

Using Windows DFS with SharePoint Remoted BLOBs

Tom E Miller
Chris Geier
Published: 2010



CONTENTS

OVERVIEW AND SCOPE	2
ABOUT STORAGEPOINT.....	2
OVERVIEW OF DFS	2
OVERVIEW OF EBS/RBS	2
BENEFITS OF USING DFS WITH EBS/RBS	3
CONFIGURING DFS FOR USE WITH EBS/RBS	4
DFS Namespaces	4
DFS Replication Groups.....	4
EBS/RBS Storage configuration and DFS	5
FAILOVER SITUATIONS.....	5
EBS/RBS DFS SCENARIOS	6
Single Datacenter	6
Multiple Datacenters	7
MONITORING THE DFS VOLUME	7
WINDOWS SERVER DFS VERSION COMPARISONS	8
Window Server 2008 DFS Features.....	8
Window 2008 DFS Performance Enhancements	9
Domain Functional Level Requirements for Windows Server 2008	9
CONCLUSION	10
APPENDIX A – DFS NAMESPACE COMPARISON	10

OVERVIEW AND SCOPE

SharePoint natively stores both content BLOBs (Binary Large Objects) and their associated metadata within SQL. Since BLOBs are exponentially larger than their metadata counterpart, they add significant overhead which impacts SQL's I/O performance and transactional operations.

Externalizing BLOB's dramatically reduces SharePoint's content database size. Externalizing SharePoint content from its native location within a SQL server content database presents IT organizations with additional backup and data protection options. This paper discusses how to set up and configure Windows Distributed File System for use with StoragePoint.

ABOUT STORAGEPOINT

Metalogix StoragePoint is an easy-to-install, remote BLOB storage and archive solution that allows you to consolidate and optimize SharePoint storage. StoragePoint improves performance, scalability and compliance requirements while decreasing overall storage and administrative costs of your growing SharePoint environment. Use StoragePoint and gain the flexibility to store unstructured SharePoint content on virtually any tier of storage device including SAN, NAS, and cloud storage platforms while improving your search, index and back up windows by 95%.

OVERVIEW OF DFS

Microsoft Windows Distributed File System (DFS) is a Windows Server technology that consolidates multiple server file shares to a logical namespace. This provides many benefits to an IT department as they can allocate additional storage, locate files closer to the users consuming them, and provide redundancy to the end user community with complete transparency. DFS has been part of the Windows Server products for a number of years and has proven to be a reliable technology. The underlying replication technology is the same technology that is used to replicate Active Directory information from one domain controller to another. There are differences to be aware of between versions of your windows OS as well as the level at which your forest/domain operate. See the end of this document for a comparison.

OVERVIEW OF EBS/RBS

EBS is an externalization component that is implemented by SharePoint in WSS 3.0 SP1 and later. The EBS provider is implemented at the lowest point in the SharePoint stack, just above SQL. When the content is ready to be written to SQL as a BLOB the EBS provider takes ownership of the BLOB and returns a token that is used to retrieve the content later. You should understand some key point for EBS

EXTERNAL BLOB STORAGE (EBS)

- ▶ Introduced with WSS 3.0 SP1, supported (but deprecated) in SharePoint 2010.
- ▶ Poorly documented COM interface, not for the faint of heart.
- ▶ Natively implemented as a Farm-wide scoped feature...all BLOBs are externalized.
- ▶ No orphaned BLOB garbage collection.
- ▶ A lot of work and a fair amount of trial and error to build something usable in a production environment.

REMOTE BLOB STORAGE (RBS)

RBS is an externalization component that is implemented by SQL 2008. RBS has nothing to do with SharePoint, except that SharePoint 2010 exposes a RBS provider interface. When enabled, all BLOBs that would normally be written to a varbinary(max) column in a SQL Server table would be handed off to the configured RBS provider. Some key points regarding RBS are:

- ▶ Capability of SQL 2008, only supported in SharePoint 2010.
- ▶ Pretty well documented .NET interface.
- ▶ Natively implemented as a content database scoped feature.
- ▶ Orphaned BLOB garbage collections present by default.
- ▶ Still a lot of work to build a production-ready solution.

BENEFITS OF USING DFS WITH EBS/RBS

EBS/RBS externalizes BLOB content that is natively stored in your SharePoint content database (SQL database) to an external BLOB store. The external BLOB store can be any file, object, or block storage platform that the “plugged in” EBS or RBS provider exposes. Using DFS in conjunction with EBS/RBS has several benefits:

- ▶ Ensure High Availability of the SharePoint BLOB store
 - BLOB location transparency
- ▶ Increase recoverability options while decreasing recovery time by replicating content offsite
- ▶ Lower total cost of ownership by leveraging automated processes and reducing or eliminating the need for magnetic tape

CONFIGURING DFS FOR USE WITH EBS/RBS

DFS NAMESPACES

A DFS namespace is a logical grouping of file shares not bound by physical location or physical servers. Administrators configure namespaces and add physical directories to the namespace. These directories can then be replicated to other physical servers. The use of namespaces provides transparency to end users as they no longer need to know a server name to access a file share. DFS provides for two types of namespaces:

Stand-alone

Stand Alone name spaces are stored only on name servers and not within Active Directory. Each stand-alone DFS namespace has its configuration information stored in the local registry of the root server. Stand-alone namespaces are subject to a single point of failure since when the root target is not available; you cannot access the complete DFS namespace. You do have the option to cluster the server that hosts the One advantage of using a Stand Alone name space is that you can host ten times more folders (up to 50,000 folders) in your namespace. The format of a Stand-alone name space UNC would be as follows:

```
\\ServerName\RootName
```

Domain-based.

An Active Directory based namespace is much more robust as the namespace information is stored and replicated via Active Directory and also cached on each namespace server. The format of a Domain based namespace UNC would be as follows:

```
\\DomainName\RootName
```

Both of these DFS names space types are virtual views of a set of shared folders. When a user views the name space, they can see all of the folders that are part of that name space in a single view. When using DFS with EBS/RBS you want to create a separate Domain-based dedicated namespace specific for use by EBS/RBS as users should not have access to the BLOB store. Configuring a separate namespace provides an additional layer of security around your externalized SharePoint content and hides irrelevant information from the view of the users. Windows Server 2008 supports access-based enumeration of files and folders. This will prevent users without the proper security from seeing any of the BLOB store shares within a DFS Namespace.

DFS REPLICATION GROUPS

A DFS Replication group defines the physical servers and folders that will be involved in replication. There are several options available when creating a replication group. Administrators can define the replication topology, replication schedule, and group types. Bandwidth can also be limited when replicating data as to not bog down networks with just replication traffic.

The replication topology determines the direction that replication occurs. There are three types of replications topologies:

- ▶ Full Mesh – In a full mesh topology every physical server replicates to every other server in the replication group.
- ▶ Hub and Spoke – This topology requires three or more members; otherwise this option is unavailable.

For each spoke member, you can choose a required hub member and an optional second hub member for redundancy. This optional hub ensures that a spoke member can still replicate if one of the hub members is unavailable. If you specify two hub members, the hub members will have a full-mesh topology between them.

- ▶ Custom – Administrators can define custom replication topologies.

Many factors affect your choice of DFS replication groups for use with EBS/RBS. These scenarios will often be dictated by resources available to each organization. See the EBS/RBS DFS Scenarios section for more information. Keep in mind that configuring replication groups and schedules will affect recoverability in the case of a disaster. Data recoverability will be dictated by the amount of time since the last replication, as you will only be able to recover data available as of the last replication. The replication schedule decision should be dictated by the Service Level Agreement you have with your business units.

An additional option in your DFS replication strategy if you are using Windows 2008 R2 is the use of read-only replicated folders. A read-only replicated folder is a folder that is replicated to a particular member on which users cannot add or change any of the files. This can be useful in contexts such as replicating to a test environment or keeping certain repositories in a pristine state. For more information on this topic see:

Make a Replicated Folder Read-Only on a Particular Member

<http://go.microsoft.com/fwlink/?LinkId=155093>

EBS/RBS STORAGE CONFIGURATION AND DFS

Because DFS shares can be addressed the same as any other file shares, with EBS you need only to point your BLOB store location at the DFS share where you want your SharePoint BLOBs to be located. If you are using SQL 2008 with the default FILESTREAM provider and SharePoint 2010 you will need to have a local DFS folder, since the FILESTREAM provider only works with local drives. In this scenario you should only be using one way DFS replication to replicate the blob store off site. This limitation may not apply if you are developing your own adapter/provider. If developing your own adapter you should build it to allow you to simply configure the adapter to use the DFS share name and SharePoint content will be written to that location. Making sure of this will ensure that you do not require any additional configuration to leverage a DFS volume when implementing EBS/RBS.

FAILOVER SITUATIONS

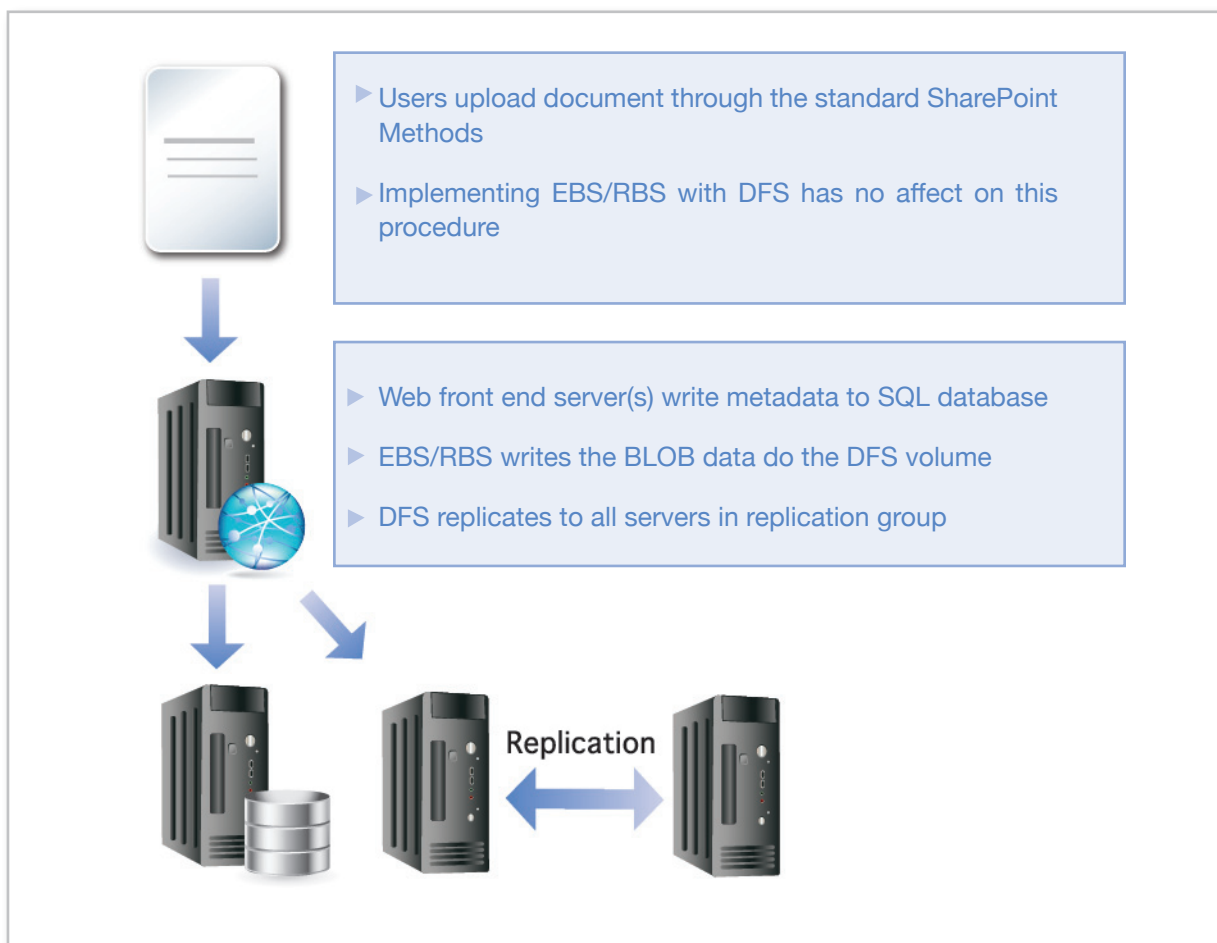
When a client connects to a DFS share they are directed to a physical server within the DFS namespace. DFS automatically redirects clients to a different physical server when the server they are attached to fails. With EBS, each SharePoint web front end server is considered the client to the DFS share. These clients will failover to the next nearest physical server based on the Active Directory Sites configuration. During this failover period, a SharePoint user may get an error message when working with externalized content. If this error occurs, a user will need to retry the operation that generated the error message. It can take up to a minute for the web front end to fail over to the next physical server. These situations will need to be planned for and communicated properly to users. One important factor when you are using the FILESTREAM provider with RBS is that the primary DFS server needs to be located on the SQL server. These blobs are replicated out from there on schedule in a 1 way replication. As such if the SQL server fails or becomes for any reason unavailable having the BLOBs externalized and replicated will not provide any advantage. However if you are writing your own RBS based provider you can overcome this limitation and gain the additional advantage of DFS. Even with the existing limitation of the FILESTREAM RBS provider there is still advantage in the speed and ease of recoverability by using DFS as you will see in the following scenarios.

EBS/RBS DFS SCENARIOS

SINGLE DATACENTER

A single datacenter is the simplest configuration option to consider when using DFS with EBS/RBS. In this scenario DFS can be used to ensure high availability of the BLOB store within the same datacenter by replicating externalized content to a separate physical server. In this scenario there is no location redundancy provided. If you need to have true disaster recovery you should not chose this option. To setup this configuration the following tasks should be completed:

- ▶ Configure a dedicated namespace in Active Directory
- ▶ Configure a replication group on two or more physical servers
- ▶ Mesh replication topology should be used
- ▶ Consider using RAID on each of the physical servers in accordance to performance needs, SLA requirements, and budget
- ▶ Be certain to backup all of the replicated data to an offsite location to ensure data is recoverable in the event of a disaster.

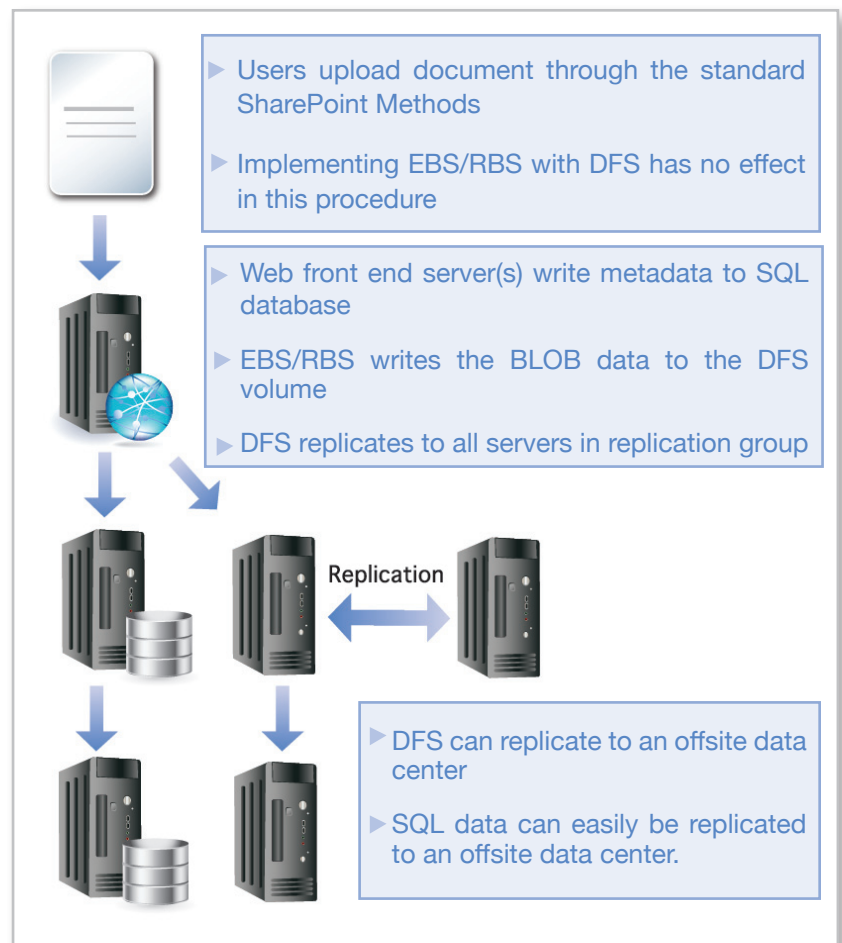


MULTIPLE DATACENTERS

Multiple datacenters can offer additional benefits as content can be replicated offsite. This has the potential to eliminate the some or all need for tape backup and can save thousands of dollars per year. Replicating content to a DR site can be scheduled and coordinated with any SQL database backups.

Microsoft recommends that each SharePoint server in a farm should be physically located in the same datacenter. Replicating the SharePoint servers as well as content off-site could be used to ensure recoverability. There are several additional options that this can be used in conjunction with to achieve high availability such as SQL log shipping since the SharePoint content databases and transaction logs are considerably smaller since the BLOB data has been externalized. To implement this scenario the following tasks should be completed:

- ▶ Configure a dedicated namespace in Active Directory
- ▶ Configure a replication group
 - Include at least one server in the remote site
 - Consider using one way replication.
 - If there are bandwidth constraints, consider setting a replication schedule that will not interfere. *Ensure that you have a proper SLA to support this*
- ▶ Consider using RAID on each of the physical servers in accordance to performance needs, SLA requirements, and budget.
- ▶ Consider enabling compression in the storage profile. This reduces the overall bandwidth requirements when replicating content off-premise



MONITORING THE DFS VOLUME

A common method of monitoring DFS and or reporting on DFS health is through the DFSSAdmin command line utility, or through the DFS Management interface. With the command line interface if DFSSAdmin you can script and automate these health reports. The recommended way to automate DFS Replication health report generation is to schedule a regular task that uses DfsrAdmin.exe to generate a health report for each replication group that an organization wishes to monitor. Beyond the out of the box functionality organizations may choose

to use a more robust infrastructure such as Microsoft System center for monitoring and reporting which is available as a management pack

<http://technet.microsoft.com/en-us/library/dd261970.aspx>

There are also performance counters that can be accessed through performance monitor:

DFS Namespace Service API Queue. Shows the number of requests (made using the NetDfs API) in the queue for the DFS Namespace service to process.

DFS Namespace Service API Requests. Shows performance information about requests (such as creating a namespace) made to the DFS Namespace service.

DFS Namespace Service Referrals. Shows performance information about various referral requests that are processed by the DFS Namespace service.

WINDOWS SERVER DFS VERSION COMPARISONS

There are several differences in DFS functionality and performance within different windows versions. Windows Server 2008 offers several improvements and additional features over Windows Server 2003 R2. The chart below details additional features available only in windows server 2008. While Windows Server 2008 is not required for use with EBS/RBS and DFS it is highly recommended due to its improvements in performance.

WINDOW SERVER 2008 DFS FEATURES

Feature	Description	EBS/RBS Benefit
Access-Based Enumeration	Users cannot see files and folders to which they do not have access.	An existing DFS namespace can be used; however it is still recommended that a dedicated namespace be created to host SharePoint BLOB information.
Improved Scalability and Performance	Windows Server 2008 offers several scalability and performance improvements over Windows Server 2003 R2. See the chart below for a comparison.	The improved performance gives administrators flexibility when choosing specific scenarios and replication topologies.
Content Freshness	DFS Replication in Windows Server 2008 has a new feature called Content Freshness, which prevents a server that was offline for a long time from over-writing fresh data when it comes back online with stale (out-of-date) data	This can be critical to ensure data integrity in a failover situation.

WINDOW 2008 DFS PERFORMANCE ENHANCEMENTS

Windows Server 2003 R2	Windows Server 2008
Multiple RPC calls	RPC Async Pipes (when replicating with other servers running Windows Server 2008)
Synchronous inputs/outputs (I/Os)	Asynchronous I/Os
Buffered I/Os	Unbuffered I/Os
Normal Priority I/Os	Low Priority I/Os (this reduces the load on the system as a result of replication)
4 concurrent file downloads	16 concurrent file downloads. *This affects replication of data and not data access by users.

Additional Windows Server 2003 DFS Information:

[http://technet.microsoft.com/en-us/library/cc787066\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc787066(WS.10).aspx)

Additional Windows Server 2008 DFS Information:

[http://technet.microsoft.com/en-us/library/cc753479\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc753479(WS.10).aspx)

DOMAIN FUNCTIONAL LEVEL REQUIREMENTS FOR WINDOWS SERVER 2008

A domain must meet these minimum requirements to use the additional DFS features of Windows Server 2008:

- ▶ The forest uses the Windows Server 2003 or higher forest functional level.
- ▶ The domain uses the Windows Server 2008 or higher domain functional level.
- ▶ All namespace servers are running Windows Server 2008.

IMPORTANT! Do not attempt to create a domain-based namespace using the Windows Server 2008 mode unless the forest functional level is Windows Server 2003 or higher. Doing so can result in a namespace for which you cannot delete DFS folders, yielding the following error message: "The folder cannot be deleted. Cannot complete this function."

CONCLUSION

The use of DFS can play a significant role in providing additional levels of availability and recoverability of file system storage. This is true with or without the use of EBS/RBS. Because EBS/RBS can store SharePoint content (BLOBs) in the file system, having DFS in place helps your SharePoint environment take advantage of this as well, resulting in an increased ability to achieve higher levels of both availability and recoverability. In the case that you want to leverage a third party EBS/RBS provider,

StoragePoint is a storage solution for Microsoft SharePoint that externalizes the content BLOBs normally stored in SharePoint content databases. It makes SharePoint easier to manage and scale, while improving performance and without impacting the end user. It provides advanced capabilities that you will not find in the RBS FILESTREAM provider and other entry level or freeware solutions. You can request a free 30 day trial at Metalogix.com.

APPENDIX A – DFS NAMESPACE COMPARISON

Characteristic	Domain-based Namespace	Stand-Alone Namespace
Path to namespace	\\NetBIOSDomainName\RootName \\DNSDomainName\RootName	\\ServerName\RootName
Where namespace information is stored	In Active Directory and in a memory cache on each namespace server	In the registry and in a memory cache on the namespace server.
Namespace size recommendations	We recommend that you keep the size of the namespace object in Active Directory to less than 5 megabytes (MB) by using fewer than 5,000 folders with targets in domain-based namespaces. (The namespace object is where namespace metadata is stored in Active Directory.)	The largest recommended namespace size for a stand- alone namespace is 50,000 folders with targets.
Supported methods to ensure namespace availability	Use multiple namespace servers to host the namespace. (The namespace servers must be in the same domain.)	Create a stand-alone namespace on a server cluster using the Cluster Administrator snap-in.
Support for using DFS Replication to replicate folder targets	Supported	Supported

GLOSSARY

BLOB – A Binary Large Object is a collection of binary data stored as a single entity in a database management system

Network Addressable Storage (NAS) – A device connected to a network to provide storage for other network clients

Storage Area Networks (SAN) – A SAN is a network of devices optimized to provide storage capacity for servers and other IT infrastructure hardware

Cloud Storage – Network storage that is accessible through API's. Cloud storage is typically provided by a third party.

ABOUT THE AUTHORS

Tom E Miller

Tom has more than 13 years of experience within the IT industry with a strong focus on ECM technologies. Since 2001, he has architected/designed dozens of SharePoint solutions as well as implemented several SharePoint-based 3rd party tools. Other areas of expertise include BPM/Workflow, Capture/Imaging, Network Architecture and Administration. Tom currently works as a StoragePoint Senior Technical Specialist

Chris Geier

Chris Geier is a 15-year veteran of the technology industry and specializes in all things Microsoft. He was introduced to SharePoint in 2001 while working for Microsoft services. Currently Chris is the StoragePoint Product Manager at Metalogix. Chris is a participant in, and advocate for, the SharePoint community, as well as a regular participant/speaker at SharePoint Saturday and other grass roots events. When he's not twiddling with the latest gadgets, software and technology, you'll probably find Chris on a bleacher cheering his budding baseball and gymnastics stars. Chris and his wife Sara have four kids and live in the suburban Chicago area.

ABOUT METALOGIX

Metalogix is the trusted provider of innovative content lifecycle management solutions for Microsoft SharePoint, Exchange and Cloud platforms. We deliver high-performance solutions to scale and cost-effectively manage, migrate, store, archive and protect enterprise content. Metalogix provides global support to thousands of customers and strategic partners and is a Microsoft Gold Partner, a managed partner in Microsoft's High Potential ISV Group and GSA provider. Metalogix is a privately held company backed by Insight Venture Partners and Bessemer Venture Partners.

METALOGIX

5335 Wisconsin Ave NW, Suite 510, Washington DC 20015
sales@metalogix.com | www.metalogix.com | 1.877.450.8667