HyperSDK

Disaster Recovery & Backup Strategy

Enterprise-grade backup automation with Changed Block Tracking, scheduled exports, retention policies, and cross-provider recovery workflows.

CBT Incremental -- Cron Scheduling -- Multi-Site Recovery

Backup Architecture

Automated, policy-driven backups with incremental exports and configurable retention.

Schedule Trigger → CBT Query → Delta Export → Storage Target → Retention Cleanup

Scheduled Exports

Cron-based scheduling with timezone support. Define backup windows to minimize production impact. Multiple schedules per VM or group.

Incremental CBT

Changed Block Tracking exports only modified disk blocks since the last backup. Reduces backup time by up to 95% and storage by 90%.

Retention Policies

Configurable retention with daily, weekly, and monthly tiers. Automatic cleanup of expired backups. GFS (Grandfather-Father-Son) rotation supported.

Multi-Target Storage

Export to local storage, S3-compatible object stores, NFS shares, or directly to another cloud provider for cross-region disaster recovery.

95%
Faster Incremental Backups

90%
Storage Savings with CBT

10+
Cloud Providers Supported

Changed Block Tracking

Dramatically reduce backup windows and storage costs with incremental delta exports.

How CBT Works

The hypervisor tracks which disk blocks have changed since the last snapshot. HyperSDK queries these changed blocks and exports only the deltas, reconstructing full images at restore time.

Base + Delta Model

An initial full backup creates the base image. Subsequent backups are lightweight deltas referencing the base. Periodic consolidation merges deltas into a new base.

Full vs. Incremental Backup Comparison

VM Size	Full Backup Time	CBT Incremental Time	Storage (Full)	Storage (Delta)
50 GB	25 minutes	1.5 minutes	50 GB	3-5 GB
200 GB	100 minutes	5 minutes	200 GB	10-20 GB
1 TB	8+ hours	25 minutes	1 TB	50-100 GB

> *With CBT enabled, our nightly backup window dropped from 6 hours to under 20 minutes for 500+ VMs. Storage costs fell by 85%.*

Backup Scheduling

Flexible cron-based scheduling with time windows, timezone support, and priority queuing.

Cron Expressions

Standard 5-field cron syntax for precise scheduling. Supports complex patterns like "every 4 hours on weekdays" or "first Sunday of each month at 2 AM."

Timezone Support

Schedules respect IANA timezone identifiers. Run backups at 2 AM local time regardless of server location. DST transitions handled automatically.

Time Windows

Define allowed execution windows to prevent backups during peak hours. Jobs queued outside windows are held until the next valid window opens.

Priority Queue

Assign priority levels to backup jobs. Critical production VMs are backed up first. Priority inheritance prevents starvation of lower-priority tasks.

Schedule Configuration Examples

```
schedules:
  - name: "nightly-production"
    cron: "0 2 * * *"           # Every day at 2:00 AM
    timezone: "America/New_York"
    priority: 10
    targets: ["prod-*"]
    type: "incremental"

  - name: "weekly-full"
```

```yaml
    cron: "0 1 * * 0"           # Every Sunday at 1:00 AM
    timezone: "America/New_York"
    priority: 5
    targets: ["prod-*"]
    type: "full"

  - name: "dev-daily"
    cron: "0 22 * * 1-5"        # Weekdays at 10:00 PM
    timezone: "UTC"
    priority: 1
    targets: ["dev-*"]
    window: "22:00-06:00"
```

Recovery Workflow

End-to-end restore process: from backup selection through format conversion to target deployment.

1. **Select Backup Point**

   Browse available backups by date, VM name, or tag. View backup metadata including size, type (full/incremental), and verification status.

2. **Reconstruct Image**

   For incremental backups, HyperSDK automatically merges the base image with all required deltas to reconstruct the full VM image at the selected point in time.

3. **Format Conversion**

   Convert the backup to the target format: VMDK, QCOW2, VHD, OVA, or raw. Automatic format detection based on the destination provider.

4. **Deploy to Target**

   Import the restored image to any supported provider. Configure networking, storage tier, and compute resources during deployment.

5. **Verify & Validate**

   Automated post-restore verification: boot test, network connectivity check, application health validation. Generate recovery report.

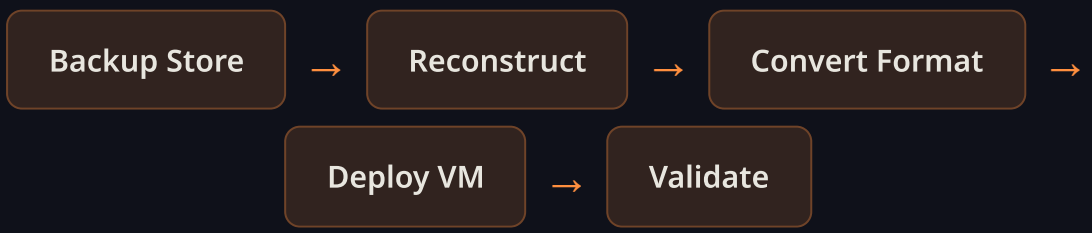

DR Best Practices

Guidelines for RPO/RTO targets, recovery testing, and multi-site backup strategies.

RPO/RTO Targets

- **Tier 1 (Critical):** RPO 1 hour, RTO 30 minutes
- **Tier 2 (Important):** RPO 4 hours, RTO 2 hours
- **Tier 3 (Standard):** RPO 24 hours, RTO 8 hours
- Align backup frequency with RPO requirements
- Use CBT incremental for tight RPO windows

Testing Recovery

- Schedule quarterly DR drills with full restore tests
- Automate restore validation with health checks
- Document recovery runbooks for each tier
- Measure actual RTO against targets
- Test cross-provider recovery paths

Multi-Site Backup

- Replicate backups to a secondary region or provider
- Use different providers for backup diversity
- Air-gapped copies for ransomware protection
- Geographic distribution for natural disaster resilience
- Verify cross-site restore latency regularly

Operational Guidelines

- Monitor backup job success rates (target: 99.5%+)
- Alert on consecutive backup failures
- Review retention policies quarterly
- Encrypt backups at rest and in transit
- Maintain backup inventory with automated auditing

Enterprise-Ready Disaster Recovery

HyperSDK combines CBT incremental backups, flexible scheduling, and cross-provider recovery to deliver a complete DR solution. Achieve your RPO/RTO targets with automated, policy-driven backup workflows.