

Advancements in Backup to Support Application Storage on a File Server

Molly Brown

Windows File Server Team

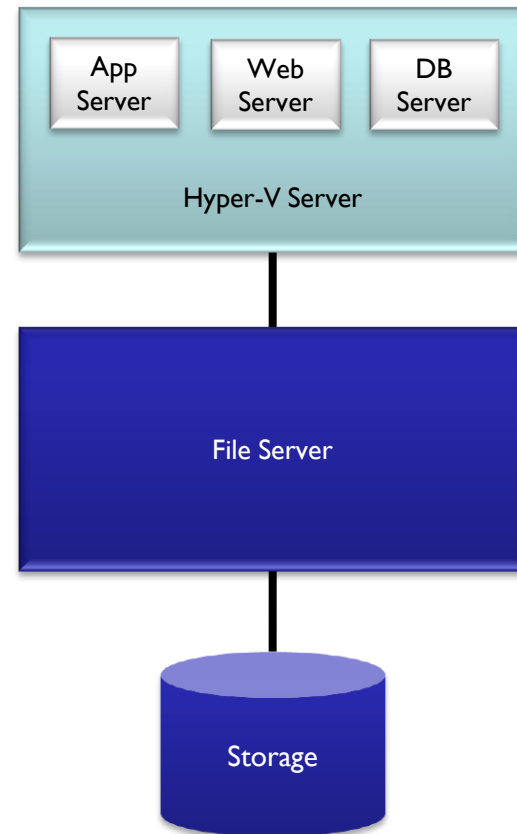
Microsoft Corporation

- ❑ Overview of using file server for application storage
- ❑ Goals for shadow copies for SMB 2.2 file shares
- ❑ VSS background and overview of integration with SMB 2.2 file shares
- ❑ Protocol overview
- ❑ Required changes for software using Volume Shadow Copy Service (VSS) APIs

Application Storage on a File Server

- ❑ What is it?
 - ❑ Server Applications, such as Hyper-V storing their data on SMB 2.2 file shares
- ❑ What is the value?
 - ❑ Added flexibility – dynamic relocation
 - ❑ Easier management – shares vs. LUNs
 - ❑ Lower cost
- ❑ What is the problem?
 - ❑ Application and data are no longer on a single server
 - ❑ Storage-level infrastructure to support consistent data backup only supports local disks
 - ❑ Need to be able to coordinate application consistent view of the data on the file server for backup

- ❑ Example of file based server application storage



□ Goals

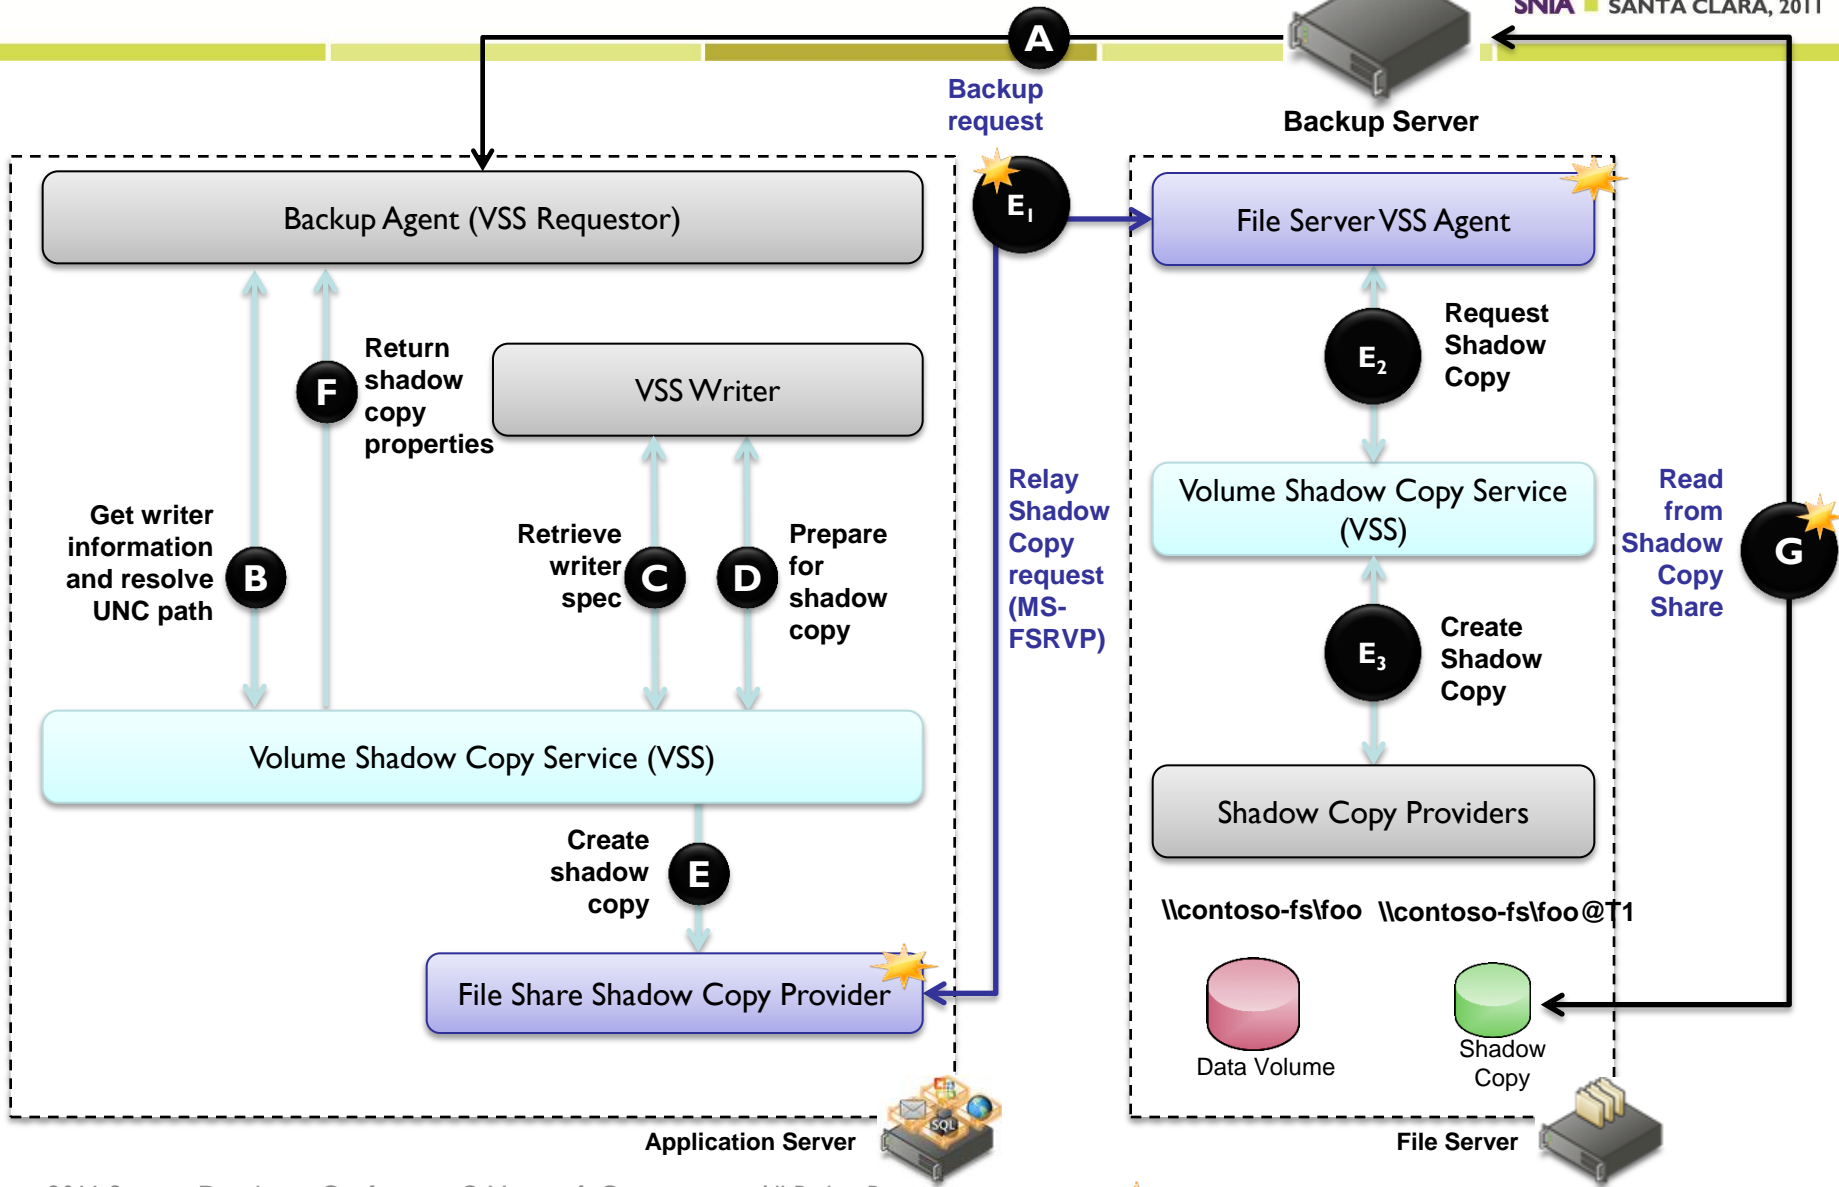
- Provide *application consistent* shadow copies for data stored on remote file shares to support backup and restore scenarios
- Integrate with existing Volume Shadow Copy Service (VSS) infrastructure in Windows
- Have minimal impact on existing VSS ecosystem
 - Backup/restore software ISVs
 - Storage hardware ISVs
 - Enterprise application ISVs

- ❑ A *Requester* is an application that uses the VSS API to request creation and management of shadow copies
 - ❑ Example: Microsoft Data Protection Manager
- ❑ A *Writer* is an application that stores persistent information in files on disk and that provides the names and locations of these files to the requestors
 - ❑ Examples: Hyper-V, SQL, Oracle, Exchange
- ❑ A *Provider* manages the shadow copy volumes or shares and creates the shadow copies on demand

VSS Terminology (continued)

- ❑ A *Shadow Copy* is a snapshot of a volume or share that duplicates all the data at a well-defined point in time*
- ❑ A *Shadow Copy Set* is a collection of shadow copies for various volumes or shares all taken at the same time*
- ❑ The *Shadow Copy Creation Sequence* operates on a shadow copy set
- ❑ (*) The “same time” in this context is not necessarily a consistent timestamp across all the files
 - ❑ Data that will be backed up must be consistent from the application’s point of view
 - ❑ No consistency guarantees are provided for other data on the volume or share

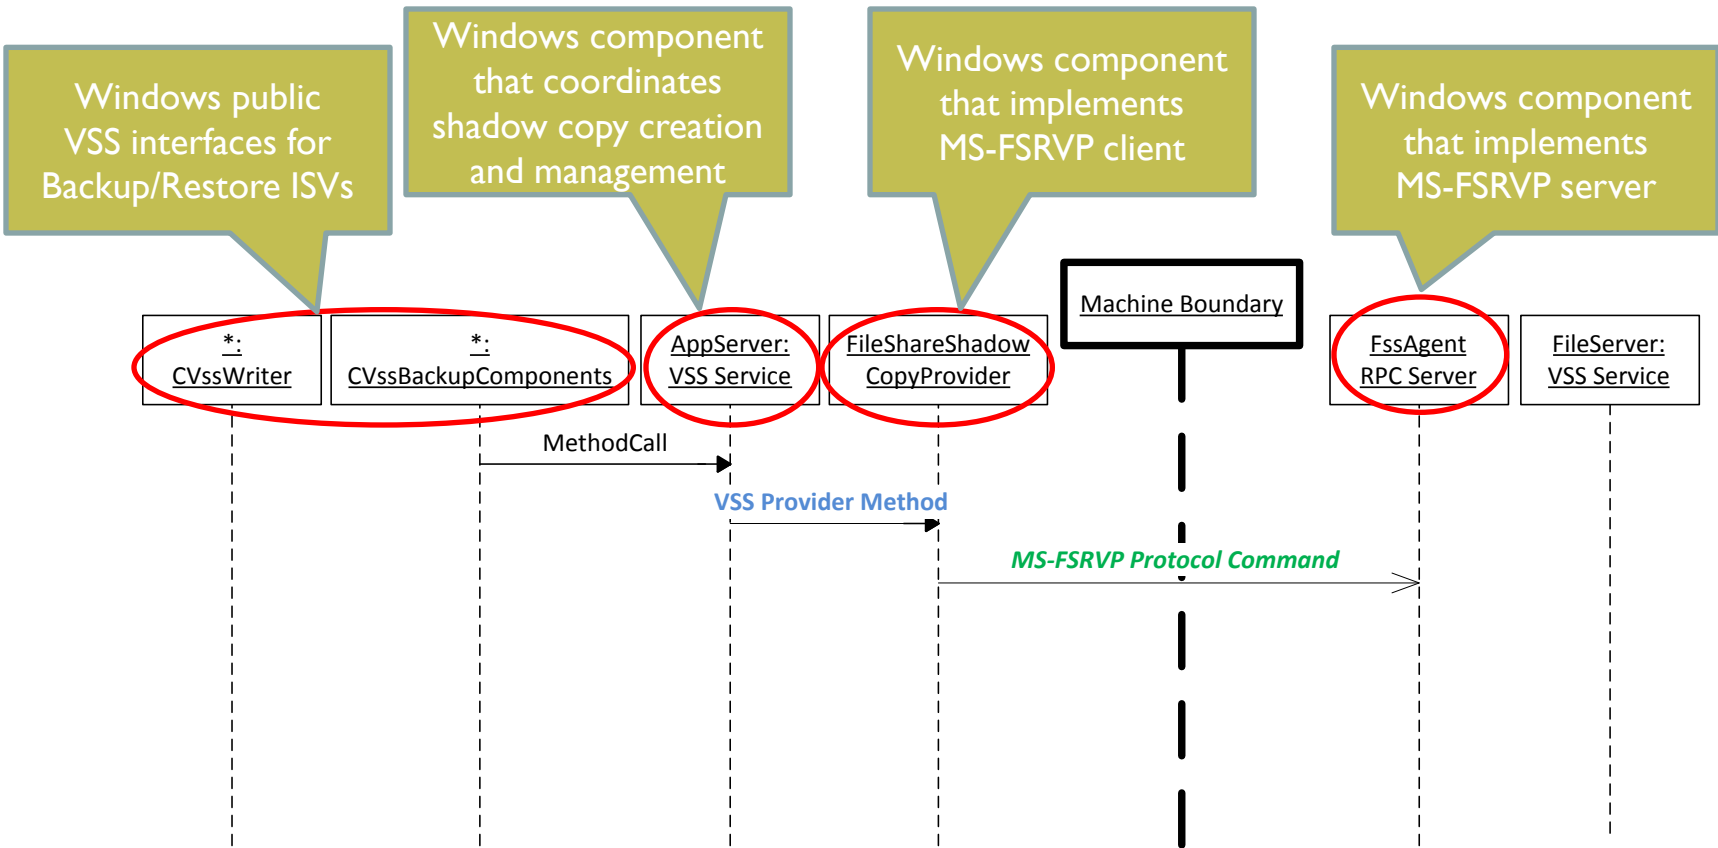
Overview



- ❑ **File Server Remote VSS Provider Protocol**
- ❑ **RPC Protocol**
 - ❑ Uses RPC over named pipes
 - ❑ Pipe name: \\pipe\FssagentRpc
- ❑ 12 commands used to drive the shadow copy creation sequence and management on the file server from the application server

- ❑ Preview protocol document will be posted this week at <http://msdn.microsoft.com/en-us/library/ee941641.aspx>

Sequence Diagram Key

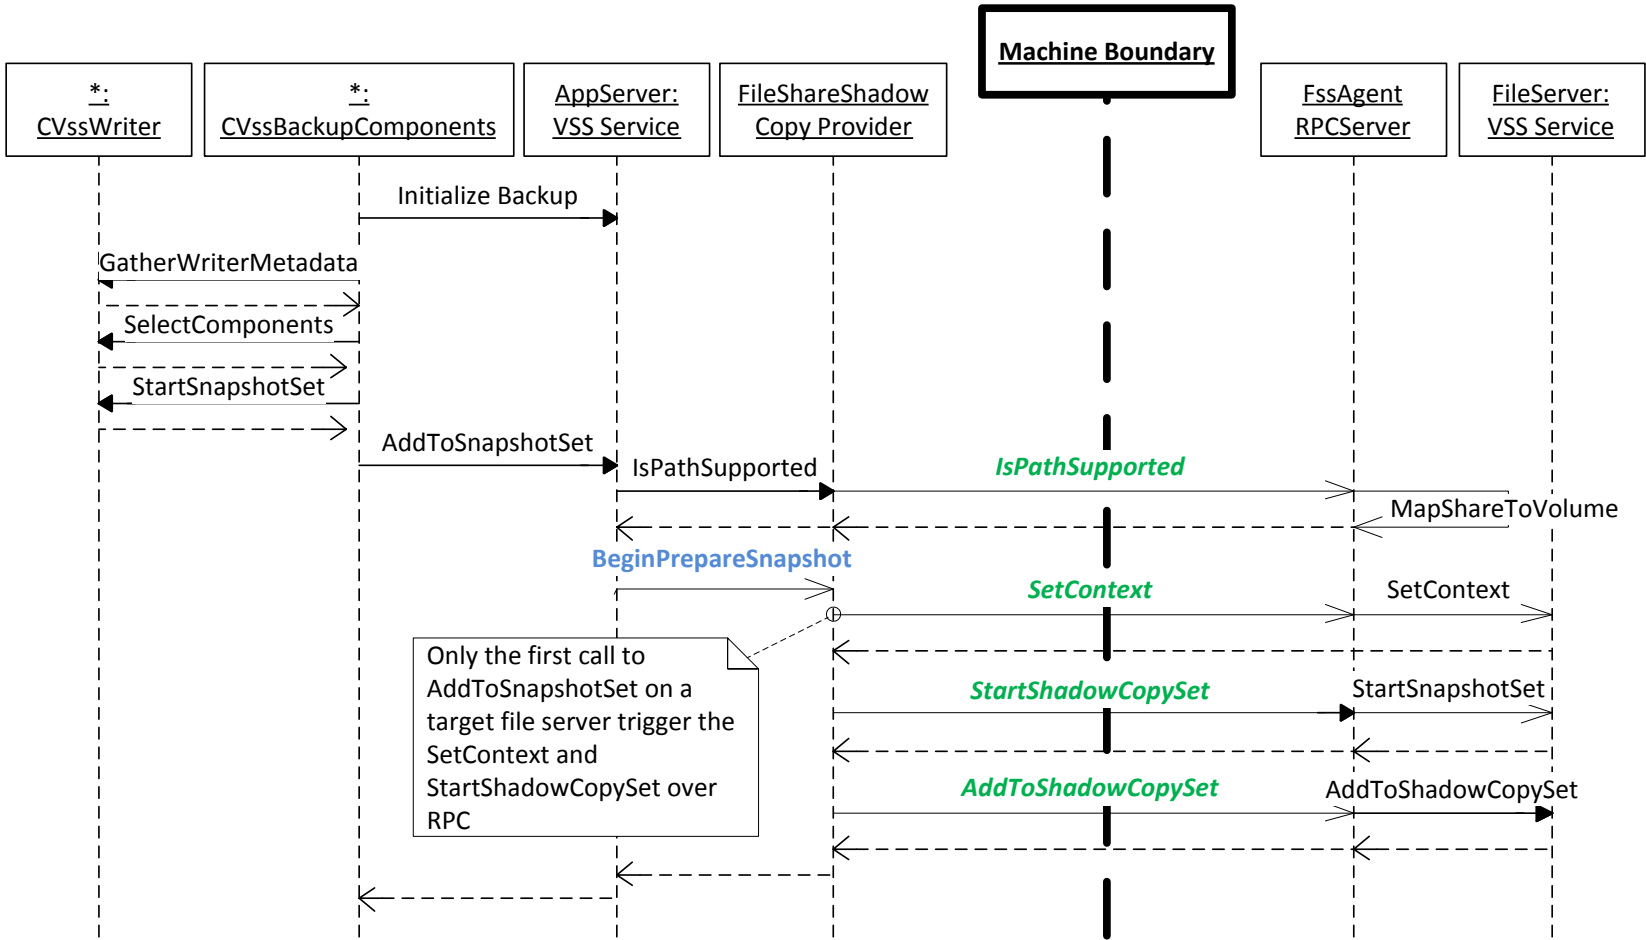


For the purpose of illustrating the flow of the protocol, the next few slides will describe the MS-FSRVP protocol commands and how they interact with the various components in the Windows implementation.

1. `IsPathSupported()`
 - ❑ Verifies that storage backing the path supports shadow copies
2. `SetContext()`
 - ❑ Sets the relevant attributes from the requestor for this shadow copy creation sequence, e.g. persistent, no writers, no auto recover etc.
 - ❑ These settings are then static throughout the rest of the sequence
3. `StartShadowCopySet()`
 - ❑ Creates the shadow copy set on the file server
 - ❑ A file server will receive this once per creation sequence
4. `AddToShadowCopySet()`
 - ❑ Adds a share to an existing shadow copy set
 - ❑ A file server will receive this once per share per sequence, but may received this command more than once if multiple shares are in the shadow copy set

Basic Command Sequence

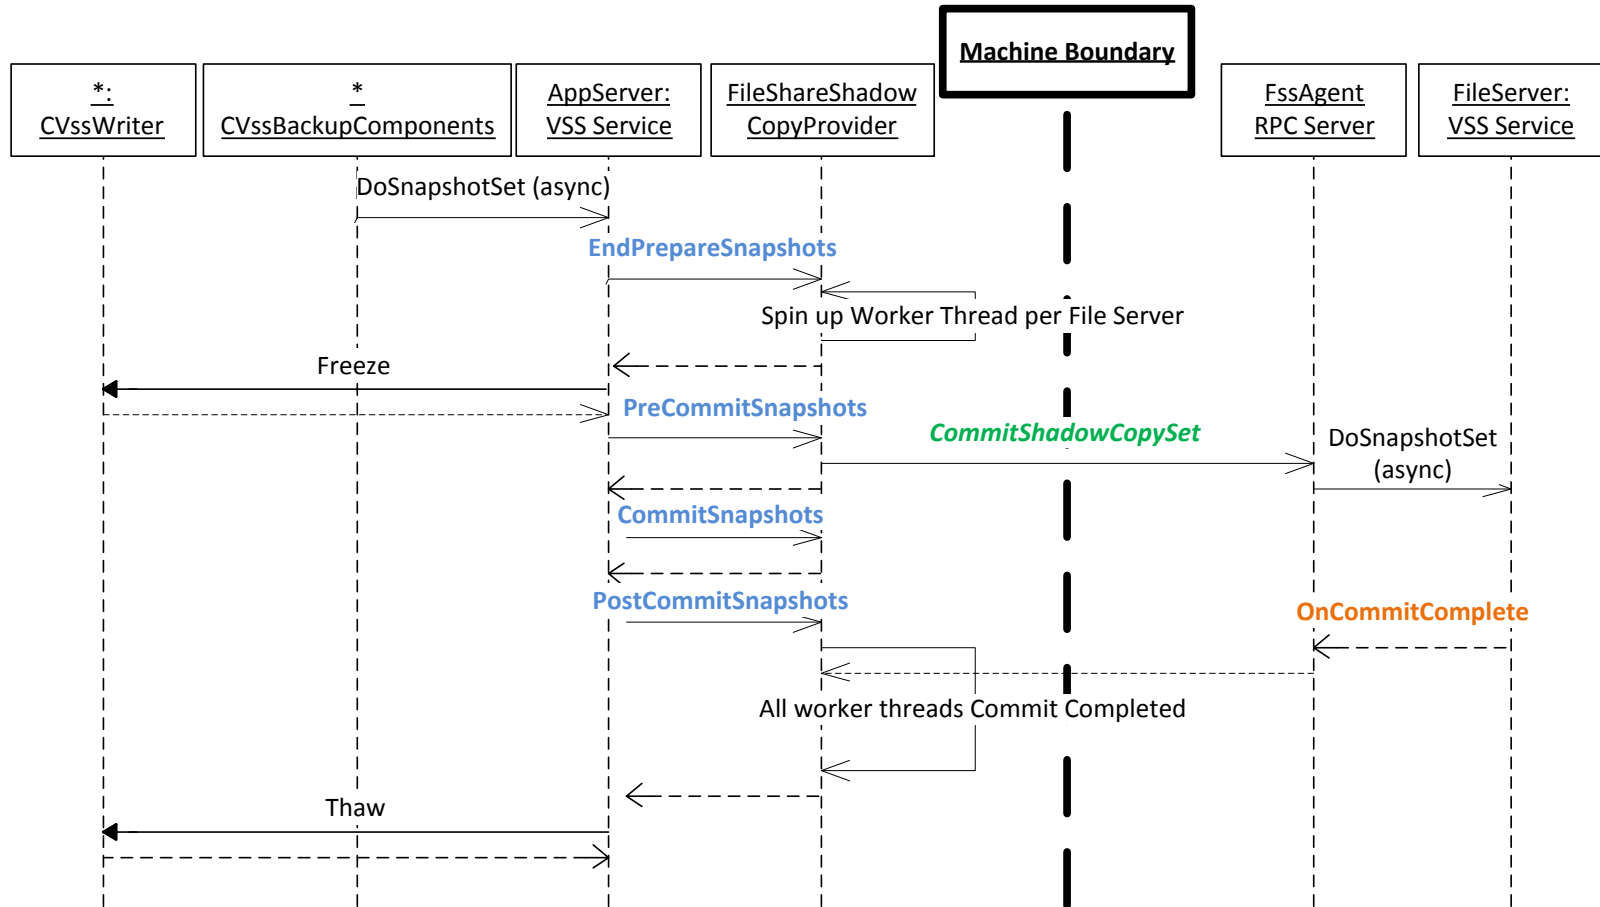
Phase I: Prepare (Sequence Diagram)



- ❑ VSS Service will only hold applications from writing for up to 60 seconds, so this processing is time sensitive
 - ❑ At this point in time, DoSnapshotSet() on the file server has not yet returned
5. CommitShadowCopySet()
- ❑ Application writes are frozen and shadow copy creation processing at the storage level begins

Basic Command Sequence

Phase II: Create (Sequence Diagram)



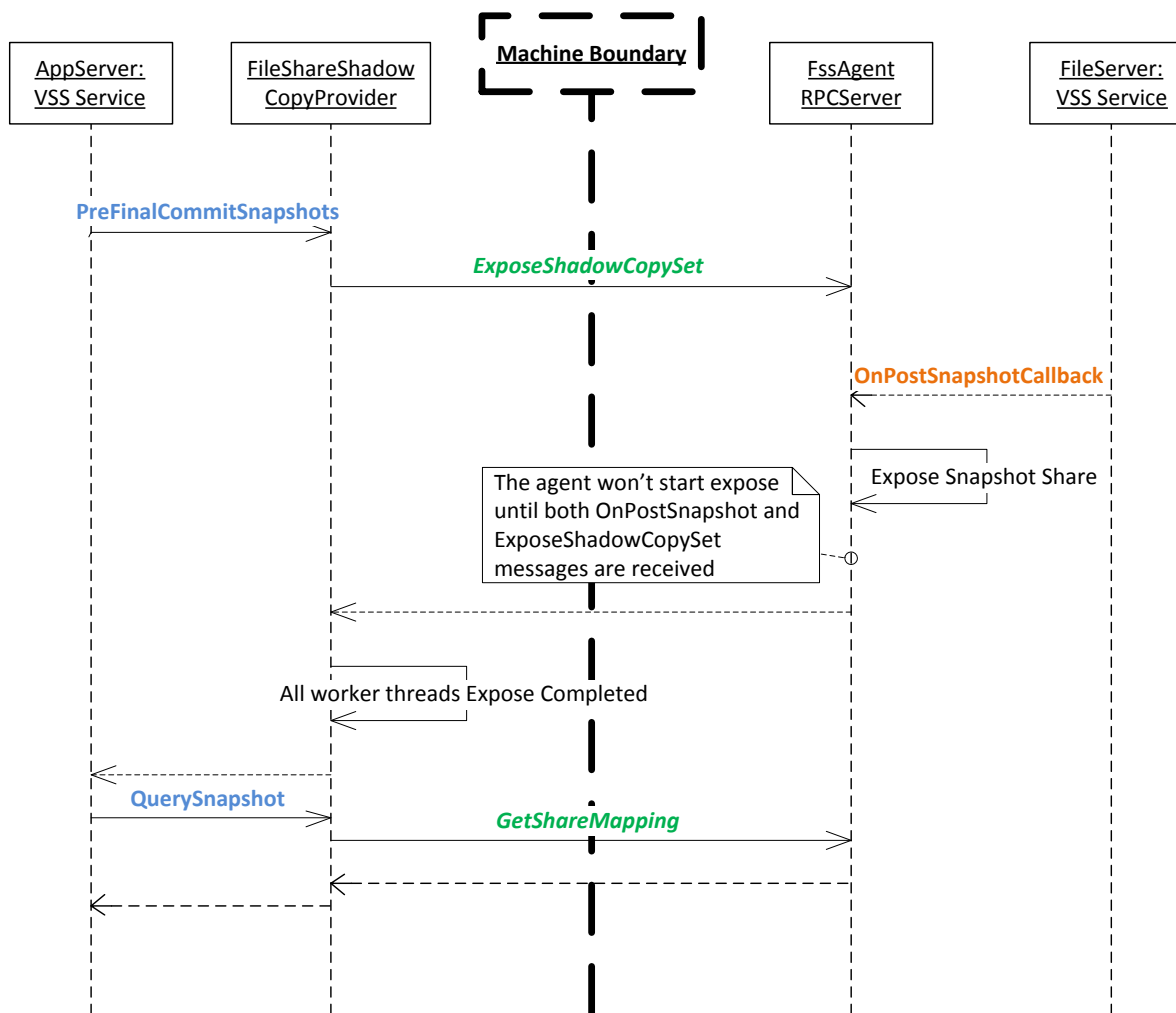
Basic Command Sequence

Phase III: Expose and Query

6. `ExposeShadowCopySet()`
 - ❑ Raises the shadow copies in the set as shares from the file server
 - ❑ Permissions on share match the original share
 - ❑ Applications are now allowed to resume writes
7. `GetShareMapping()`
 - ❑ Returns shadow copy share information that maps to file server shadow copy

Basic Command Sequence

Phase III: Expose and Query (Entity Diagram)



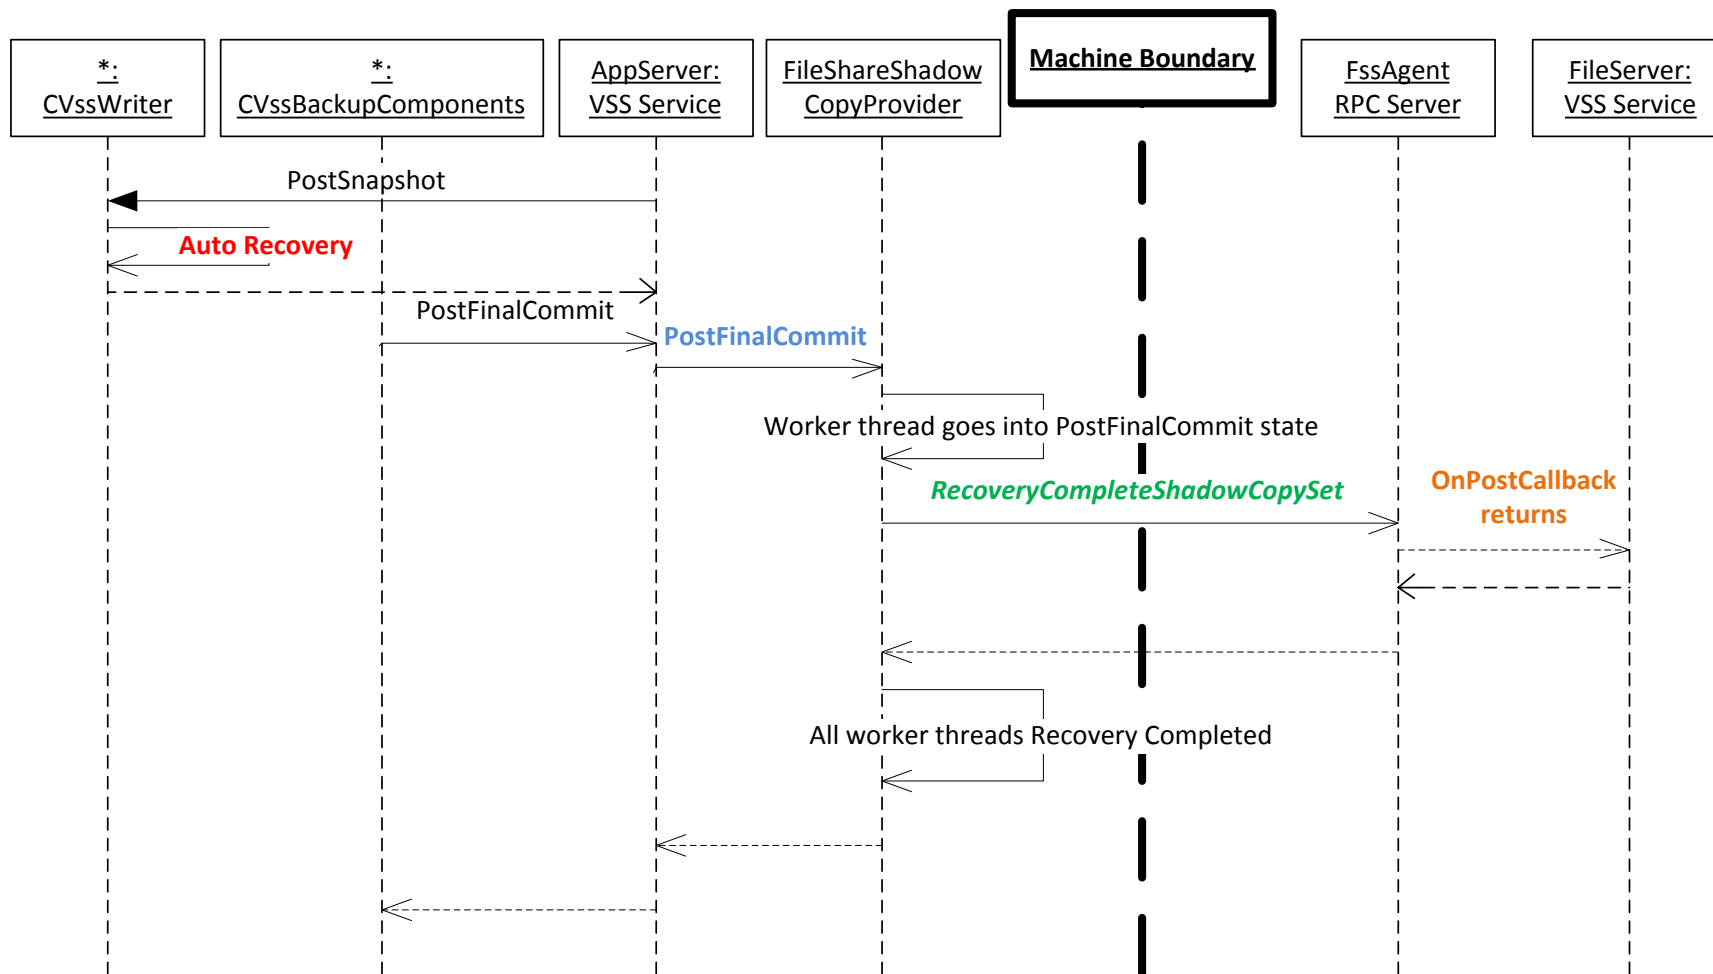
Basic Command Sequence

Phase IV: AutoRecovery

- ❑ AUTO_RECOVERY context flag is set by VSS when it detects that the writer needs to patch the files in the shadow copy
 - ❑ Shadow copy share is exposed with backing shadow copy in read/write mode to allow the VSS Writer to adjust the shadow copy data as required to ensure consistency
-
8. RecoveryCompleteShadowCopySet()
- ❑ Transitions the shadow copy into read-only state and shadow copy creation sequence is complete

Basic Command Sequence

Phase IV: AutoRecovery (Sequence Diagram)



- ❑ `AbortShadowCopySet()`
 - ❑ Aborts the shadow copy creation sequence on the file server
- ❑ `IsPathShadowCopied()`
 - ❑ Determines whether there is existing shadow copies for the specified share
- ❑ `DeleteShareMapping()`
 - ❑ Deletes the shadow copy from a shadow copy set and deletes the exposed share

- ❑ Support shadow copies on shares:
 - ❑ Exposed through DFS-N link targets
 - ❑ Exposed from stand-alone file server
 - ❑ Exposed from a Windows failover or scale-out clustered file server
 - ❑ Exposed from a file server that supports SMB 2.2
- ❑ Support application servers running in stand-alone or clustered configurations
- ❑ Only one shadow copy creation sequence can be running on a file server at a time
 - ❑ From `SetContext()` to `RecoveryCompleteShadowCopySet()`
 - ❑ If server sees a request to start another shadow copy creation sequence, an error will be returned
- ❑ Shadow copies cannot be created on loopback shares

- ❑ On Windows file server, shadow copy shares exposed as:
\\ServerName\ShareName@{ShadowCopyId GUID}
 - ❑ Security for the share should match that of the share the source of the shadow copy
 - ❑ On a clustered file server, the shadow copy share is exposed from one node and is not continuously available nor guaranteed to be accessible after a failover
 - ❑ Even if the target share was accessed via a DFS namespace, the shadow copy share will be exposed from the target machine
- ❑ Multiple data sets can be stored on the same share
 - ❑ A given shadow copy is consistent for the application data set associated with the application instance orchestrating the shadow copy creation sequence

- ❑ Multiple shares can be added to the same shadow copy set
 - ❑ Shadow copy creation sequence across targeted file servers is parallelized as much as possible
- ❑ Mount points nested within file system namespace exposed by share will not be included in shadow copy
- ❑ Shadow copy cleanup
 - ❑ For auto-release shadow copies, deletion is triggered by the VSS on the application server when the backup sequence is complete
 - ❑ For non-auto-release shadow copies, deletion is triggered by the requestor when it is no longer needed
 - ❑ File server has background garbage collection processing to cleanup long-lived shadow copies

Summary of Supported VSS Capabilities

VSS Capabilities	Local Shadow Copies	Remote Shadow Copies
Persistent	✓	✓
AutoRelease	✓	✓
AutoRecovery	✓	✓
Client Accessible	✓	x
Transportable	✓	x
Differential	✓	✓ (if storage supports)
Plex	✓	✓ (if storage supports)
Imported	✓	x
TxF Recovery	✓	x

Required Changes in Software Using VSS APIs

- VSS Requestors
 - Identify a unique set of file shares (UNC paths) that cover the backup set in a similar manner to what is done today to build a unique set of volumes
 - New API to make this processing symmetric for local and UNC paths:
 - `IVssBackupComponentsEx4::GetRootAndLogicalPrefixPaths`
 - Enable impersonation when initializing COM
- VSS Writers
 - Report file shares (UNC paths) as path of interest in their backup set
- VSS Providers
 - No change needed
- No changes required if backup or application ISV do not desire to support UNC paths

Questions?