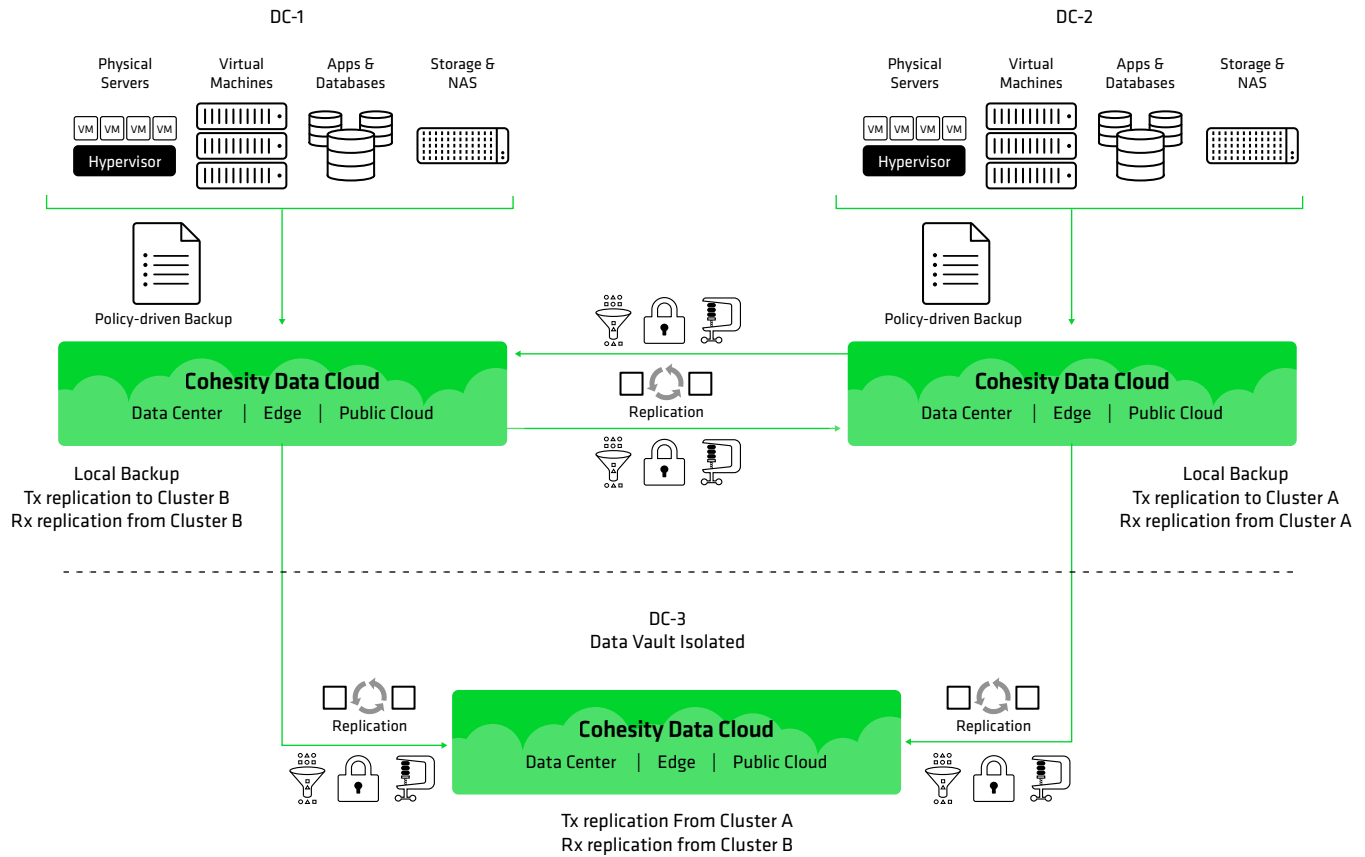
Enhanced: E1 - Backup, Replication & Archive



Backup and replica are roughly aligned (e.g. twice daily for 90 days), while the archive can cover months or even years. Recovery from a backup or replica is a one-step process, while recovery from archive recovers two steps: a read from the archive, and then a restore of the data. The archive can be to an on/off-prem private cloud or to a public cloud

such as AWS, Google Cloud, Microsoft Azure, Oracle Cloud, or any S3/NFS compatible cloud service. FortKnox would also be an excellent choice for this topology given its added isolation and security. Note that with the Cohesity Data Cloud, an archive can be restored via either a primary or secondary cluster.

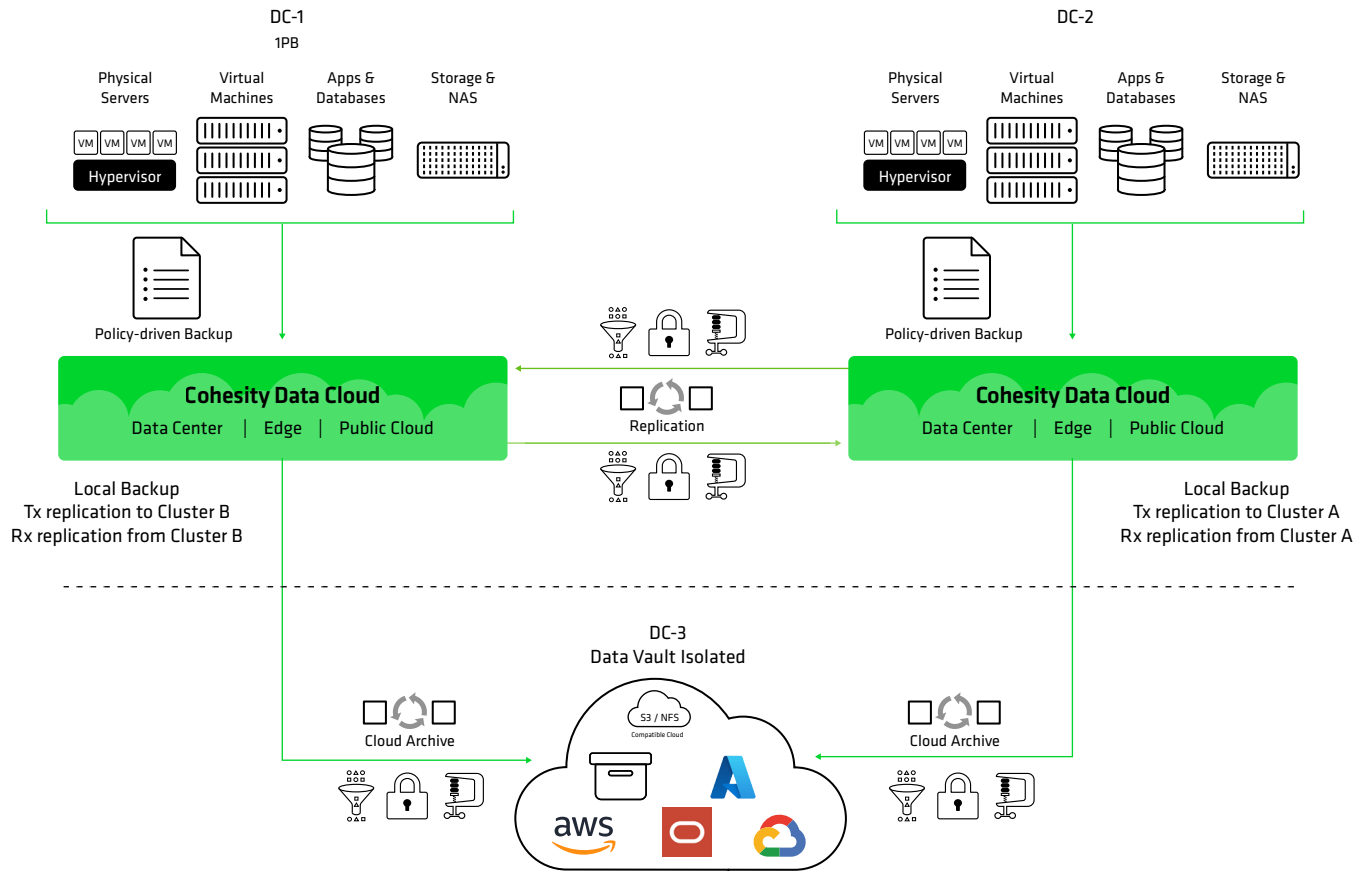
Enhanced: E1 - Active-Active with a Data Vault



This topology is another **active-active** approach. Here, the cross-replicated data centers share a single isolated data vault. The isolation is physical, with the data vault disconnected from the network when not in use. Note that

the vault is a replica, allowing for a one-step recovery of either data center from the replica. This architecture can be extended as well, with multiple active-active pairs all using the same data vault.

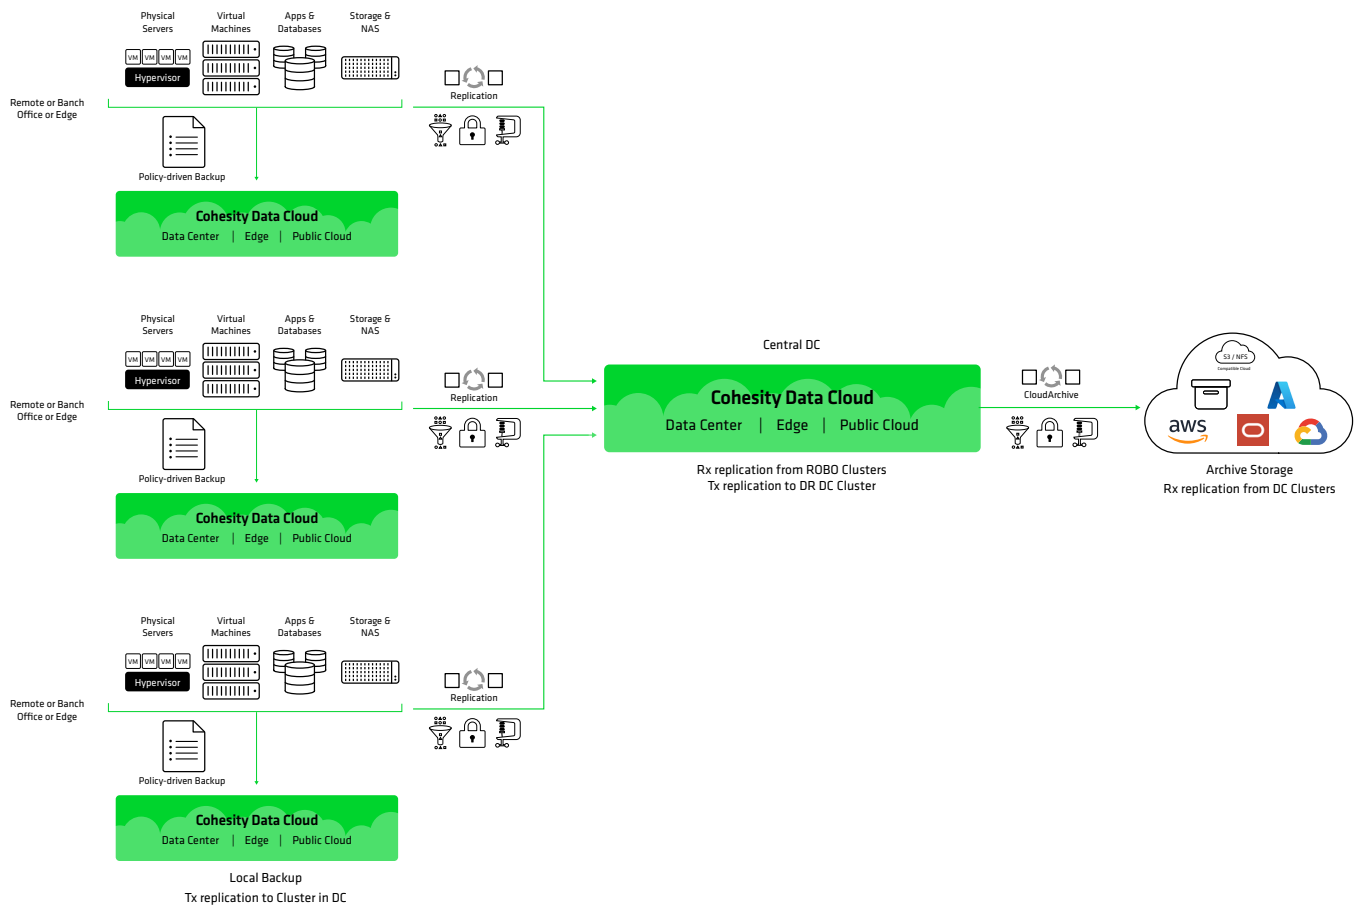
Enhanced: E1 - Active-Active with an Isolated Archive



The graphic above is another **active-active** approach. In this case, the cross replicated data centers shared a single isolated Archive. This use case would work well with our FortKnox archive approach given its isolation and added

security. This architecture can be extended as well, with multiple active active pairs all using the same isolated archive. With the Cohesity Data Cloud, an archive can be restored to any of the backups or replicas.

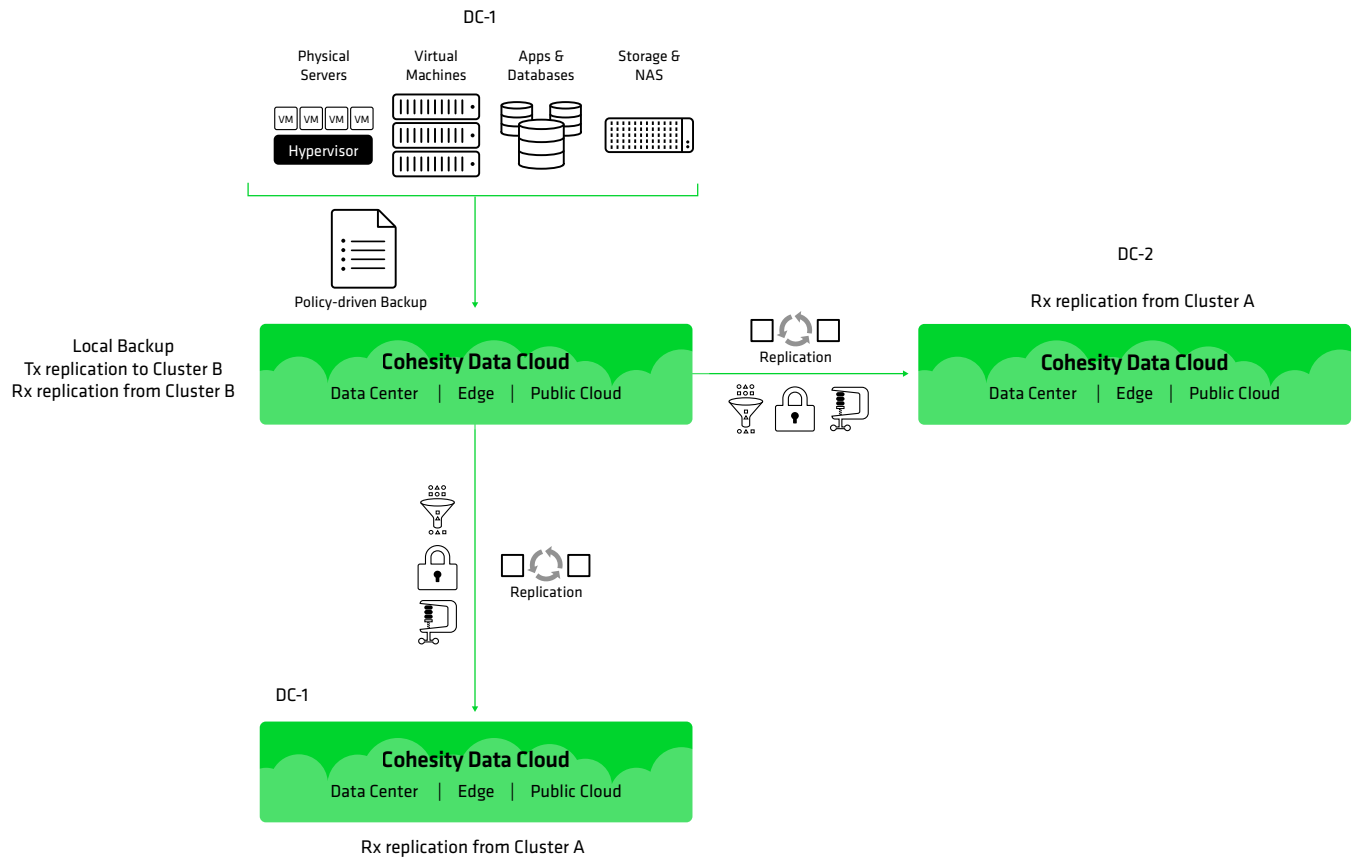
Enhanced: E1 - Backup, Replication & Archive (Hub and Spoke)



Many topologies can also be extended to **hub and spoke**, which is also known as a fan-in topology. In the hub and spoke model, individual branches have their own backups, and those backups are replicated to a single integrated Cohesity Data Cloud in a central data center. From the central data center, the replica is archived via FortKnox,

or to another private or public archive. FortKnox would be a great choice here, as it gives complete isolation in the event that both the backups and replica were compromised. With the Cohesity Data Cloud, an archive can be restored via a primary or secondary cluster.

Enhanced: E2 - Backup & Dual Replication



This topology is for those cases where the two-step restore process from an archive doesn't provide a sufficiently low RTO. All three copies (the backup and

both replicas) can be used to restore the data in a one-step process in this configuration.

Enhanced: E2 - Hub and Spoke with Active-Active Hubs



This topology combines several different models. Remote branches each do their own backups, and those backups are replicated in a central data center. This central data center has a disaster recovery data center as a backup as well.

This is all mirrored, with the primary data center on the left serving as the disaster recovery site for the data center on the right and vice versa.

Mission Critical

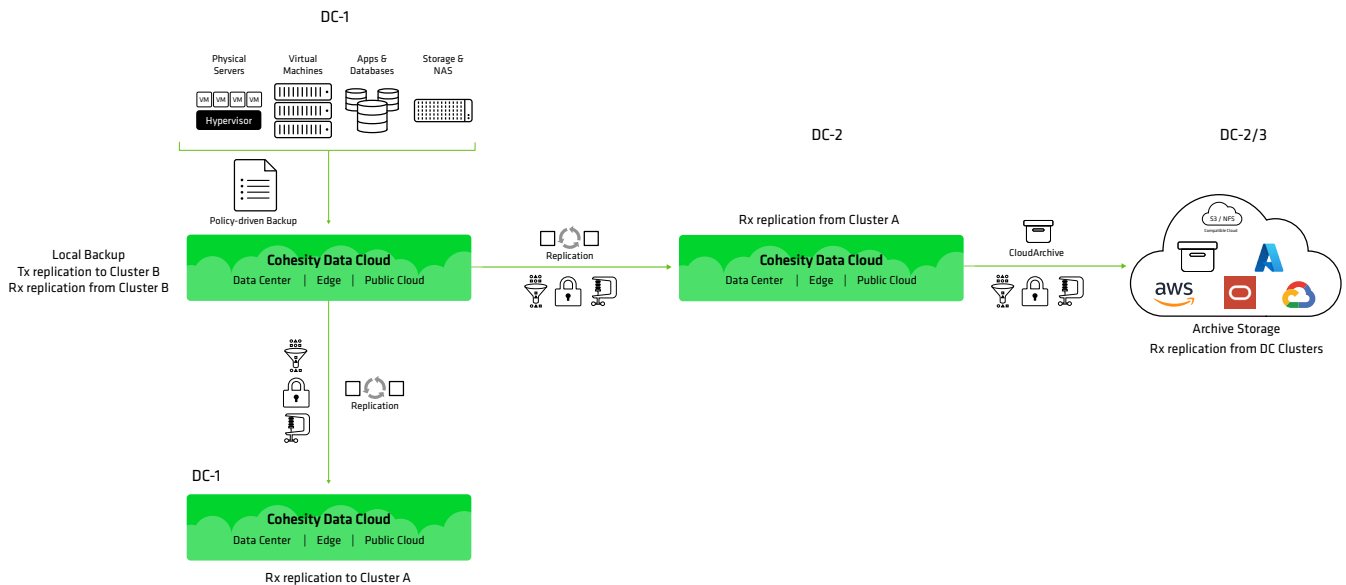
Topology	Single Data Center	Active-Active	Hub & Spoke
Backup & Dual Replica with Archive	✓		✓
Backup, Replica & Dual Archive	✓	✓	
Backup, Dual Replica & Dual Archive			✓

Mission Critical is emerging as a topology for the most valuable data at a given firm. This is the data required to run the Minimum Viable Company (MVC).

What is your Minimum Viable Company?

An MVC is the collection of applications, infrastructure, and processes that must be restored for the business to function at a minimally viable level. These systems must be brought back online first; all other systems are a secondary priority. IT leaders must employ MVC when planning their incident response and recovery strategies—and their data topology.

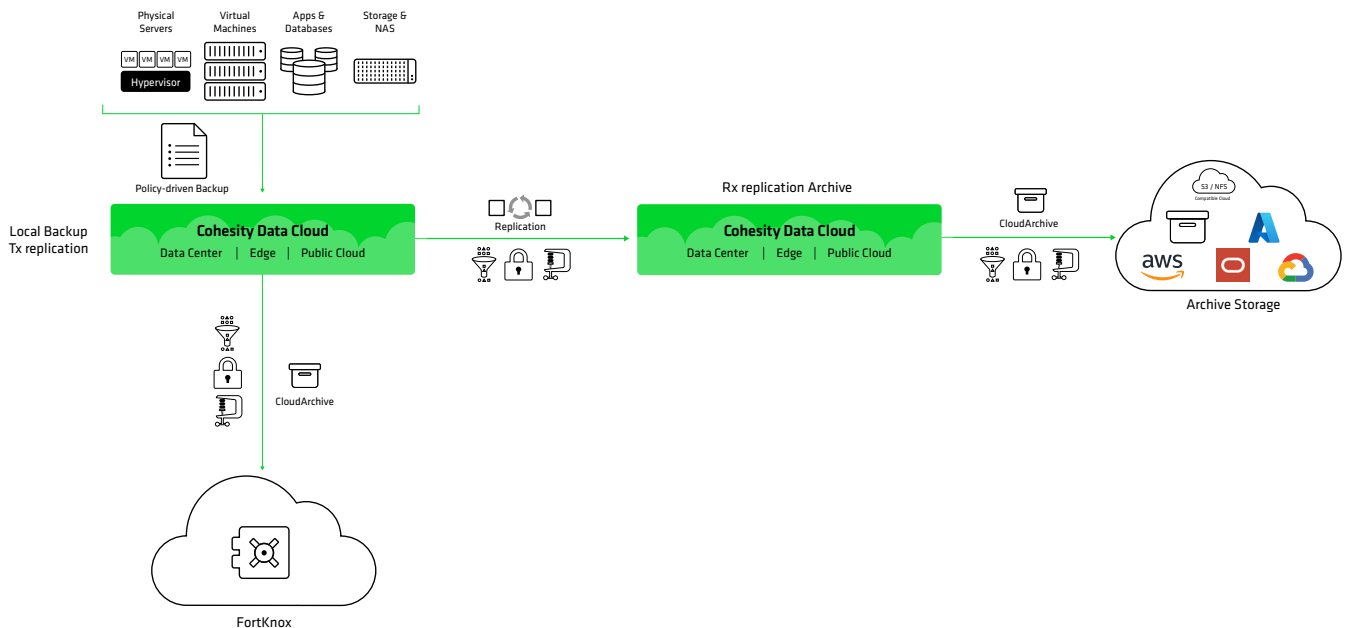
Mission Critical: M1 - Backup & Dual Replication with Archive



This fault-tolerant architecture meets stringent RTO and RPO requirements and is coupled with a long-term retention requirement. The second replica provides additional disaster recovery and ransomware protection. With the Cohesity Data

Cloud, an archive can be restored via a primary or secondary cluster. Adding an air gap to the second replica provides additional resilience.

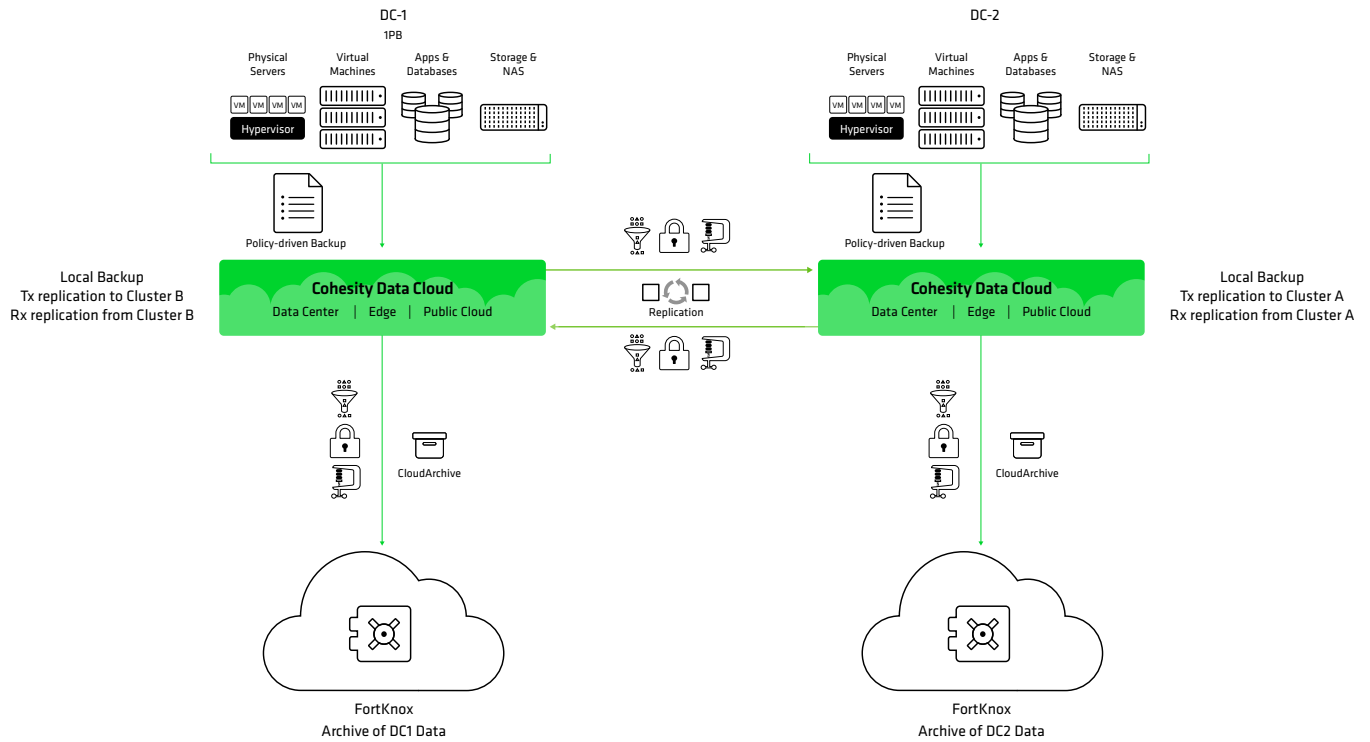
Mission Critical: M2 - Backup & Replication with Dual Archive Using FortKnox (from local backup)



This fault-tolerant architecture uses FortKnox in lieu of a second replica as FortKnox provides additional security and isolation. The first archive can be used for compliance

activities, while the FortKnox archive provides added anti-ransomware resilience. With the Cohesity Data Cloud, an archive can be restored via a primary or secondary cluster.

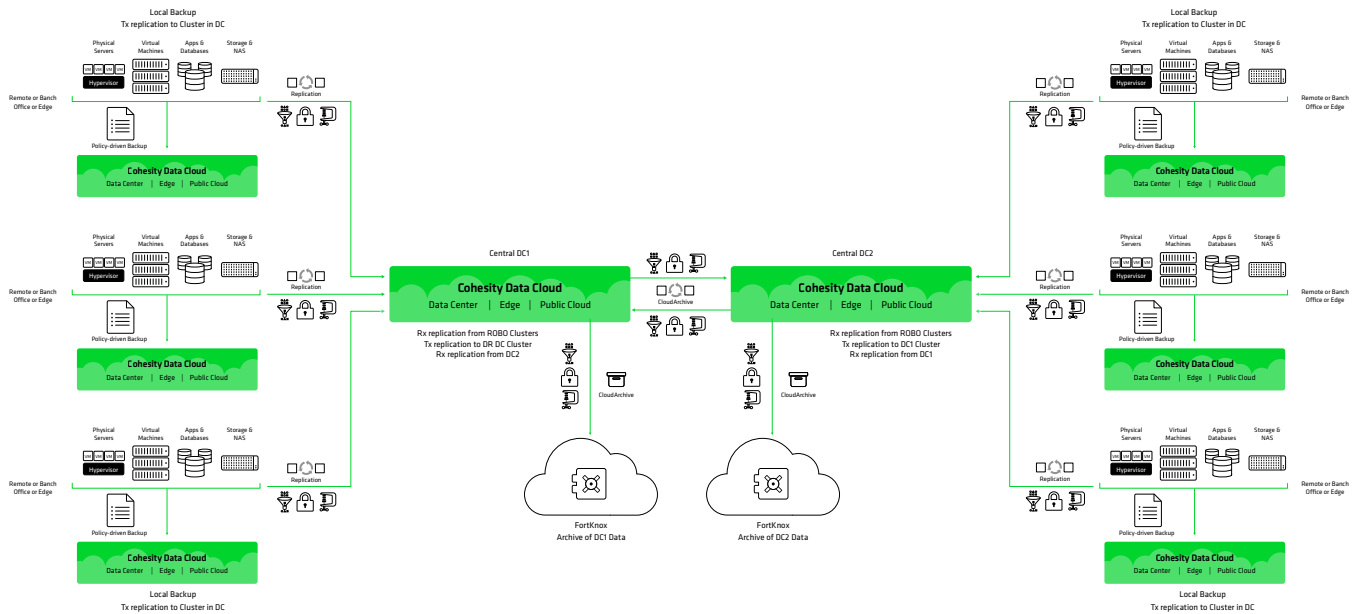
Mission Critical: M2 - Cross Replication & Archive from Local with FortKnox



This is the **Active-Active** model we saw in the **Basic** topologies, with the addition of dual FortKnox Archives. Since each cluster contains both DC1 and DC2 copies, each

FortKnox instance contains DC1 and DC2 copies as well. With the Cohesity Data Cloud, an archive can be restored via a primary or secondary cluster.

Mission Critical: M3 - Hub and Spoke with Active-Active Hubs and FortKnox Archive



This is similar to what we saw in the **Enhanced** types, but in this case each of the replicas is connected to FortKnox as a long-term archive. Since the replicas have copies of both the left and right spokes, each of the

FortKnox archives have both sets of copies as well. With the Cohesity Data Cloud, an archive can be restored via a primary or secondary cluster.

Conclusions and next steps

Many enterprise leaders seek to strengthen the protection of their critical data. As we discuss new approaches with these decision-makers, blueprints are critical. The designs discussed in this white paper will help executives understand what their peers have done in similar situations with similar data protection requirements.

“More” is not always better when it comes to copies of data. Both Cohesity and our customers understand that adding more copies of data adds operational, licensing, and often hardware costs. In some cases, we don’t advocate adding another copy, but instead encourage the use of different types of copies.

We often recommend the use of archives for company compliance activities, and to promote cyber resilience. Therefore, our customers often choose to keep the same

number of copies of data but change the type of the copies they used. They may, for example, replace an onsite, nonsecure archive with an isolated archive such as FortKnox to provide copies that can be used both for compliance and for ransomware resilience purposes.

Blueprints are powerful because they allow you to review all relevant, proven options, and make an informed decision as to which of these options you will use in your deployment .

The chart below provides a simplified, aggregated view of the benefit of each topology as it relates to failure domains, acts of God, and cyber protection.

Type	Basic		Enhanced	Mission Critical	
Copies	1	2	3	4	5
Topology	Backup Only	Backup & Repository (Replica or Archive)	Backup & Dual Repository (Replica and Archive, or Dual Replica)	Backup & Dual Replica & Archive	Backup & Dual Replica & Dual Archive
Protection from HW & SW failure domains	★	★★	★★★	★★★★	★★★★★
Protection from “Acts of God”		★	★★★	★★★★	★★★★★
Cyber protection	★	★★	★★★	★★★★	★★★★★

The journey to modern data security and management may seem daunting. We assembled this blueprint information to make it easier and help you accelerate how quickly you can achieve better business outcomes while reducing risk and cost.

From here, we recommend these next steps:

1. Determine which blueprints are most relevant to your situation.
2. Assess the ROI and TCO of a modern data platform with respect to your incumbent solution. Key points of comparison should be:
 - a. Data protection efficiency
 - b. Operational efficiency
 - c. Risk and compliance

3. Select your solution, based on product demonstrations, proven ROI and TCO calculations, and support for roadmap priorities.
4. Deploy your chosen solution following the most relevant blueprints, and proceed to execute the roadmap from the prior step.

Once your modern platform is in place, generate an initial set of KPIs with respect to cyber resilience, and regularly measure your progress against this baseline. From there, you'll know when to advance to the next phase of your journey.

About Cohesity

Cohesity is the leader in AI-powered data security. Over 13,600 enterprise customers, including over 85 of the Fortune 100 and nearly 70% of the Global 500, rely on Cohesity to strengthen their resilience while providing Gen AI insights into their vast amounts of data. Formed from the combination of Cohesity with Veritas' enterprise data protection business, the company's solutions secure and protect data on-premises, in the cloud, and at the edge. Backed by NVIDIA, IBM, HPE, Cisco, AWS, Google Cloud, and others, Cohesity is headquartered in Santa Clara, CA, with offices around the globe. To learn more, follow Cohesity on [LinkedIn](#), [X](#), and [Facebook](#).

Learn more at [Cohesity](#)

© 2025 Cohesity, Inc. All rights reserved.

Cohesity, the Cohesity logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an "AS IS" basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.

COHESITY

[cohesity.com](#)

1-855-926-4374

2625 Augustine Drive, Santa Clara, CA 95054

2000050-002 EN 4-2025