

Hyper-V Replica

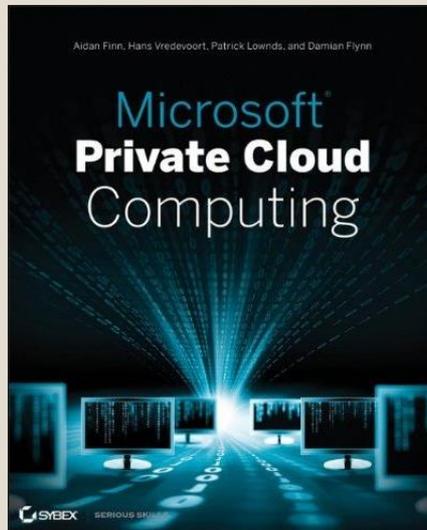
Aidan Finn

About Aidan Finn

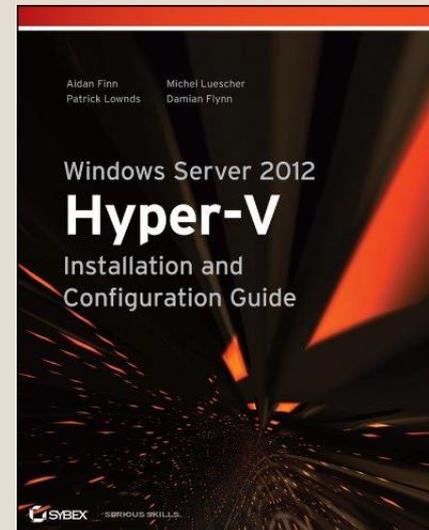


- Technical Sales Lead at MicroWarehouse (Dublin)
- Working in IT since 1996
- MVP (Virtual Machine)
- Experienced with Windows Server/Desktop, System Center, virtualisation, and IT infrastructure
- @joe_elway
- <http://www.aidanfinn.com>
- <http://www.petri.co.il/author/aidan-finn>
- Published author/contributor of several books

Books



System Center
2012 VMM



Windows Server
2012 Hyper-V

Agenda

- DR basics
- Introducing Hyper-V Replica
- Adding functionality
- Automation
- Ongoing management
- Replica IP addressing
- And the answers to the inevitable questions

DR Basics

Some Essential Terminology

- Business Continuity Planning (BCP)
 - A plan for disaster recovery
 - Formed at the business level
 - Includes IT's role in the plan, communications, services access, etc.
 - Should also include long term plan for business recovery
- Recovery Time Objective (RTO)
 - How long will it take to get services back online?
- Recovery Point Objective (RPO)
 - How much data will be lost by bringing services online in the DR site?

What is Backup?

- A process of copying data to offline storage
- Think of it as an archive
- Best practice:
 - Keep short-term data local
 - Quick restoration of most
 - Store long-term data off-site
 - Recover from total data loss
- Sometimes used to restore from disasters
 - S-l-o-w: Long RTO
 - Infrequent restore points: Long RPO
 - Unreliable: have you tested backups/tapes?
 - Where are the couriers during disaster?

What is High Availability (HA)?

- An automated process
- Deals with infrastructure failure
- Example:
 1. Host fails
 2. Virtual machine fails over to a healthy host
 3. Virtual machine starts running
 - No human intervention
- Very short RPO and RTO
- Intended for single logical site (such as campus):
 - Very complex and expensive to stretch to multi-site
 - WS2012 R2 HA is not intended for DR

What is Disaster Recovery (DR) Replicaion?

- A process for replicating services & data to another location:
 - Other side of campus
 - Another office or data center
 - A service provider
- Frequent replication
- Designed to get services back online very quickly
- Automatic invocation?
 - Yes: Very short RTO, but risk of false positive that can cause more genuine damage than less possible disaster
 - No: Longer RTO but allows business leaders to judge the scenario. Best practice almost all of the time.

Synchronous Replication

- Data is replicated live from production site
- Changes are not acknowledged until committed in both sites
- Requires
 - Expensive storage
 - Sometimes requiring additional licensing
 - Extremely low latency on replication link (usually < 2 MS)
 - Secondary site must be relatively close to primary site

Asynchronous Replication

- Data is replicated frequently instead of constantly
- Changes are acknowledged after commitment in primary site
 - Replicated some time afterwards
- Benefits
 - Cheaper storage
 - Low latency links
 - Allows DR site to be further away

Examples of Disasters

- Hurricane Sandy – East coast USA, 2012
- Typhoon **** - Phillipines, 2013
- Floods – Australia, ****
- Hurricane Katrina - USA Gulf Coast, *****
- Earthquake – San Francisco, 19**
- Fire
- Chemical leak
- 2 days running of sunshine - Ireland, any time of year

Every Business Should Have a BCP?

- Every business is at risk
- Why doesn't every business have DR systems?
 - Too expensive
 - Dark fiber networking
 - Replicating SANs + licensing
 - Expensive software
 - Too complex
 - Servers, clustering, storage & networking
 - One of the most expensive infrastructure solutions
- Even large corporations have struggled
 - How can Small/Medium Enterprises (SMEs) have a chance?
 - SMEs have the same need to survive a disaster

Introducing Hyper-V Replica

Every Business Should Have a BCP?

- Every business is at risk
- Why doesn't every business have DR systems?
 - Too expensive
 - Dark fiber networking
 - Replicating SANs + licensing
 - Expensive software
 - Too complex
 - Servers, clustering, storage & networking
 - One of the most expensive infrastructure solutions
- Even large corporations have struggled
 - How can Small/Medium Enterprises (SMEs) have a chance?
 - SMEs have the same need to survive a disaster

Hyper-V Replica

- A free DR replication solution built into Windows Server since 2012
 - Improved in 2012 R2
- Replicate VMs from one site to another
- Asynchronous replication
- No extra licensing required
- One of the most popular features in WS2012 Hyper-V
- Granular control - done on a:
 - Per VM and ...
 - ... per virtual hard disk
 - ... basis
- Designed for the SME but the Fortune 1000's want it too

Enabling Inbound Replication

- Configured in Host Settings on replica host
- 2 types of authentication:
 - HTTP/Kerberos: trusted networks with mutual AD authentication
 - HTTPS/SSL: untrusted networks with no common AD
- Configure source hosts:
 - Allow all authenticated hosts
 - A single storage location for all replica VMs
 - Fine a very small configuration, such as a lab
 - Per source authorization
 - Configure replica VM location on a per-source basis
- Enable HTTP or HTTPS HVR firewall rule

Replication Process

- You select which VMs & virtual hard disks to replicate
 - Tip: put the guest OS paging file in non-replicated disk
- Hyper-V maintains a Hyper-V Replica Log (HRL) file with/for each replicated virtual hard disk
 - Tracks changes to the virtual hard disk
- HRL file is swapped out at the end of each interval
 - Old HRL is replicated (final changes only and compressed by default)
 - New HRL continues to track changes
- Interval options:
 - 30 seconds
 - 5 minutes (only option in WS2012)
 - 15 minutes

Historical Copies

- Not limited to just the last replication
- Hyper-V Replica offers another option
- Can retain historical copies of a VM
- Choose to maintain hourly copies of a VM in the secondary site
- You can choose to use one of the historical copies when you failover the VM in the secondary site
 - Maybe data corruption is a part of the disaster
 - Travel back in time

Test Failover

- Like backups, a BCP that isn't tested is worthless
- Educate and test
- Hyper-V offers a test failover process:
 - You can configure a test virtual switch
 - Isolate the test from production systems
 - Creates a linked copy of replica VM
 - Won't interfere with ongoing replications
 - Rapidly test DR, but can also use it to do other kinds of testing

Unplanned Failover

- Used when an unexpected disaster happens:
 - Fire
 - Earthquake
- You have lost the hosts in the production site
- Business decides to invoke the BCP
- Start up VMs in secondary site
 - Some data loss – uses data from last replication interval
 - Maximum RPO = replication frequency interval
 - Time to get back online = time to start up VMs & services
 - Tiny RTO

Planned Failover

- Used to avoid an expected disaster:
 - A hurricane
 - Flooding
- You act before you lose the primary site
 - Shut down the VMs in the production site
 - Failover VMs to secondary site
 - Last bit of data is replicated
 - Replication is automatically reversed!
- There is
 - No data loss – zero RPO
 - Very quick failover – short RTO

Adding Functionality

Replicating To/From Clusters

- Imagine the possible complexity:
 - Configure up to 64 destination hosts
 - Configure inbound replication for up to 64 hosts
- Instead, we enable the Hyper-V Replica Broker once on each cluster
- Configures a computer account & FQDN for each source and/or destination cluster
- Gives you a single configuration point & identity for each source and/or destination cluster

Using HTTPS

- No need for common AD
 - Authenticate using X.509 certificates
- Can use untrusted networks
 - Traffic encrypted using HTTPS
- Opens up possibilities:
 - Service provider offering hosted DR-as-a-Service
- Certificate requirements:
 - A common PKI – trusted root authority and certificate revocation list (CRL)
 - Install computer certificate for the public name of the host or Hyper-V Replica Broker
 - Select the certificate when configuring replication

Extended Replication

- In WS2012 R2 Hyper-V you are not limited to SiteA – SiteB replication
- You can extend replication from SiteB to SiteC
- Scenarios:
 - Production in one side of campus, local DR in other side, and remote DR in another city/state
 - Production in branch office, DR in head office, and head office has a DR site too
 - Production on premise in SME, DR in service provider, and service provider has a DR site
- There is no A-B and A-C replication
 - Extended replication = A – B – C replication
 - Replication is extended from SiteB to SiteC

Automation

Manual Invocation

- How do you start the VMs after decision is made to invoke the BCP?
- Do you need to start VMs in specific order?
- Remember: the day a disaster happens will be the worst and most chaotic day of your career
 - People might be missing or worse
 - Staff probably will focus attention on safety of family first
 - Executives will be screaming at you as if it's all your fault
 - Murphy's Law
- A small business might be able to do this manually, but medium or larger will require process automation

PowerShell

- You can put together basic or complex automation scripts using PowerShell
- I have (in Windows Server 2012 Hyper-V book) a PowerShell script:
 - Document failover configuration in a CSV file
 - Script reads in the CSV file
 - Starts up VMs in correct order
 - Uses a pre-configured timeout to know when to stop waiting for a VM to start

System Center (Option 1)

- Model the failover using System Center Orchestrator
 - Will require some clever scripting to avoid non-recommended complex runbooks
- Could start a master runbook manually
- A service-oriented company could publish the process via Service Manager

System Center (Option 2)

- Run a production cloud and manage it using System Center Virtual Machine Manager (SCVMM) 2012 R2
- Run a secondary cloud and manage it using another SCVMM 2012 R2 deployment
- Sign up (requires per-replicated-VM payment) Hyper-V Recovery Manager (an Azure service)
 - Downloads agents to both SCVMM installations
 - Manages direct site-site replication
 - Allows you to create complex runbooks in HRM for automating failover

Ongoing Management

Intended Market

- Hyper-V Replica was designed for the SME
 - Simple to setup and own
 - Designed for commercial broadband (including outages)
 - Simple to monitor
- Replication information is displayed in the console
- You can dig deep using PowerShell

Replica IP Addressing

VMs Useless in DR Site If Not On Network

- IP address virtualization
 - Really only an option for massive corporations and telcos
- Stretched VLANs
 - Simple for server admin
 - Painful for network admins
 - Not an option for service providers – every customer using 192.168.1.0/24
- Hyper Replica IP address injection
 - OK for small implementations
 - Might cause issues for IP-centric (not DNS) services
- Hyper-V Network Virtualization
 - VM IP address (Customer Address) does not change
 - Great for multi-tenant service providers

I Delayed This Long
Enough

Bandwidth

- There is no “right” amount of bandwidth
- You need whatever is enough to replicate your change generated during the replication interval
 - If 50 GB is created in 5 minutes, then you need to replicate 50 GB within the next 5 minutes
- Use the Capacity Planner For Hyper-V Replica to monitor your environment and determine needs
- Ideally you have dedicated line for replication traffic
- Alternatively, use QoS to guarantee other systems their required bandwidth to/from Internet

Gotchas

- Hyper-V Replica will not replicate:
 - VM configuration changes (do this manually)
 - Virtual hard disk size changes (do this manually)
 - Passthrough disks (which are the work of the devil)
 - iSCSI or Fibre Channel LUNs that a VM connects to directly