CONTINUITY HUB • QUICK CARD

Disaster Recovery Site Selection: Hot, Warm, Cold, and Cloud Architecture

QUICK REFERENCE
Current version — Nov 2, 2026. Update history: https://continuityhub.org/disaster-recovery-site-selection-hot-warm-cold-cloud-architecture/

CORE SELECTION PRINCIPLE & ISO 22301 STANDARD

Disaster Recovery Site Selection evaluates, designs, and provisions physical or virtual infrastructure hosting recovered IT systems during and after disruption. The selection decision is driven by **RTO and RPO requirements from the Business Impact Analysis (BIA)** and balances recovery speed against cost, geographic risk diversification, and operational complexity. **ISO 22301 standard:** derive targets from BIA, drill failover on schedule — an untested hot site offers less real protection than a regularly drilled warm one.

THE FOUR SITE ARCHITECTURES: TRADE-OFF SPECTRUM

Site Type	Typical RTO	Typical RPO	What's Pre-Provisioned	Relative Cost & Best For
Hot Site	**Seconds to minutes** (often near zero)	**Near zero** (continuous / synchronous)	Full duplicate environment running live, data synchronized in real time, licenses and staff in place.	**Highest** (5–10x cold; $200k–$500k/yr for mid-size enterprise Tier 1). Best for: revenue-critical, real-time payments, healthcare, emergency dispatch.
Warm Site	**Hours** (typically 4–24 hours)	**Seconds to minutes** (scheduled sync 4–24h)	Scaled-down but functional stack with hardware, network, and recent data; scaled up and synced at failover.	**Moderate** (30–60% less than hot; ~$125k/yr vs ~$300k/yr directional). Best for: Tier 2 internal systems (email, collaboration, non-real-time ERP, reporting).
Cold Site	**Days** (can extend to weeks)	**Equals backup interval** (often 24h+)	Facility only: space, power, cooling, connectivity; no live servers or current data.	**Lowest** (80–90% less than hot). Best for: non-critical or archival workloads (Tier 3/4), catastrophic fallback of last resort. Displaced by on-demand cloud.
Cloud / DRaaS	**Under 1 min to several hours**	**Seconds to minutes** (tunable by strategy)	Varies by strategy: data-only "pilot light" to full active-active serving live traffic. Standby capacity at storage rates (cents/GB).	**Pay-as-you-go** ($50–$200/VM/mo replication; $500–$2,000/mo small-biz warm DRaaS). Best for: most organizations; tunable RTO/RPO without 2nd data center.

THE FOUR AWS / CLOUD DR FRAMEWORK STRATEGIES

1. Backup and Restore — Cold Analog

Mechanism: Cloud analog of a cold site. Replicates data and backups to cloud storage. Compute resources spun up on-demand during recovery.
RTO: Hours | **RPO:** Backup interval (e.g. 24h).
Cost: Cheapest cloud tier; pays primarily for raw storage rates.

2. Pilot Light — Data Warm / Compute Off

Mechanism: Core database and data continuously replicated; compute servers, app tiers, and caches kept switched off. Spun up at failover.
RTO: Tens of minutes to a few hours | **RPO:** Seconds/minutes.
Cost: Standby storage rates + minimal replication infrastructure.

3. Warm Standby — Scaled-Down 24/7

Mechanism: Runs scaled-down version of full stack 24/7 that can handle limited traffic. Scaled up rapidly during failover.
RTO: Roughly minutes | **RPO:** Near zero / seconds.
Cost: Always-on reduced compute + storage + replication links.

4. Multi-Site Active/Active — Full Live Traffic

Mechanism: Serves live production traffic across two or more regions simultaneously using DNS or intelligent load balancers.
RTO: Under 1 minute | **RPO:** Seconds to zero.
Cost: Highest cloud cost; duplicates compute, licenses, continuous bandwidth.

CONTINUITY HUB • QUICK CARD

Disaster Recovery Site Selection: Hybrid Strategy & Geographic Rules

HYBRID RECOVERY STRATEGY: TIERING BY BUSINESS IMPACT ANALYSIS

Mature organizations avoid one-size-fits-all spending. They assign different recovery site types per system tier to align investment with business impact:

Tier	RTO Target	Architecture Strategy	Representative Workloads (from Article)
Tier 1	Near-zero RTO	Hot-site replication or cloud-native active-active	Financial trading systems, real-time payment processing, emergency dispatch, clinical healthcare, revenue-critical customer portals.
Tier 2	4–24 hours	Cloud-based warm recovery or warm standby (scheduled sync)	Email systems, internal collaboration platforms, ERP systems for non-real-time functions, reporting and analytics environments.
Tier 3	24–72 hours	Cloud-based cold recovery with daily backups	Non-critical operational systems, batch processing, developmental environments, secondary records.
Tier 4	72+ hours	Backup restoration to on-demand cloud infrastructure	Archival data, legacy reporting systems, historical compliance stores.

GEOGRAPHIC SITE SELECTION RULES

Physical & Hazard Separation

- **100–200 Miles Minimum:** Standard minimum distance for protection against regional natural disasters.
- **Hurricane Corridors:** 500+ miles separation required to clear widespread coastal storm fronts.
- **Seismic Zones:** Must be separated across different tectonic fault systems.
- **Flood Plains & Grids:** Must sit in a different hydrological flood plain and separate regional electrical power grid.

Data Sovereignty & Latency

- **Latency Boundary:** Must remain close enough for synchronous replication if near-zero RPO is required; asynchronous replication needed if distance introduces latency.
- **Regulatory Jurisdiction:** Organizations subject to GDPR, HIPAA, or national residency laws must ensure recovery site is in a compliant legal jurisdiction.
- **Cloud Region Selection:** Select regions in different availability zones/regions within the same country for compliance.

MAJOR CLOUD PROVIDERS & PROVIDER CONCENTRATION RISK

Provider / Technology	Native DR Capabilities	Key Strength / Sweet Spot
Amazon Web Services (AWS)	AWS CloudEndure, Elastic Disaster Recovery (EDR)	Continuous replication with automated failover; highest automation maturity.
Microsoft Azure	Azure Site Recovery (ASR)	Azure-to-Azure and on-premises-to-Azure; strongest hybrid on-prem integration.
Google Cloud Platform (GCP)	Asynchronous Persistent Disk (PD) replication, regional failover	High-performance storage replication; cost advantages for data-heavy workloads.

MITIGATING CLOUD PROVIDER CONCENTRATION RISK

Replicating within a single provider creates concentration risk — a global provider outage (e.g., 2024 CrowdStrike incident) disables both production and recovery. **Mitigations:** (1) Multi-cloud recovery architecture, (2) Air-gapped offline backups independent of any cloud provider, (3) Ensuring critical recovery documentation and runbooks are accessible out-of-band without cloud connectivity.