

# Cloud Seeding

## Some Practical Methods for Moving Data between the Enterprise and the Cloud

Chris Dunder and Gradimir Starovic

Symantec, SAMG

- ❑ Why we are looking at clouds
- ❑ Amazon AWS storage services
- ❑ Symantec products we used
- ❑ Cloud storage as a tier
- ❑ Replication and failover between on-premise and cloud
- ❑ Things we learned while using cloud storage
- ❑ Things to try next

# Why cloud storage

- ❑ There is obviously a wave of interest in clouds among users and product/service providers
- ❑ Some of the potential benefits
  - ❑ on demand capacity, massive scale
  - ❑ pay as you go
  - ❑ reduced need for IT staff and premises

# Why cloud storage

- And some concerns, such as
  - security
  - availability
  - performance
  - integration with the legacy applications and data

- ❑ Cloud-based computing and storage will co-exist with other, traditional ways of provisioning
  
- ❑ Options for applications and data, to be
  - ❑ both on-premise or both in cloud
  
  - ❑ one on-premise and the other in cloud
  
  - ❑ in two different clouds

- ❑ About ongoing prototyping work
- ❑ Exploring feasibility and finding the limits
  - ❑ applications on-premise and using cloud storage
  - ❑ failover between on-premise and cloud
  - ❑ using existing products (Veritas Storage Foundation and Cluster Server) to migrate or replicate data, monitor/failover applications
  - ❑ using AWS storage

# Using Amazon AWS storage

- ❑ For pragmatic reasons, seems to be more mature and widely used at the moment than other public clouds
- ❑ Looking at both Elastic Block Storage (EBS) and Simple Storage Service (S3)
- ❑ Would like to do similar experiments with other clouds

- ❑ Storage Foundation (**SF**), enterprise-class management for online data
  - ❑ Veritas Volume Manager (**VxVM**), host-based disk virtualization - Windows & UNIX/Linux
  - ❑ Veritas Volume Replicator (**VVR**), network data replication for SF - Windows & UNIX/Linux
  - ❑ Veritas File System (**VxFS**), advanced-feature file system - UNIX/Linux
  - ❑ Veritas Dynamic Storage Tiering (**DST**) – Unix/Linux



- ❑ Veritas Cluster Server (**VCS**)
  - ❑ Shared storage clusters supporting application failover within a data center and campus-wide
  - ❑ Global Cluster Option (**GCO**), extending failover across distances spanned by WAN
  - ❑ Windows & UNIX/Linux

# Storage Foundation (SF) in the cloud

# Storage Foundation in the cloud

- ❑ Single machine in cloud
- ❑ Includes VxVM and VxFS; Basic functionality ok
- ❑ Xen kernel
- ❑ EBS devices respond to only some SCSI commands
- ❑ Single path to all EBS devices
  - ❑ No Dynamic Multipathing (DMP)
  - ❑ DMP is irrelevant to performance w/ EBS (160-400Mbps w/ EBS, and higher on the local scratch drive)

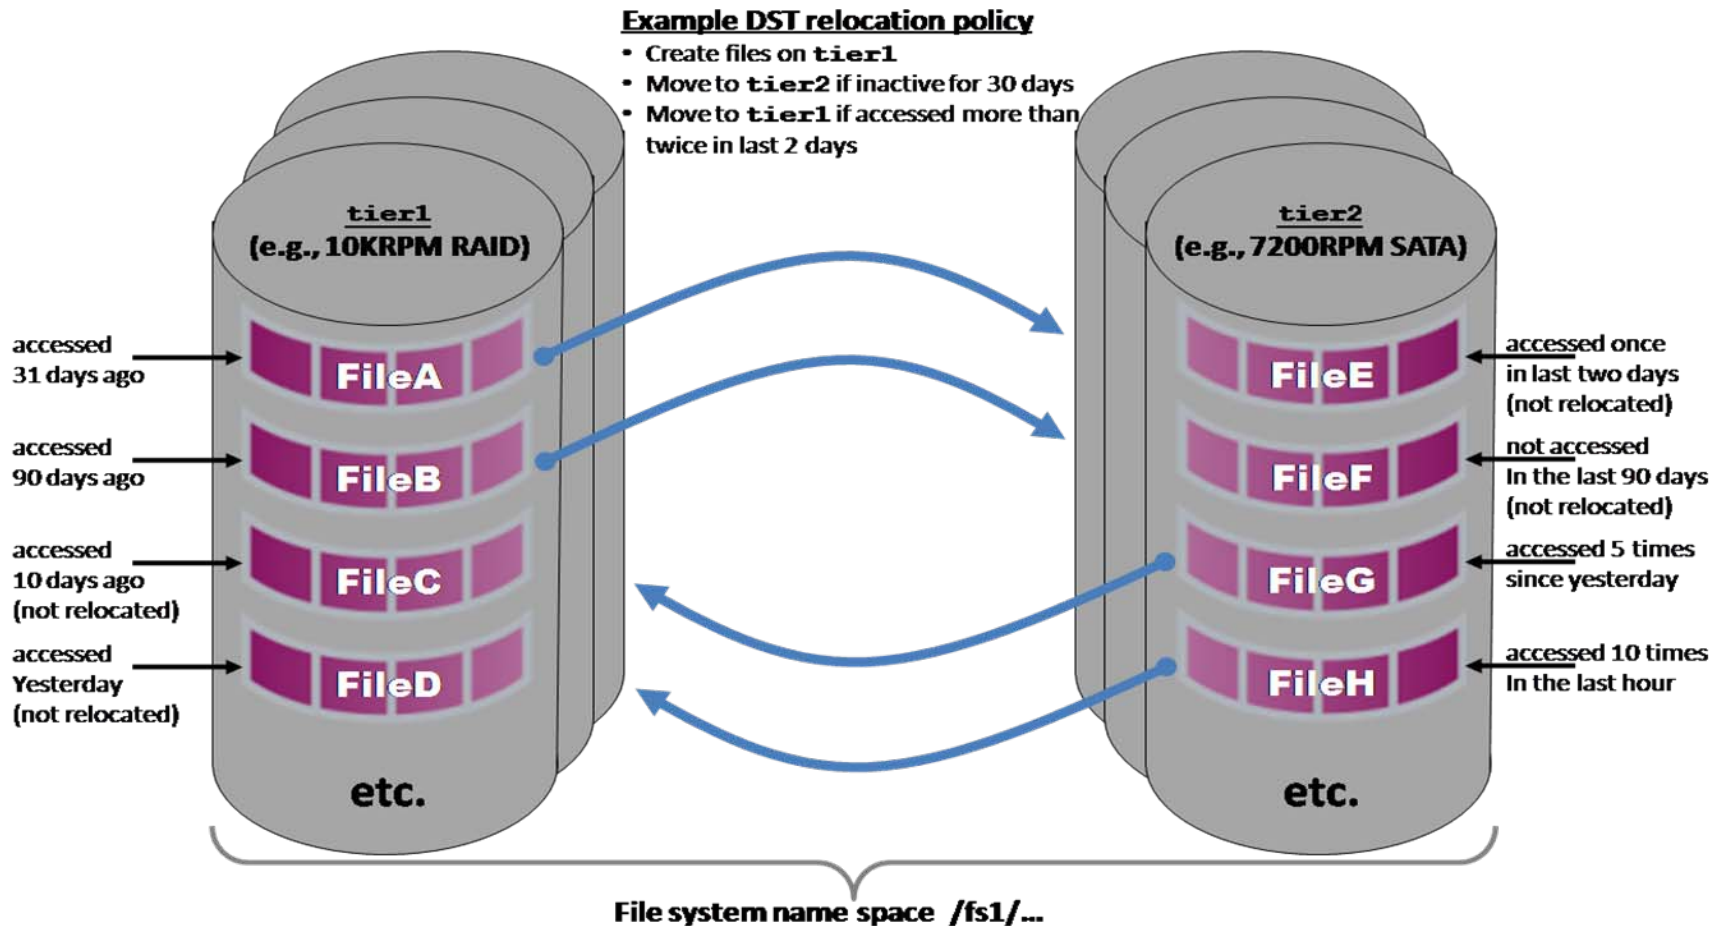
# Cloud storage as a tier

- ❑ Storage with availability, performance and cost characteristics different from other available storage (other tiers), and the difference is such that it justifies movement of data between them
- ❑ In cloud
  - ❑ storage classes (EBS, S3)
  - ❑ cost is space and usage based
  - ❑ performance varies between the storage classes, between regions and availability zones

# Veritas SF dynamic storage tiering

- ❑ A mechanism for placing data on the suitable tier
- ❑ Location transparency
- ❑ Policy controlled initial file placement (on creation), migration and removal
- ❑ Based on: file name and size, directory, time of last access, access frequency, ownership
- ❑ Multi-volume file systems, for performance, availability and migration

# Example: dynamic storage tiering




# EBS as a tier - process

- ❑ In cloud
  - ❑ using a public Linux AMI provided by Amazon
  - ❑ export iSCSI target
  - ❑ storage provided by EBS volumes
  
- ❑ On-premise server
  - ❑ create a multi-volume filesystem, some volumes are iSCSI targets in EC2
  - ❑ create DST policy
  - ❑ attach the policy to the file system



# Cloud storage as a tier & issues



- ❑ Requires 24x7 instance
- ❑ Access through firewall

# Cloud storage as a tier – status

- ❑ Currently tiers inside and outside of cloud, utilizing VPN to avoid firewall issues
- ❑ Placement rules: time of last access, access frequency, name pattern
- ❑ Investigated use of S3 w/ FUSE and EBS w/ iSCSI

# S3 as a tier – details

- ❑ Provides file-level access only
- ❑ PUT: 32-64Mbps; GET: 20-96Mbps  
(cloudstatus.org averages for week of Sep-6-09)
- ❑ Snapshots: Not provided
- ❑ Replication: Automatic (3 copies in 3 locations)

- ❑ FUSE-based products provide only a file system view
- ❑ Elastic Drive provides a block level device, but it still doesn't respond to any SCSI commands
- ❑ VxVM won't work without at least some working SCSI commands
- ❑ Result: Poor performance when treated as a pseudo-block-level device
  - ❑ S3 is optimized for file-level access

# EBS as a tier – details

- ❑ Provides block-level access to devices
- ❑ EBS devices: 160-400Mbps
- ❑ AWS internal network: 250-1000Mbps  
(S3 to EC2 large file transfers)
- ❑ Snapshots: low performance if compared against VxVM or VxFS snapshots, but expected since it is using S3
- ❑ Replication: Automatic, but transparent

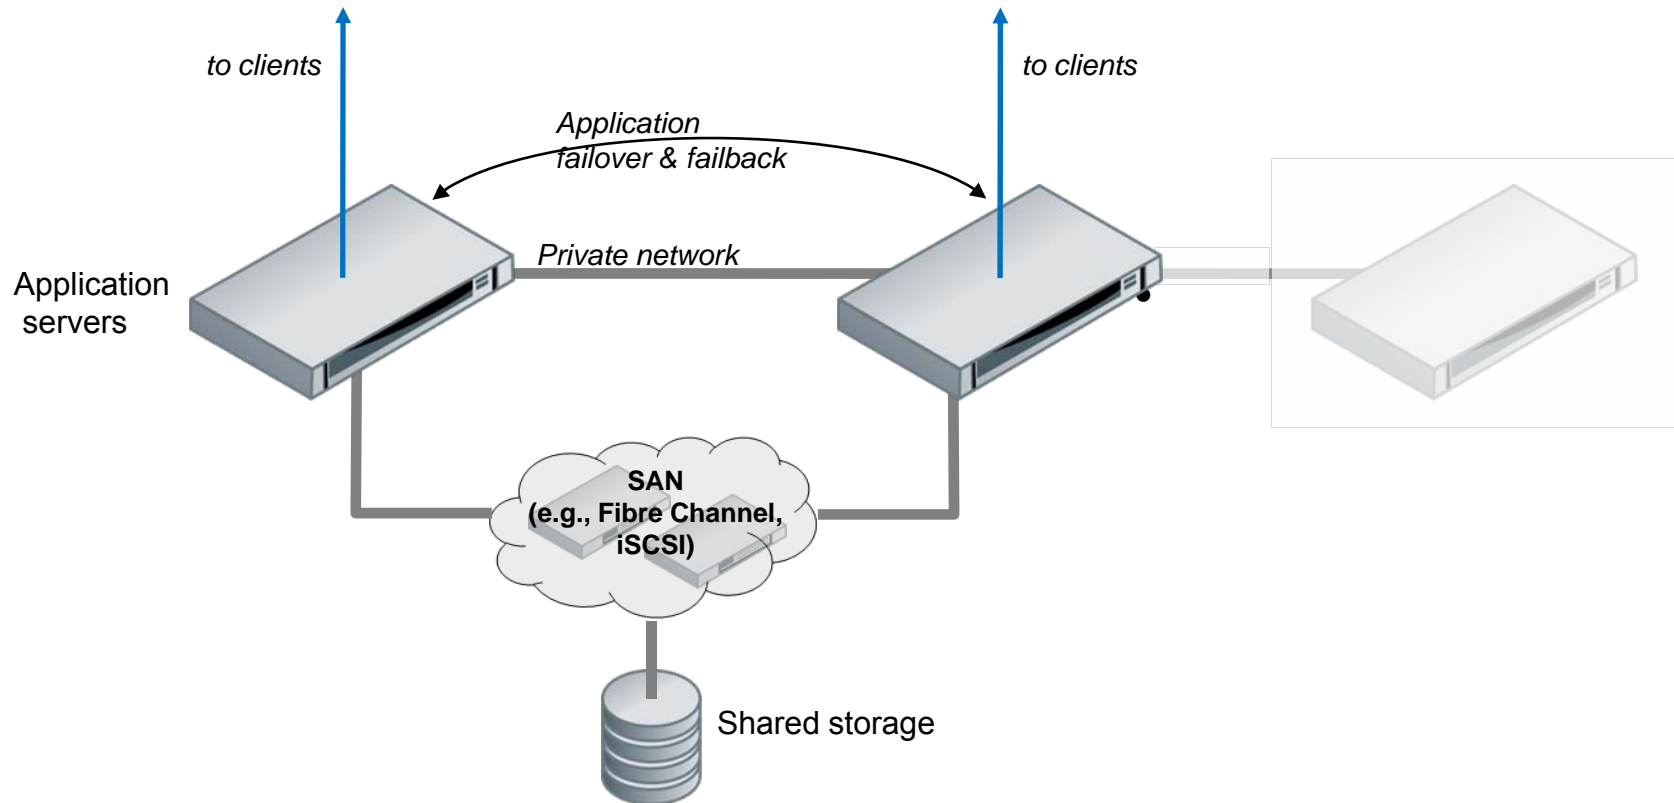
# EBS as a tier – details

- ❑ Costly
  - ❑ Requires EC2 instance (24x7)
  - ❑ \$75-\$600/mo not including data transfer costs
- ❑ EBS volumes act as pseudo-SCSI devices which limits some functionality (e.g., no DMP)
- ❑ Limit on the size and number of EBS volumes (can be increased by request)

# Replication and failover using cloud

# Background: Veritas Cluster Server

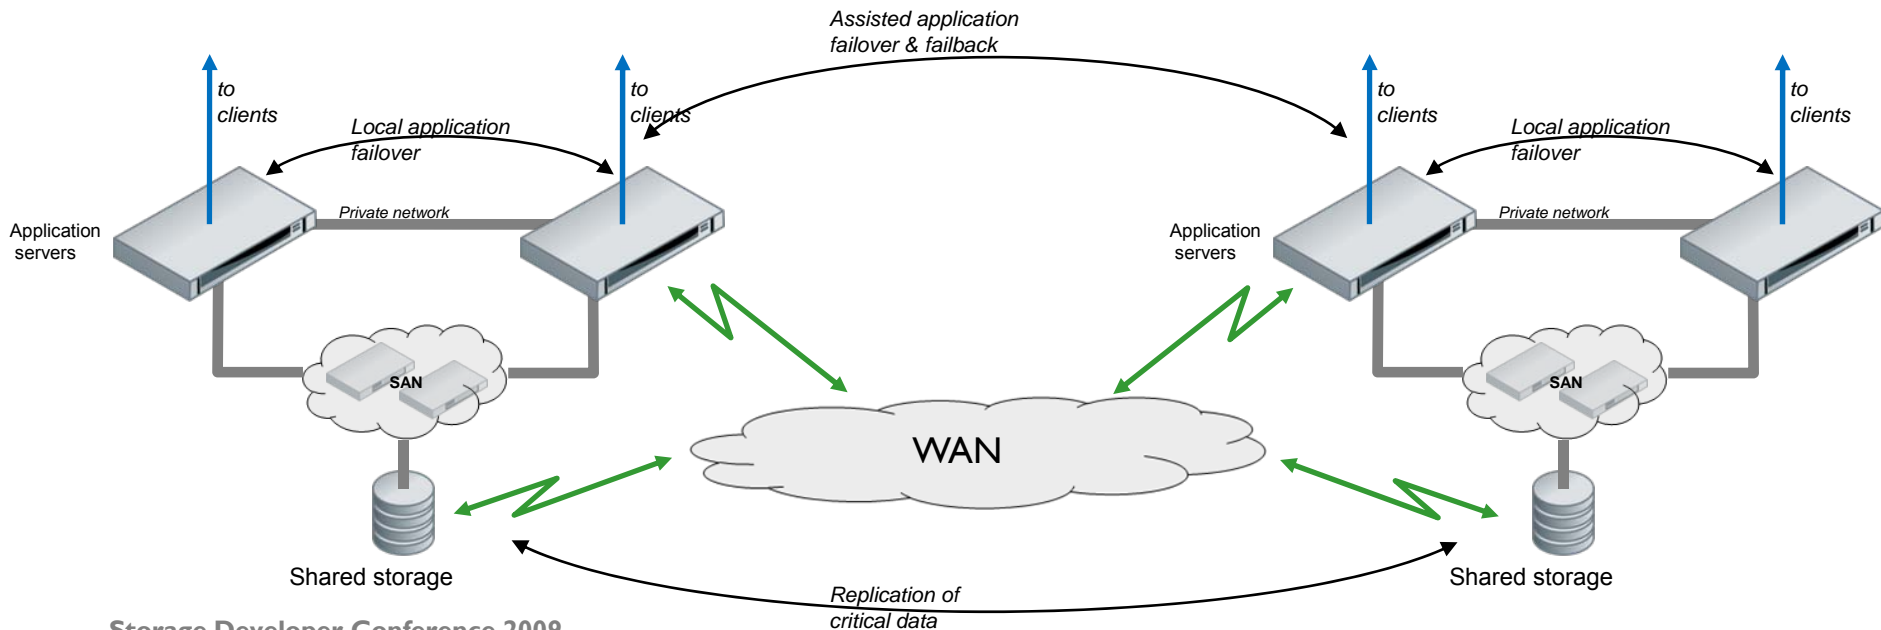
- 2-32 application servers
- Shared storage devices





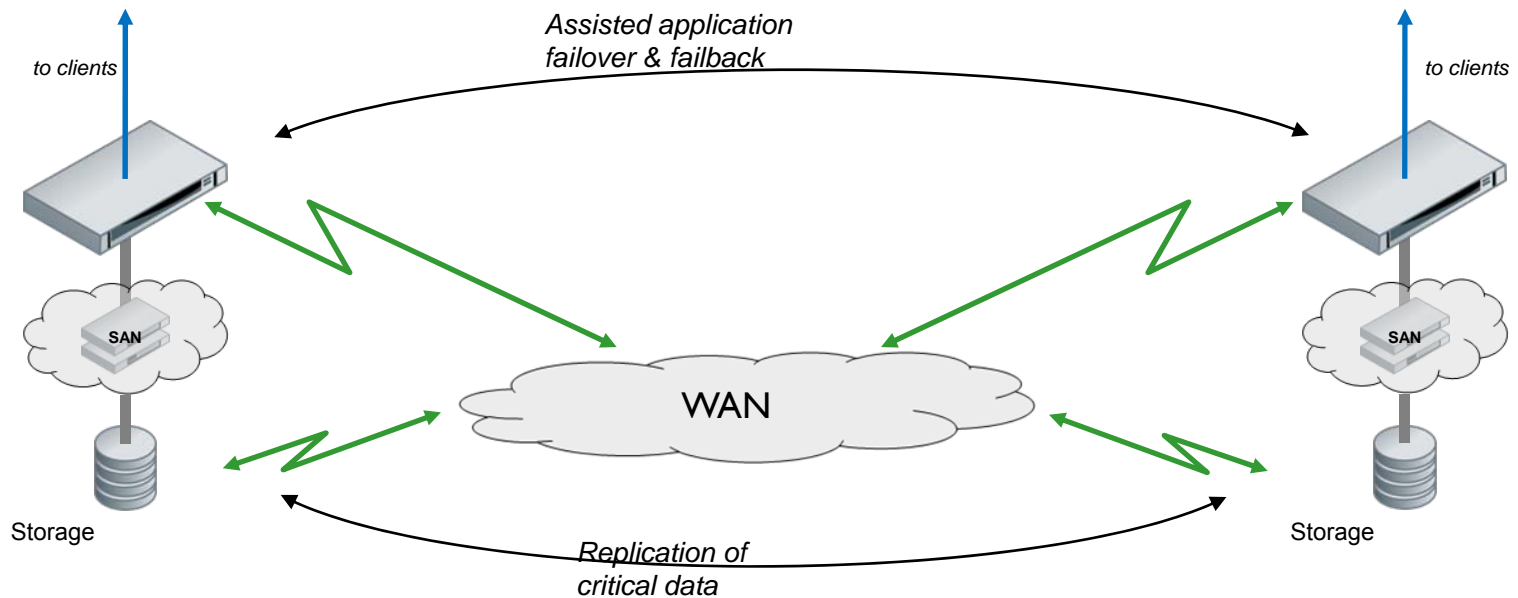
# VCS global cluster option (GCO)

- ❑ Two geographically separated VCS clusters
- ❑ Unreliable “heartbeat” network
- ❑ Long-distance replication of critical data
- ❑ Can be asynchronous

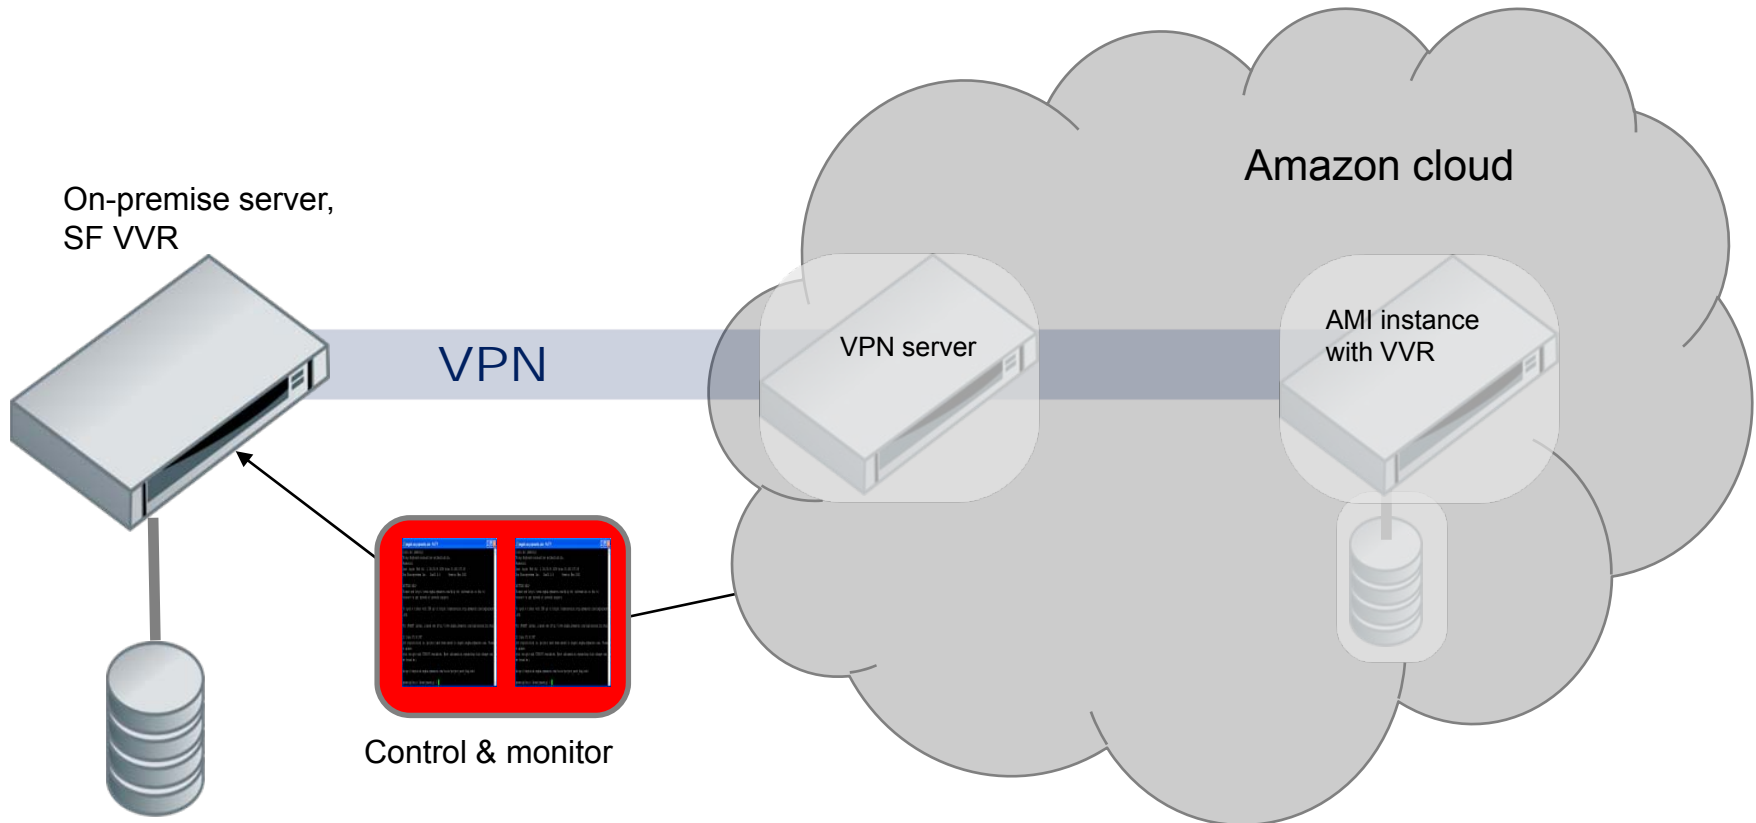


# VCS GCO variation

- ❑ Single-node “clusters” at each location
- ❑ Failover now only over WAN



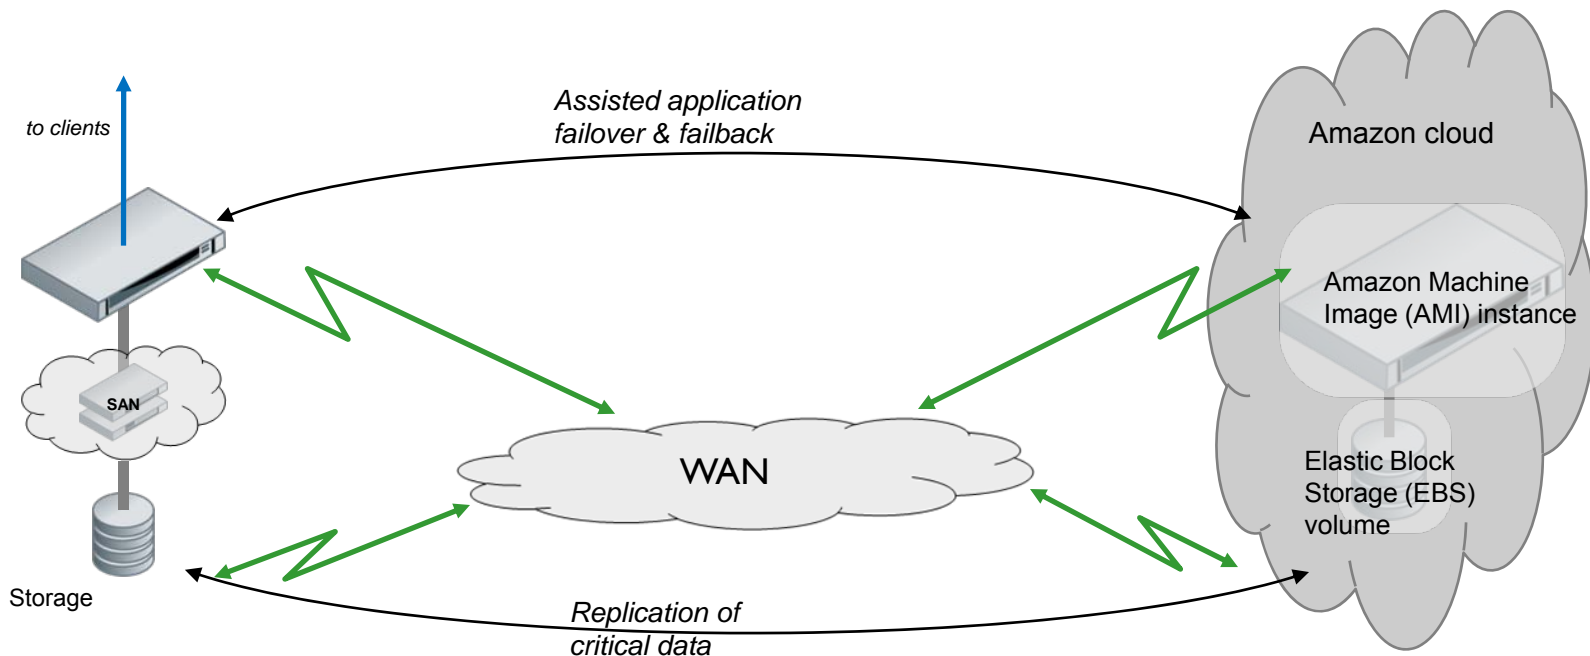
# VVR replication - prototype



- ❑ Veritas Volume Replicator (VVR) for replication to and from the cloud
- ❑ VCS/GCO for failing over an application between on-premise machine and AMI instance in the cloud
  - ❑ Single-AMI cluster is failover target in the cloud
  - ❑ Can serve as the DR solution for multi-node on-premise cluster

# Cloud as a DR site

- ❑ Main data center on user premises
- ❑ Based on a public Linux AMI provided by Amazon
- ❑ Storage provided by EBS volumes



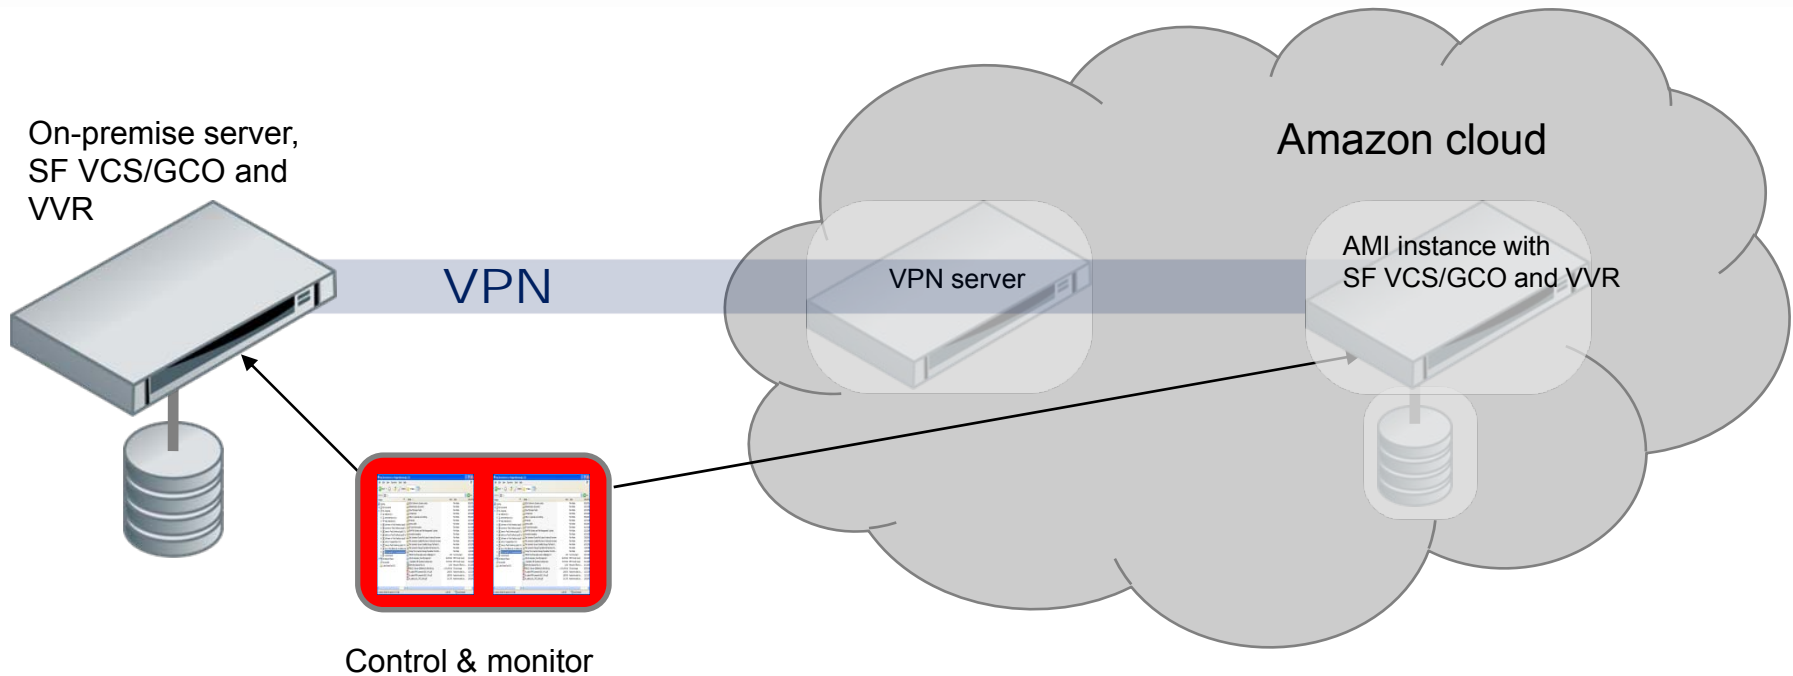
# Steps for VCS/GCO failover

- ❑ GCO heartbeat (usually ICMP ping) times out
- ❑ Cluster in the cloud declares on-premise cluster down
- ❑ Replication target volume in the cloud becomes primary data
- ❑ VCS mounts file system on former replication target volume
- ❑ VCS restarts application on AMI instance in the cloud
- ❑ VCS updates DNS
- ❑ Manual steps: clear DNS cache or application-specific state

# Cloud as a DR site – current status

- Replication
  - from on-premise machine to AMI instance in the cloud
  - from AMI instance in the cloud to on-premise machine
  
- GCO failover with VVR-replicated live data
  - from on-premise I-node cluster to I-node cluster in the cloud
  - “Repair” on-premise machine and fail back from cloud AMI instance to on-premise machine
  
- On Linux and Windows Server platforms

# VCS/GCO failover prototype - issues



- ❑ Corporate firewall: requires open port or use of VPN
- ❑ Requires 24x7 AMI instance for heart-beating
- ❑ EC2 instances support one NIC
- ❑ Clearing stale DNS cache entries or other application-specific non-persistent state



# Cloud as a DR site – performance

- ❑ Difficult to be exact
- ❑ Limit is local connection in our case, better results inside the cloud
- ❑ Sync VVR: ~150 MB/min
- ❑ Amazon: 250+ Mbps, but no guarantees

# In future would like to do

- ❑ More tests using tiering, replication and failover
  - ❑ between on-premise and cloud
  - ❑ in the cloud
- ❑ EBS volumes shared by multiple instances
- ❑ I/O fencing for shared volumes
- ❑ Using other clouds

# Prototyping team

- Matthew Cheung
- Manav Deshmukh
- Chris Lin
- Chitrak Ojha
- Paul Massiglia
- Niranjana Pendharkar
- Ajay Salpekar

**Thank You!**  
**Questions?**