ComDirect

Disaster Recovery Planning Guide

A Step-by-Step Guide to Protecting Your Business from Data Loss and Downtime

Technology Guide

ComDirect USA
comdirectusa.com

Table of Contents

Chapter 1

Why Every Business Needs a Disaster Recovery Plan

Disasters are not theoretical. Hardware fails, ransomware encrypts, data centers flood, ISPs go dark, and employees accidentally delete production databases. The question is never whether your organization will face a disruptive event, but whether you will recover from it quickly or slowly, cheaply or expensively, completely or with permanent data loss.

A disaster recovery (DR) plan is a documented, tested set of procedures that enables your organization to restore critical IT systems and data after a disruption. Without one, recovery depends on improvisation under pressure, and improvisation under pressure is how organizations lose data, lose customers, and lose money.

The Cost of Not Planning

- The average cost of IT downtime for mid-market organizations is estimated at $5,600 per minute
- 93% of companies that experience a major data loss without a recovery plan are out of business within five years, according to widely cited industry studies
- Ransomware attacks that destroy backups leave organizations choosing between paying the ransom (with no guarantee of recovery) and rebuilding from scratch
- Regulatory penalties for data loss can be severe: HIPAA fines up to $2.1 million per violation category, PCI-DSS fines of $100,000 per month

What a DR Plan Is Not

A DR plan is not a backup schedule. Backups are one component of disaster recovery, but a plan without tested restoration procedures, defined roles, communication protocols, and recovery priorities is just a list of tapes or snapshots. Similarly, a DR plan is not a business continuity plan (BCP), though the two are closely related. BCP covers the full scope of maintaining business operations during a disruption; DR focuses specifically on restoring IT systems and data.

Key Distinction

Business Continuity answers: "How do we keep operating during a disruption?" Disaster Recovery answers: "How do we restore our IT systems after a disruption?" Both are essential. This guide focuses on the DR component.

Chapter 2

Business Impact Analysis (BIA)

The Business Impact Analysis is the foundation of every effective DR plan. It answers a deceptively simple question: what happens to the business when each system goes down, and for how long can the business tolerate it?

Conducting the BIA

A BIA is not a technical exercise conducted by IT alone. It requires input from business stakeholders who understand the operational and financial consequences of system unavailability.

- Identify all business processes that depend on IT systems
- Map each business process to the specific IT systems, applications, and data stores that support it
- For each system, determine the financial impact of downtime per hour (lost revenue, idle labor, contractual penalties, regulatory exposure)
- Determine the maximum tolerable downtime (MTD) before the impact becomes existential or unacceptable
- Identify dependencies between systems: which systems must be restored first because other systems depend on them?

BIA Output: System Criticality Tiers

Group systems into tiers based on their impact and MTD:

Tier	MTD	Examples	Recovery Priority
Tier 1: Mission-Critical	< 1 hour	ERP, Email, Phone	Immediate
Tier 2: Business-Critical	4-8 hours	CRM, File Shares	High
Tier 3: Important	24-48 hours	Dev/Test, Intranet	Medium
Tier 4: Non-Critical	72+ hours	Archive, Training	Low

This tiering drives every subsequent decision in your DR plan: how much you invest in protection, what backup frequency you need, whether you deploy replication or failover, and in what order you restore systems during a disaster.

Common BIA Mistake

IT teams often classify systems based on technical complexity rather than business impact. A simple spreadsheet that drives a $2M revenue process is more critical than a complex application that supports an internal convenience. Let the business stakeholders determine criticality; IT determines the technical recovery strategy.

Chapter 3

Defining RPO and RTO for Every System

Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are the two metrics that translate business requirements into technical specifications for your DR architecture.

Recovery Point Objective (RPO)

RPO answers: "How much data can we afford to lose?" It is measured in time. An RPO of 4 hours means you can tolerate losing up to 4 hours of data. This directly determines your backup or replication frequency:

- RPO = 0 (zero data loss): Requires synchronous replication. Every write is committed to both primary and secondary simultaneously. Expensive and latency-sensitive.
- RPO = 15 minutes: Requires near-continuous backup or asynchronous replication with frequent snapshots.
- RPO = 1 hour: Hourly incremental backups or replication intervals are sufficient.
- RPO = 24 hours: Nightly backups meet the requirement. Acceptable for non-critical systems.

Recovery Time Objective (RTO)

RTO answers: "How quickly must we restore the system?" It is measured from the moment of failure to the moment the system is operational again. RTO determines your recovery architecture:

- RTO = minutes: Requires hot standby with automatic failover (clustering, load balancing, or active-active replication).
- RTO = 1-4 hours: Requires warm standby with pre-provisioned infrastructure that can be activated quickly.
- RTO = 8-24 hours: Bare-metal restore from backup to replacement hardware or cloud instance.
- RTO = 48-72 hours: Cold recovery from offsite backup, potentially requiring hardware procurement.

RPO/RTO by Tier

Tier	Typical RPO	Typical RTO	Recovery Method
Tier 1	< 15 min	< 1 hour	Replication + Failover
Tier 2	1-4 hours	4-8 hours	Frequent Backup + Warm DR
Tier 3	24 hours	24-48 hours	Daily Backup + Cloud Restore
Tier 4	24-48 hours	72+ hours	Offsite Backup

Cost vs. Recovery Tradeoff

Tighter RPO and RTO cost more. Zero data loss with instant failover for every system is technically possible but financially impractical for most organizations. The BIA ensures you invest heavily only where the business impact

justifies it.

Cost vs. Recovery Tradeoff

Tighter RPO and RTO cost more. Zero data loss with instant failover for every system is technically possible but financially impractical for most organizations. The BIA ensures you invest heavily only where the business impact

justifies it.

Chapter 4

Building Your System and Data Inventory

You cannot recover what you have not inventoried. A complete, accurate system and data inventory is the operational backbone of your DR plan. During a disaster, there is no time to figure out what exists, where it lives, and how it is configured.

System Inventory Requirements

For every system in scope, document:

- System name, IP address(es), and physical or virtual host location
- Operating system and version, with patch level at time of last documentation update
- Applications installed, versions, license keys, and configuration files
- Dependencies: what other systems, services, or network paths does this system require?
- Data stores: databases, file shares, and local storage, with volume sizes and growth rates
- Authentication: how users and services authenticate (AD-joined, local accounts, certificates, API keys)
- Owner: the business owner (who decides priority) and technical owner (who performs recovery)

Data Classification and Location

- Identify all locations where business data resides: on-premise servers, cloud storage, SaaS applications, employee endpoints, and email
- Classify data by sensitivity (public, internal, confidential, regulated) and by recovery priority (mirrors system tiering)
- Map data flows between systems to understand which restores must happen in sequence
- Identify data that exists only in a single location with no backup (this is your highest-priority gap)

Network and Infrastructure Dependencies

- Document all network paths required for system operation: VLANs, firewall rules, VPN tunnels, DNS records, and load balancer configurations
- Identify external dependencies: ISP circuits, cloud provider regions, SaaS services, and third-party APIs
- Document authentication infrastructure dependencies: Active Directory domain controllers, certificate authorities, MFA providers

The Dependency Trap

The most common DR failure is restoring a system only to discover it cannot function because a dependency was not restored first. Map dependencies explicitly and build your recovery sequence around them. If the domain controller is not up, nothing that requires AD authentication will work.

Chapter 5

Backup Strategies and Architecture

Backups are the safety net beneath every other DR control. Even organizations with sophisticated replication and failover need backups as a last resort against data corruption, ransomware, and logical errors that replication faithfully copies to the secondary site.

The 3-2-1-1 Rule

The classic 3-2-1 backup rule has evolved to address modern threats:

- 3 copies of your data (production + 2 backups)
- 2 different storage media types (disk + cloud, or disk + tape)
- 1 copy offsite (geographically separated from production)
- 1 copy immutable or air-gapped (cannot be modified or deleted by ransomware that compromises your network)

Backup Types and When to Use Them

- Full backup: Complete copy of all data. Simplest to restore but slowest to create and most storage-intensive. Typically run weekly.
- Incremental backup: Copies only data changed since the last backup of any type. Fast and storage-efficient. Restoration requires the last full backup plus all subsequent incrementals in sequence.
- Differential backup: Copies all data changed since the last full backup. Faster to restore than incremental (needs only the last full + last differential) but uses more storage.
- Continuous Data Protection (CDP): Captures every write operation in real time. Provides granular recovery to any point in time. Resource-intensive but essential for Tier 1 systems with near-zero RPO.

Backup Targets

- Local disk or appliance: Fast backup and restore. Vulnerable to site-level disasters (fire, flood) and ransomware that spreads across the network.
- Cloud storage (S3, Azure Blob, etc.): Offsite by default. Cost-effective for long-term retention. Restore speed depends on bandwidth.
- Immutable cloud storage: Object lock prevents modification or deletion for a defined retention period. Critical defense against ransomware.
- Tape: Air-gapped by nature when removed from the library. Slow to restore but inexpensive for long-term archival retention.

Ransomware and Backups

Modern ransomware specifically targets backup infrastructure. Attackers look for backup servers, delete shadow

copies, and encrypt backup repositories before triggering the primary encryption. Your backup architecture must include at least one copy that is unreachable from the production network: immutable cloud storage, air-gapped tape, or an isolated backup network with separate credentials.

Ransomware and Backups

Modern ransomware specifically targets backup infrastructure. Attackers look for backup servers, delete shadow

copies, and encrypt backup repositories before triggering the primary encryption. Your backup architecture must include at least one copy that is unreachable from the production network: immutable cloud storage, air-gapped tape, or an isolated backup network with separate credentials.

Chapter 6

Replication and Failover Design

For Tier 1 and Tier 2 systems where RTOs measured in hours are unacceptable, replication and failover provide near-instant recovery without relying on backup restoration.

Replication Types

- Synchronous replication: Every write is committed to both primary and secondary before the application receives confirmation. Zero data loss (RPO = 0) but adds latency. Best for sites within 100km of each other.
- Asynchronous replication: Writes are committed to the primary first, then replicated to the secondary with a short delay. Minimal performance impact but accepts a small data loss window (typically seconds to minutes).
- Application-level replication: Databases like SQL Server Always On, Oracle Data Guard, and PostgreSQL streaming replication handle their own replication. Often provides more granular control than storage-level replication.

Failover Architectures

- Active-Passive: Secondary site is pre-provisioned but idle. Failover is triggered manually or by monitoring automation. RTO depends on startup and validation time (typically 15-60 minutes).
- Active-Active: Both sites handle production traffic simultaneously. If one fails, the other absorbs the full load. Near-zero RTO but requires application support for multi-site operation and conflict resolution.
- Pilot Light: Minimal infrastructure running at the DR site (domain controllers, DNS, core networking). Full environment is spun up from snapshots or images on demand. RTO of 1-4 hours at lower ongoing cost than active-passive.
- Warm Standby: Scaled-down replica of the production environment running continuously. Scales up to full capacity during failover. Balances cost and RTO.

Failover Triggers and Automation

Define clear criteria for when failover is triggered:

- Automated failover for well-defined failure modes (server crash, storage failure) with health-check-based detection
- Manual failover for ambiguous scenarios (partial outages, intermittent issues) where premature failover could cause split-brain
- Documented decision tree: who has authority to declare a disaster and initiate failover?
- Failback procedures: how do you return to the primary site once it is restored, and how do you resynchronize data?

Failback Is Harder Than Failover

Most organizations focus on failover and neglect failback. Returning to the primary site after a disaster requires resynchronizing data accumulated at the DR site, validating that the primary is stable, and coordinating another cutover. Plan and test failback with the same rigor as failover.

Failback Is Harder Than Failover

Most organizations focus on failover and neglect failback. Returning to the primary site after a disaster requires resynchronizing data accumulated at the DR site, validating that the primary is stable, and coordinating another cutover. Plan and test failback with the same rigor as failover.

Chapter 7

Cloud-Based Disaster Recovery (DRaaS)

Disaster Recovery as a Service (DRaaS) has transformed DR economics. Instead of building and maintaining a secondary data center, organizations can replicate to cloud infrastructure that is paid for on demand, scaling from a trickle of replication traffic to a full production environment only when disaster strikes.

How DRaaS Works

- Lightweight agents on production servers continuously replicate data and system state to cloud storage
- During normal operations, only replication bandwidth and storage costs are incurred (no running compute instances)
- When a disaster is declared, recovery instances are launched from the replicated data in minutes
- Once the primary site is restored, data is synchronized back and DR instances are shut down

DRaaS Advantages

- No capital expenditure for secondary hardware: pay only for storage during normal operations and compute during failover
- Geographic flexibility: replicate to any cloud region, providing distance from your primary site without building or leasing a facility
- Scalability: DR environment can match production capacity without maintaining idle servers year-round
- Simplified testing: spin up DR instances for testing without disrupting production, then shut them down

DRaaS Considerations

- Bandwidth: initial seeding of large datasets may require a physical data transfer device rather than network replication
- Network reconfiguration: during failover, DNS, VPN tunnels, and firewall rules must redirect traffic to cloud-hosted instances. Automate this.
- Licensing: verify that your software licenses permit operation in a cloud environment during DR events
- Data sovereignty: ensure the cloud region you replicate to meets any data residency or regulatory requirements
- Vendor lock-in: understand the DRaaS provider's data portability options and contract exit terms

DRaaS Is Not Set-and-Forget

DRaaS replication must be monitored continuously. If replication falls behind or fails silently, your RPO is not being

met. Treat DRaaS monitoring with the same rigor as backup monitoring: daily verification, alerting on lag or failure, and periodic test failovers.

DRaaS Is Not Set-and-Forget

DRaaS replication must be monitored continuously. If replication falls behind or fails silently, your RPO is not being

met. Treat DRaaS monitoring with the same rigor as backup monitoring: daily verification, alerting on lag or failure, and periodic test failovers.

Chapter 8

The Disaster Recovery Plan Document

The DR plan document is the operational playbook that your team follows during a disaster. It must be clear, complete, and accessible to people who are working under stress and time pressure.

Document Structure

A well-organized DR plan includes the following sections:

- 1. Plan Overview: Purpose, scope, assumptions, and plan maintenance schedule
- 2. Roles and Responsibilities: DR team members, their roles, and contact information (personal cell phones, not office numbers)
- 3. Activation Criteria: What constitutes a disaster, who has authority to declare one, and the decision process
- 4. Communication Plan: Internal notification tree, external stakeholder communication, customer notification templates
- 5. System Recovery Procedures: Step-by-step instructions for restoring each system, organized by recovery priority tier
- 6. Network Recovery: Procedures for restoring network connectivity, DNS, VPN, and firewall configurations
- 7. Validation Checklists: How to verify each system is operational after recovery (not just running, but functional)
- 8. Failback Procedures: How to return to normal operations once the primary site is restored
- 9. Appendices: System inventory, vendor contact list, license keys, network diagrams, credential vault access procedures

Writing Effective Recovery Procedures

Recovery procedures must be written for someone who has never performed them before. During a disaster, your most experienced engineer may be unavailable. Write procedures that a competent but unfamiliar technician can follow:

- Number every step sequentially. Do not use ambiguous language like "configure the server appropriately."
- Include exact commands, file paths, configuration values, and expected outputs
- Include verification steps after each major milestone ("After step 7, verify that the database responds to queries by running...")
- Include rollback instructions for steps that could fail or produce unexpected results
- Include screenshots where visual confirmation is required

Accessibility

Your DR plan must be accessible when your primary infrastructure is down. If the plan is stored only on a file server that was destroyed in the disaster, it is useless. Maintain copies in at least three locations: printed binder in a known location, cloud storage accessible from personal devices, and with your MSP or DR provider.

Accessibility

Your DR plan must be accessible when your primary infrastructure is down. If the plan is stored only on a file server that was destroyed in the disaster, it is useless. Maintain copies in at least three locations: printed binder in a known location, cloud storage accessible from personal devices, and with your MSP or DR provider.

Chapter 9

Communication and Escalation Procedures

Technical recovery is only half the battle. The other half is communication: making sure the right people know what happened, what is being done, and what they should expect. Poor communication during a disaster amplifies confusion, damages trust, and leads to bad decisions.

Internal Communication

- Notification tree: define who is notified first (DR team lead and executive sponsor), who is notified second (department heads), and who is notified third (all staff). Use automated notification tools, not manual phone trees.
- Status cadence: commit to regular status updates at defined intervals (every 30 minutes during active recovery, hourly once stable). Even "no change" updates reduce anxiety.
- Communication channels: define primary and backup channels. If email is down, use a pre-established group text, messaging app, or bridge phone line.
- Decision authority: document who can authorize spending (emergency hardware purchases, overtime), who can communicate externally, and who can approve changes to the recovery sequence.

External Communication

- Customer notification: prepare templates in advance for different scenarios (brief outage, extended outage, data breach). Have legal and PR review these templates before a disaster occurs.
- Vendor escalation: maintain a vendor contact list with escalation paths and account numbers. During a disaster is not the time to search for your ISP's emergency support number.
- Regulatory notification: know your breach notification obligations and timelines (HIPAA: 60 days, many state laws: 30-60 days, PCI: immediately). Have legal counsel on speed dial.
- Insurance notification: notify your cyber insurance carrier as soon as a qualifying event occurs. Many policies require prompt notification as a condition of coverage.

Post-Disaster Communication

After recovery, communicate what happened, what was done, and what will change:

- Internal post-incident report: timeline, root cause, actions taken, lessons learned, and remediation plan
- Customer communication: transparent summary appropriate to the audience, avoiding technical jargon
- Regulatory filings: as required by applicable laws and regulations
- Insurance claim documentation: costs incurred, timeline, and evidence of controls that were in place

Chapter 10

Testing Your DR Plan

An untested DR plan is a hypothesis, not a capability. Testing reveals gaps that are invisible on paper: procedures that assume a step that no longer works, recovery times that exceed the RTO, dependencies that were not documented, and team members who do not know their roles.

Types of DR Tests

Plan Review (Annually, minimum)

- Walk through the plan document with the DR team and verify that all information is current: contact lists, system inventory, procedures, and vendor details
- Low effort, catches documentation drift, but does not validate technical recovery

Tabletop Exercise (Semi-annually)

- Gather the DR team in a room and walk through a disaster scenario step by step. Each participant describes what they would do at each stage.
- Reveals role confusion, communication gaps, and procedural ambiguity without touching production systems

Component Test (Quarterly)

- Test individual recovery components: restore a server from backup, failover a database to its replica, activate a VPN tunnel to the DR site
- Validates that specific technical capabilities work but does not test the full recovery workflow

Full Simulation (Annually)

- Execute the complete DR plan against a realistic scenario. Restore systems in priority order, verify functionality, and measure actual RPO and RTO achieved.
- The only test type that validates end-to-end recovery capability. Requires significant planning and coordination.

Measuring Test Results

- Actual RTO vs. target RTO for each system
- Actual RPO vs. target RPO (how much data was lost in the test scenario?)
- Number of procedural failures (steps that did not work as documented)
- Number of undocumented dependencies discovered during testing
- Time spent on unplanned troubleshooting vs. following documented procedures

The Test You Skip

Organizations commonly skip DR testing because it is disruptive, time-consuming, and politically inconvenient. Every organization that has survived a real disaster without significant data loss credits one thing above all others: they

tested. Every organization that suffered catastrophic data loss during a disaster says the same thing: we thought our backups worked.

The Test You Skip

Organizations commonly skip DR testing because it is disruptive, time-consuming, and politically inconvenient. Every organization that has survived a real disaster without significant data loss credits one thing above all others: they

tested. Every organization that suffered catastrophic data loss during a disaster says the same thing: we thought our backups worked.

Chapter 11

Maintaining and Updating Your Plan

A DR plan is a living document. The environment it protects changes continuously: new systems are deployed, old ones are decommissioned, staff members change roles, vendors are replaced, and the threat landscape evolves. A plan that is not maintained degrades rapidly.

Maintenance Schedule

- Monthly: Verify backup job success rates, replication lag, and monitoring alert health. Review any infrastructure changes made during the month and update documentation.
- Quarterly: Conduct component recovery tests. Update contact lists and vendor information. Review and update the system inventory.
- Semi-annually: Conduct a tabletop exercise. Review RPO/RTO targets against current business requirements. Update procedures for any systems that changed.
- Annually: Conduct a full DR simulation. Re-execute the Business Impact Analysis. Review and renew vendor contracts and insurance policies. Present DR posture to executive leadership.

Change-Triggered Updates

Certain events should trigger an immediate DR plan review:

- New system deployment or major application upgrade
- Data center migration or cloud infrastructure change
- Merger, acquisition, or significant organizational restructuring
- Change in regulatory requirements affecting data retention or recovery
- DR team member departure or role change
- Real disaster or near-miss incident

Version Control and Distribution

- Maintain the DR plan under version control with change history
- Ensure all copies (printed, cloud, MSP) are updated when changes are made
- DR team members must acknowledge receipt and review of updated versions
- Archive previous versions for audit trail purposes

Ownership

Assign a single DR plan owner who is accountable for keeping the plan current. Without clear ownership, maintenance tasks are deferred indefinitely. The DR plan owner does not need to do all the work, but they must ensure it gets done and track completion.

Chapter 12

DR Planning Checklist

Use this checklist to track your progress in building or improving your disaster recovery program.

Foundation

- Business Impact Analysis completed with business stakeholder input
- All systems classified into recovery tiers (Tier 1-4)
- RPO and RTO defined for every system based on BIA results
- Complete system and data inventory documented
- Dependencies mapped between all critical systems

Backup and Recovery

- 3-2-1-1 backup rule implemented for all critical data
- At least one backup copy is immutable or air-gapped
- Backup monitoring in place with daily verification and failure alerting
- Test restores performed quarterly for Tier 1 and Tier 2 systems
- Bare-metal restore or full DR failover tested annually

Replication and Failover

- Replication configured for all Tier 1 systems meeting RPO targets
- Failover procedures documented with step-by-step instructions
- Failover triggers and decision authority defined
- Failback procedures documented and tested

Plan Documentation

- DR plan document complete with all required sections
- Recovery procedures written for someone unfamiliar with the environment
- Plan accessible from at least three independent locations
- Contact lists current with personal phone numbers for DR team

Communication

- Internal notification tree defined and automated
- Customer notification templates pre-approved by legal/PR
- Vendor escalation contacts documented with account numbers
- Regulatory notification requirements and timelines documented

Testing and Maintenance

- DR testing schedule established (tabletop semi-annually, full simulation annually)
- Test results documented with gap analysis and remediation tracking
- DR plan owner assigned with accountability for maintenance
- Change-triggered review process defined and followed

About ComDirect USA

ComDirect USA is a full-service technology partner specializing in enterprise infrastructure, cybersecurity, cloud architecture, and managed IT services. With offices in Irvine, California, and Meridian, Idaho, we serve organizations across healthcare, legal, financial services, manufacturing, and logistics.

Need help building your DR plan? Contact us.

Phone: 949.208.8800 | Toll Free: 800.733.7030

Email: info@comdirectusa.com | Web: comdirectusa.com