

SteelFusion™ with AWS Hybrid Cloud Storage

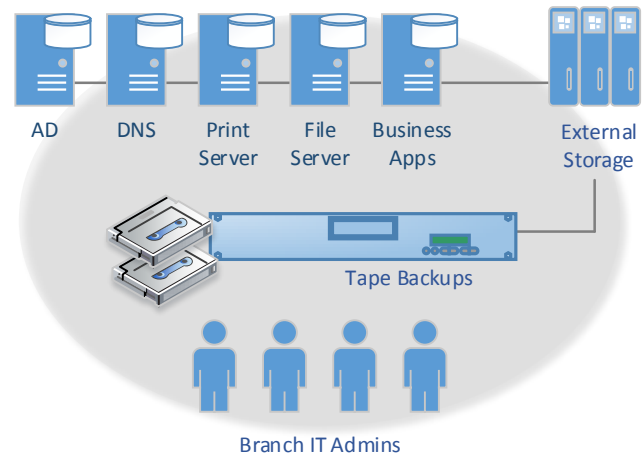
March 2016



The Challenge

According to IDC, to meet the demands of global customer and global talent requirements, companies have to maintain remote offices.

These islands of distributed branch infrastructure have been necessary to meet local performance and reliability needs to ensure the productivity of these remote offices; however, they are costly and difficult to maintain. Centralizing and consolidating data is key to eliminating these issues.



Branch Infrastructure

Branches are generally the revenue generator for most businesses. It is imperative that companies make sure that their branch operations are always up and running. In addition to protecting and securing data, they should have the ability to recover from disasters in minimum amount of time and least effort from personnel. Since branches and data center resources are generally located over the WAN with long distances, companies today deploy each of these services at every branch, either as physical or virtual servers. Generally, the following services are required to keep a branch up and running:

- Active Directory (AD)
- Domain Name System (DNS)
- Dynamic Host Configuration Protocol (DHCP)
- File Server
- Business Applications

Business Challenges

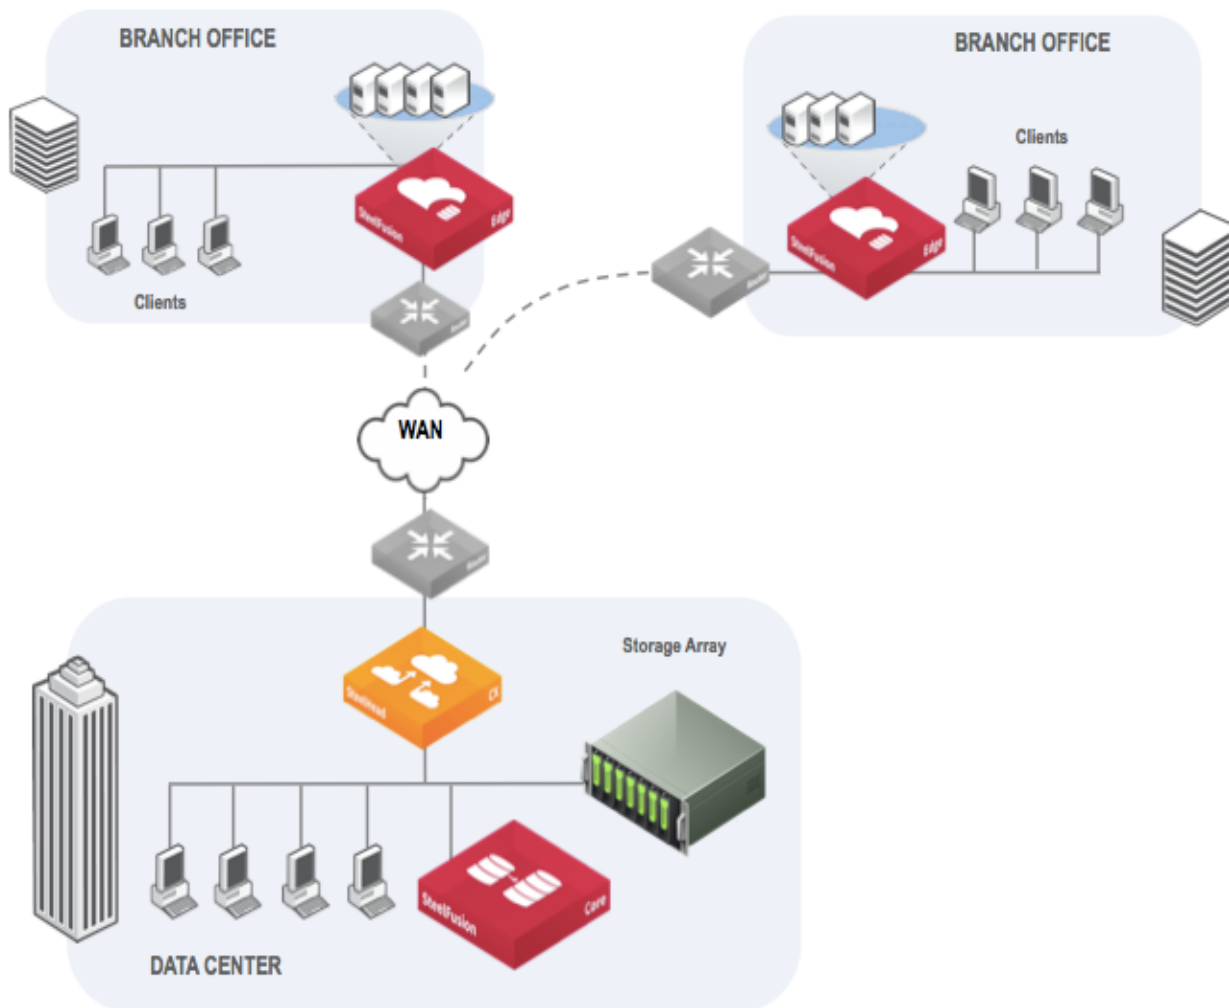
According to IDC, to meet the demands of global customer and global talent requirements, companies have to maintain remote offices. Companies are spending over \$4 billion on remote office IT support. These islands of distributed branch infrastructure have been necessary to meet local performance and reliability needs to ensure the productivity of these remote offices; however, they are costly and difficult to maintain. Because of these inefficiencies, companies rarely can afford to have the expertise in branches to maintain operation and protect data in such distributed infrastructure.

Therefore when branch offices suffer outages due to natural or manmade disasters, productivity and data are compromised and company revenue is impacted. Centralizing and consolidating data is key to eliminating these issues; however only 8 percent of branch offices consolidate data in the data center, which increases companies' exposure to data theft, data loss and downtime due to outages or natural disasters.

Chapter 1: Solution Overview

Riverbed SteelFusion is a branch converged infrastructure solution, encompassing server, storage, networking, and WAN optimization. SteelFusion can be fully deployed and administered from a central location, eliminating need for remote onsite IT. SteelFusion allows for the consolidation of all data/storage (server OS's, application, and data) into the datacenter, where it can be managed and protected (backed up), while projecting it out the branch. This is accomplished without sacrificing any of the benefits (performance & access) of having servers and data at the Edge, close to end-users. SteelFusion allows near-instant continuance of operations at any branch or remote location, guarantees the highest operational levels with little to no loss in productivity. With SteelFusion, businesses can restore operations in a matter of minutes vs. days, centrally protect and secure data, and significantly lower the TCO of branch and remote offices. SteelFusion appliances can help consolidate distributed data, improve security, and reduce administration for managing remote / branch office environments. When utilized with data center storage, SteelFusion can expose storage from data center to remote branch offices to deliver all branch services.

A SteelFusion solution consists of two components, Core (for Data Center) and Edge (for branch). SteelFusion presents block storage to applications and hosts at the branch across a WAN, while the underlying physical storage is provisioned at the data center. To an application server and file system running at the Edge, data center storage mapped on the SteelFusion appliance in the branch appears just like a local block-storage device.



Chapter 2 SteelFusion with AWS

Industry Trend

As customers look to leverage Hybrid cloud from a datacenter perspective, SteelFusion (i.e. SF) plays a key role in extending the hybrid cloud to the branch locations. SF helps customer's leverage the new IT models enabled by cloud. Customers can leverage cloud as a storage tier where they can move less critical data to the Cloud. They can leverage cloud as a compute farm where the compute comes from the cloud but customers retain the data in their datacenters. Or they can "move" their datacenter to cloud such that the branches connect directly to the cloud.

Customer Trend

Recently there was a survey done asking about 1000 organizations about their Cloud adoption. Here are some of the highlights

- Hybrid Cloud remains the preferred strategy - 82% of the organizations had a hybrid cloud strategy in place (up from 70% in 2014)
- Public Cloud leads in Breadth of Enterprise Adoption while Private Cloud leads in workloads - 88 percent of enterprises are using public cloud while 63 percent are using private cloud. 13 percent of enterprises run more than 1,000 VMs in public cloud, while 22 percent of organizations run more than 1,000 VMs in private cloud
- Plenty of headroom for more Cloud workloads - 68 percent of enterprises run less than a fifth of their application portfolios in the cloud. 55 percent of enterprises report that a significant portion of their existing application portfolios are not in cloud, but are built with cloud-friendly architectures.
- Enterprise Central IT Teams make the key Cloud decisions - 62 percent of enterprises report that central IT makes the majority of cloud spending decisions. 43 percent of IT teams are offering a self-service portal for access to cloud services, with an additional 41 percent planning or developing a portal.

About 80% of all organizations now have a hybrid cloud strategy. Organizations are looking to the public cloud as a differentiator for many different use cases - consuming cheap and deep storage in the cloud to augment their on premises datacenter(s), replacing their on premises datacenter with the public cloud, using the public cloud as a compute farm. These customers are looking to SteelFusion to enable their ROBO locations to leverage the several benefits of the public cloud. It makes even more business sense to store your branch data in the cloud.

(Source: Cloud Computing Trends: 2015 State of the Cloud Survey)

SteelFusion Solution

SteelFusion has a very unique value proposition of consolidating data from the branch locations back to the Datacenter (without compromising on availability or performance for applications running at the branch locations). No other competitive solution can offer that value.

For the Cloud use cases, competitors only offer customers the option to consolidate branch data to the Cloud to use it as cheap and deep storage using a cloud gateway in every branch location which makes the solution – cost ineffective and an operational burden. SteelFusion provides customers the opportunity to leverage the Cloud as primary storage or an augmentation to their primary storage on premise datacenter.

With public cloud integration SteelFusion will provide customers the flexibility as to where they want to store their data - on premise datacenter and/or Cloud. SteelFusion wants to enable customers to store their data wherever it makes the most business sense. Thus it is not a technology decision but a business decision, which will enable customers to succeed in their cloud journey.

Customer wants to augment their current on premise datacenter with public cloud storage. Seamlessly move data to the Cloud from the on premise datacenter leveraging robust infrastructure at the Datacenter. User has the ability to project public cloud storage (in addition to datacenter storage) from their Datacenter to their many branch locations. User has all the capabilities (redundancy, availability, data protection, security etc.) in this environment (similar to on premise datacenter storage to the branch)

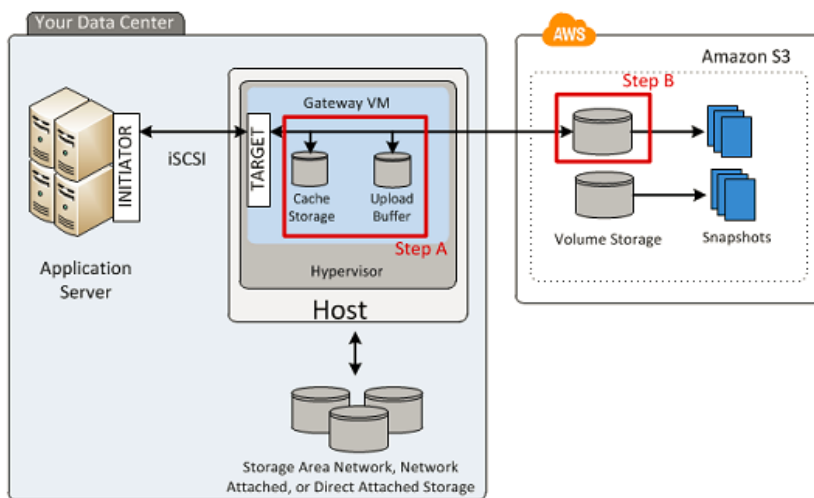
Amazon Web Services Storage Gateway Storage Concepts Review

The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. The service allows you to securely store data in the AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data on-premises while securely storing all of your data encrypted in Amazon Simple Storage Service (Amazon S3) or Amazon Glacier.

The AWS Storage Gateway delivers 2 configurations that are compatible with SteelFusion:

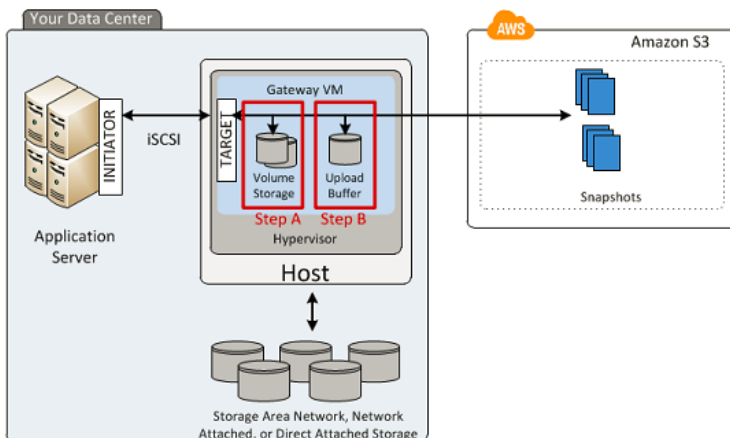
Gateway-Cached Volumes:

You can store your primary data in Amazon S3, and retain your frequently accessed data locally. Gateway-Cached volumes provide substantial cost savings on primary storage, minimize the need to scale your storage on-premises, and retain low-latency access to your frequently accessed data, as shown below.



Gateway-Stored Volumes:

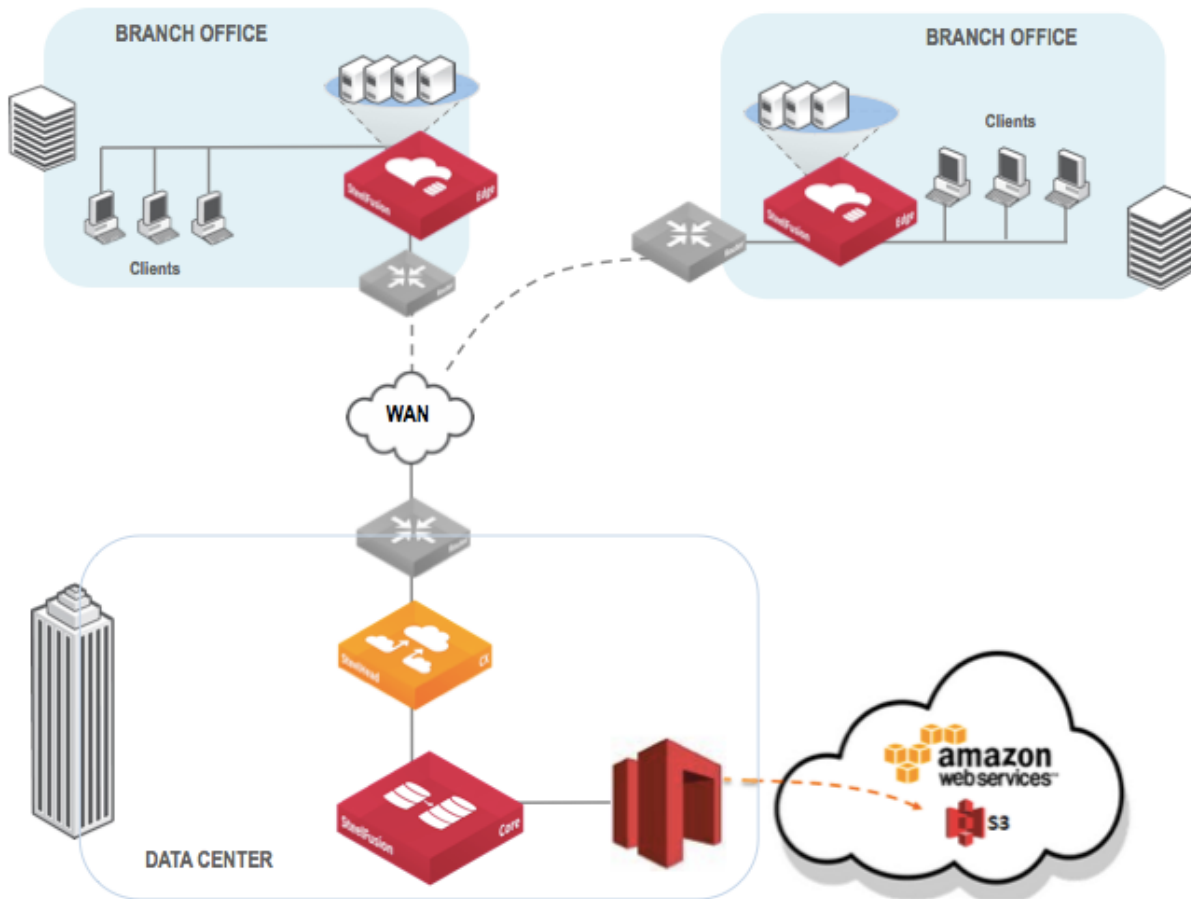
In the event you need low-latency access to your entire data set, you can configure your on-premises gateway to store your primary data locally, and asynchronously back up point-in-time snapshots of this data to Amazon S3 as shown below. Gateway-Stored volumes provide durable and inexpensive off-site backups that you can recover locally or from Amazon EC2 if, for example, you need replacement capacity for disaster recovery.



A New Paradigm For Delivering Hybrid Cloud Branch Services

SteelFusion, when used in conjunction with an AWS Storage Gateway and Amazon cloud storage, delivers unparalleled data availability, flexibility and protection for branch infrastructure, while simultaneously reducing data center storage requirements and improving data availability and recoverability by leveraging cheap, elastic cloud storage delivered via Amazon cloud storage services. By using an AWS Storage Gateway with SteelFusion, businesses can begin the transition to a storage-less data center, where all storage is served by the cloud on demand, through SteelFusion, to applications and users.

Branch offices with constrained network links now enjoy all the same cloud storage benefits that data centers enjoy through the use of SteelFusion and Amazon without compromising data, performance, and application availability. The problems of delivering cost effective storage to branches that can be managed and controlled centrally has been improved, allowing for a storage tier that can efficiently and effectively scale with business requirements. Rebuilding damaged or destroyed branch office environments is no longer a days-long recovery time objective, but one that can be compressed to as little as an hour, even though data resides two links away within Amazon cloud storage. And with the ability to utilize Amazon Direct Connect capabilities, SteelFusion can consolidate storage services away from the data center into managed public or private clouds, allowing businesses to access cloud storage at high bandwidth speeds and without the need to have a true data center storage environment.



Chapter 3: Deploying SteelFusion Appliances with AWS Storage Gateway

This section provides step-by-step instructions on configuring the entire backup solution.

Deployment Prerequisites

- An Amazon Web Services Storage Gateway licensed and configured in either gateway cache or gateway stored mode.
- Administrator access to the Amazon Web Services Storage Gateway and SteelFusion appliances to make changes such as enabling iSCSI, adding initiator groups, etc.
- SteelFusion Core and SteelFusion Edge appliances installed and powered up.

AWS Storage Gateway Deployment Steps

1. From the [Amazon web services storage gateway page](#), download and install the AWS Storage Gateway virtual machine (VM), as shown below. Install in either a Gateway-cached or Gateway-Stored configuration. Refer to the directions on the following page for details about downloading the VM: <http://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStartedDownloadVM-common.html>

Setup and Activate Gateway Close



Choose How You Want to Run Your Gateway

Please select a configuration for running your gateway

- Gateway-Cached volumes:** Maintain local, low-latency access to your most recently accessed data while storing all your data in Amazon S3.
- Gateway-Virtual Tape Library:** Durably store data on Amazon Glacier while leveraging your existing tape-based processes.
- Gateway-Stored Volumes:** Schedule off-site backups to Amazon S3 for your on-premises data.

[Learn More About Gateway Configurations](#)

Continue

Setup and Activate Gateway Close



You will need a host in your datacenter to deploy the gateway virtual machine (VM). Pick a host that meets these [minimum requirements](#).

[← Back](#)

Continue

- Once the VM is installed you will need to size and configure local storage volumes for the AWS Storage Gateway to use as either cache and buffer volumes (Gateway-Cached configuration), or as buffer and storage volumes (Gateway-Stored configuration), as shown below. Refer to the directions on the following page for details: <http://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStartedPLDSMain-vm-common.html>

Setup and Activate Gateway Close



Step 2 of 2 (Allocate Disks to Cache Storage)

Again using your VMware vSphere client, allocate one or more local disks to your gateway VM to cache recently accessed data on-premises. These disks, called cache storage, are used to provide low-latency access to data you actively access.

Step-by-Step Instructions

Calculator

Enter the size of your upload buffer: GBs

For corporate file sharing workloads, enter 20% of your current on-premises file share storage.
For other workloads like backup, leave blank: GBs

Recommended storage for caching data: -- GBs

[< Back](#)

Continue

- Next you will need to activate the AWS Storage Gateway, as shown below. Refer to the directions on the following page for details: <http://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStartedActivateGateway-common.html>

Setup and Activate Gateway Close



Using your VMware vSphere client, right mouse-click on your deployed gateway VM and select Power On. Next, click on the Summary tab and retrieve the IP Address of your VM (it may take a couple of minutes for the IP Address to appear once you've powered on your VM). Type the IP Address into the box below.

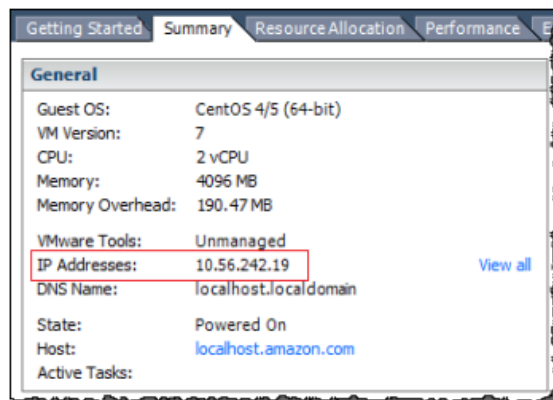
Clicking "Proceed to Activation" will redirect you to the activation page (your browser must be running on a machine with network connectivity to your local gateway host).

Step-by-Step Instructions

Enter IP Address Below:

[< Back](#)

Proceed to Activation



Screenshot showing the Summary tab.

4. Configure storage in Amazon S3 for the AWS Storage Gateway, as shown in below. If using a Gateway-Cached configuration, specify the volume size you want to create. If using a Gateway-Stored configuration, specify the local volume in the VM you created previously that will be assigned for use. Refer to the directions on the following page for details:
<http://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStartedCreateVolumes.html>

Configure Your Activated Gateway close

✓ ○

CONFIGURE LOCAL STORAGE CREATE VOLUME

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to this volume over an iSCSI interface. [Learn More.](#)

Capacity: TBs (Max: 32 TBs)

iSCSI Target Name: TBs
GBs

Based on Snapshot ID:

Host IP:

Port:

Configuring SteelFusion Core

1. Click on Add an iSCSI Portal on SteelFusion Core and enter the details of the on premise ASG Device.

iSCSI Portal Configuration

▼ Add an iSCSI Portal

Hostname or IP Address:

Port:

Authentication: None

2. Once the Portal is successfully added, SteelFusion Core will auto discover and add all the MPIO Portals (if available) and will add Targets.

10.5.42.203 3260 Connected

Status/Settings Offline LUNs

Portal Status: **Connected**

MPIO Portals: 10.5.42.203:3260 **Rescan for Portals**

Add Portal

Port:

Authentication:

Update iSCSI Portal

Targets:

+ Add a Target

Target	Status	Remove
iqn.1997-05.com.amazon:myvolume5	Connected	<input type="checkbox"/>

Status/Settings Offline LUNs

Target Status: **Connected**

Port:

Update Target

3. Launch the LUN Mapping Wizard under the Configure Tab.

DASHBOARD	CONFIGURE
STORAGE ARRAY	MANAGE
ISCSI, Initiators, MPIO	LUNs
CHAP Users	SteelFusion Edges
NETWORKING	POOL MANAGEMENT
Host Settings	Edit Pool
Management Interfaces	REST API Access
Data Interfaces	
WIZARDS	BACKUPS
Initial Setup	Snapshots
LUN Mapping	Handoff Host
Import	Snapshot Schedule
Export	Data Protection
REPLICATION	FAILOVER
Set Up / Monitor	Fallover Configuration

4. You can click Next for the first two screens.
5. On the Specify Portals page, click Select from known Portals and click Next.

LUN Mapping Wizard ?

iSCSI Initiator Name	<h3>Specify Portal</h3> <p>Select or add an iSCSI Portal.</p> <p><input checked="" type="radio"/> Select from known Portals</p> <p>Hostname or IP Address: <input type="text" value="10.5.164.51"/></p> <p>Port: <input type="text" value="3260"/></p> <p>Authentication: <input type="text" value="None"/></p> <p><input type="radio"/> Add new Portal</p>
Start	
Specify Portal	
Manage Targets	
Mount LUNs	
Specify SteelFusion Edge	
Map LUNs to Edge	
Summary	

6. Select the volume listed under the LUNs discovered section.

LUN Mapping Wizard ?

iSCSI Initiator Name	<h3>Mount LUNs</h3> <p>Select the LUNs to be configured from Discovered LUNs List. If the LUNs are already configured you may proceed.</p> <p>LUNs discovered from known Targets:</p> <div style="border: 1px solid #ccc; padding: 5px; min-height: 40px;">SS-VOL-27e84798-928c-419d-a456-c2494c1087eb</div> <p style="text-align: center;">Show Known LUNs ▼</p>
Start	
Specify Portal	
Manage Targets	
Mount LUNs	
Specify SteelFusion Edge	
Map LUNs to Edge	
Summary	

7. From the next page, select Add new SteelFusion Edge: and give it a unique edge identifier and click Next.

LUN Mapping Wizard ?

iSCSI Initiator Name	<h3>Specify SteelFusion Edge</h3> <p>Select or add a SteelFusion Edge.</p> <p><input type="radio"/> Select from known SteelFusion Edges:</p> <p><input checked="" type="radio"/> Add new SteelFusion Edge:</p> <p>SteelFusion Edge Identifier: <input type="text" value="SFBranch"/></p> <p>Blockstore Encryption: <input type="text" value="No encryption"/></p>
Start	
Specify Portal	
Manage Targets	
Mount LUNs	
Specify SteelFusion Edge	
Map LUNs to Edge	
Summary	

8. Select the LUN from the unmapped LUNs section and click Next.

LUN Mapping Wizard ?

iSCSI Initiator Name	<h3>Map LUNs to Edge</h3> <p>Select LUNs which are to be mapped to the SteelFusion Edge.'SFBranch' from Unmapped LUNs List. If the LUNs are already mapped you may proceed.</p> <p>Unmapped LUNs:</p> <div style="border: 1px solid #ccc; padding: 5px;"><div style="background-color: #f0f0f0; padding: 2px;">SS-VOL-27e84798-928c-419d-a456-c2494c1087eb</div></div> <p style="text-align: right;">Show Mapped LUNs ▼</p>
Start	
Specify Portal	
Manage Targets	
Mount LUNs	
Specify SteelFusion Edge	
Map LUNs to Edge	
Summary	

9. Click exit at the end of the wizard.

Now we have to establish the connection between edge and core.

1. On the SteelFusion Edge appliance, navigate to Storage > Storage Edge configuration and enter the IP of the Core along with the Edge Unique identifier which we had created earlier and click Add Core.

• pod3-3100a / SteelFusion Edge

Storage Edge Configuration Storage > Storage Edge

Connect to a SteelFusion Core

Hostname/IP:

SteelFusion Edge Identifier: *(Case Sensitive)*

Local Interface:

Add Core

- After the connection is established, you should see the LUN listed under the LUNs section.

SteelFusion Settings

Configured SteelFusion Core Hostname/IP:	10.5.59.101 (edit)	
Current Active Core:	oak-wa114	
SteelFusion Edge Identifier:	SFBranch	
Local Interfaces:	primary - 192.168.5.16 ⊗ aux - 10.33.195.34 ⊗ Add Interface	

Remove Core

SteelFusion Core Connection
 Connected to SteelFusion Core
[Show all Connections](#)

Blockstore Allocation	SteelFusion Core Connections	Target Details	Initiators	Initia
-----------------------	------------------------------	----------------	------------	--------

LUN Alias (Serial) ▾	Type ▾	Status ▾	LUN ID ▾	Size ▾
▶ alias-SS-VOL-27e84798-928c-419d-a456-c2494c1087eb (SS-VOL-27e84798-928c-419d-a456-c2494c1087eb)	iSCSI	Connected	1	50.00 GB

- At this point the LUN is mapped to SteelFusion Edge and now we will have to expose it to the hypervisor running on Edge.
- In order to get the IQN of the hypervisor on SteelFusion Edge appliance, navigate to the Hypervisor page and copy the IQN

Hypervisor

Status Management IP Address
Limited support - VSP managed 10.33.195.38

Software VMware vSphere
Version 6.0.0-1.17.3029758
Uptime 1 week, 6 days
IQN iqn.1998-01.com.vmware:568c4870-a57d-0f60-3397-000eb6b40ec4-06ff082e

5. Back on SteelFusion Core, navigate to SteelFusion Edges page and under Initiators, click Add an Initiator.

+ Add a SteelFusion Edge

SteelFusion Edge	Connection	Duration
▼ SFBranch	Connected	18h 34m 9s

Status Target Settings Initiators Initiator Groups LUNs Prepopulation

▼ Add an Initiator

Initiator Name: [Add Discovered Initiator](#)
Add to Initiator Group: [New Group](#)
Authentication:

Add Initiator

Once the initiator is added, you can log into the hypervisor via a vsphere client and create a datastore on the Storage you added and then create a virtual machine on it.

Chapter 4 Branch Disaster Recovery

In the unlikely event there is a significant disruption at a branch (such as a natural disaster), SteelFusion can be used to quickly recover branch data and services, either to another branch location, or within the data center itself. Since the data is stored within a LUN provided by the ASG device, SteelFusion core can redirect the LUN from the original branch to another branch with SteelFusion. Alternatively, you can directly mount the LUN to a VMware ESXi host in the data center.

Note that a branch outage will result in a crash consistent LUN state to exist for the LUN delivered by the ASG device, to SteelFusion, so the amount of data loss will depend on how much pending data SteelFusion was in the process of synchronizing to the AWS Storage Gateway LUN at the time the branch outage occurred.

Configure the LUN to a New SteelFusion Edge at a Different Branch

Reconfiguring a LUN for access by a new or different SteelFusion Edge is straightforward. Because a branch outage disconnects the ASG device, LUN from the original SteelFusion Edge, SteelFusion Core will need to perform the following steps in order to associate the LUN for use by a new or different SteelFusion Edge:

Login to the SteelFusion Core GUI, and select Configure > Manage > LUNs and identify the LUNs that are associated with the original SteelFusion Edge that is experiencing the outage, as shown in below figure. Note the LUN name that it has been given (also known as the LUN alias name).

+ Add a LUN

LUN	Type	Status	Size
▼ SFBranchLUN (SS-VOL-27e84798-928c-419d-a456-c2494c1087eb)	iSCSI	Degraded	50.00 GB

Establish the connection between your new Edge device and the SteelFusion Core. Please use the same Edge Identifier as your previous edge – the one which you have lost for reasons such as natural disaster or bad device.

• pod3-3100a / SteelFusion- Edge

Storage Edge Configuration

Storage > Storage Ed

Connect to a SteelFusion Core

Hostname/IP:

SteelFusion Edge Identifier: (Case Sensitive)

Local Interface:

Add Core

SSH to the SteelFusion Core management interface, and login. From the command prompt, issue the following commands to disassociate the previous edge with this serial.

```
en
conf t
edge modify id <Original_SteelFusion_Edge_Name> clear-serial
```

After this you should be able to see the LUN on the new edge.

SteelFusion Settings


Configured SteelFusion Core Hostname/IP: 10.5.59.101 ([edit](#))

Current Active Core: oak-wva114

SteelFusion Edge Identifier: SFBranch

Local Interfaces: aux - 10.33.195.34 [Add Interface](#)

Remove Core



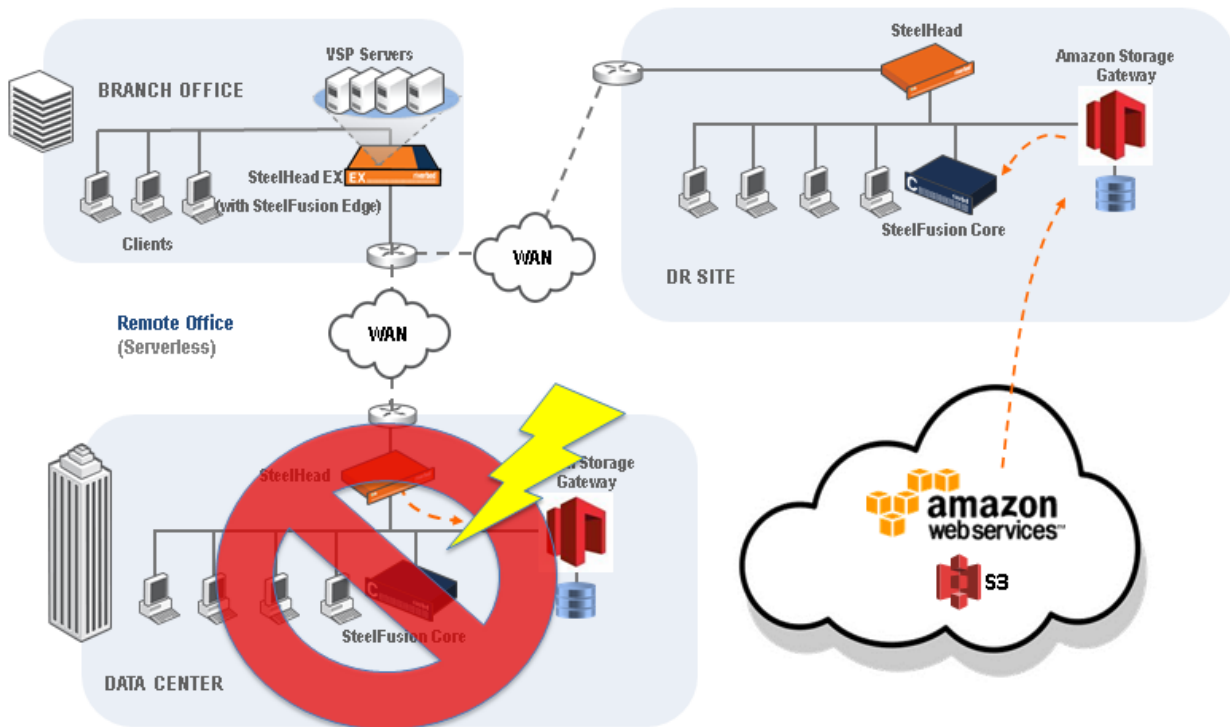
SteelFusion Core Connection Status
 Connected to SteelFusion Core
[Show all Connections](#)

Blockstore Allocation	SteelFusion Core Connections	Target Details	Initiators	Initiator Groups	LUNs
-----------------------	------------------------------	----------------	------------	------------------	------

LUN Alias (Serial) ▾	Type ▾	Status ▾	LUN ID ▾	Size ▾	Cached Data ▾
▶ SFBranchLUN (SS-VOL-27e84798-928c-419d-a456-c2494c1087eb)	iSCSI	Connected	1	50.00 GB	4.2 MB (0%)

Deploying a new instance of ASG:

Performing disaster recovery (DR) with an AWS Storage Gateway differs slightly from traditional storage disaster recovery since the data or data snapshots are stored in Amazon cloud storage, rather than on another SAN storage device at a DR site. The benefits of having cloud based snapshots available for recovery is that you do not have to maintain power, networking, and other infrastructure until such time that a DR recovery is needed. When a DR event occurs, you can deploy a new virtual instance of an AWS Storage Gateway and SteelFusion to perform recovery of required services for branch offices, as shown below This deployment can be done either within a new data center, or within a specialized facility that can use the [Amazon Direct Connect](#) features to connect to Amazon cloud storage services directly over high bandwidth networks.



DR Recovery with the AWS Storage Gateway and SteelFusion

1. Deploy a new AWS Storage Gateway at the DR site, similar to the steps outlined in Chapter 2 above. During step 4, in which you deploy a new volume for the AWS Storage Gateway, you will instead configure the new volume as a recovered volume from a previous snapshot. Refer to the directions on the following page for details: <http://docs.aws.amazon.com/storagegateway/latest/userguide/RestoringSnapshotVolume.html>

Create Storage Volume
close

Disk: Preserve existing data

iSCSI Target Name:

Based on Snapshot ID:

Size:

Host IP:

Port:

Configure Your Activated Gateway close

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to this volume over an iSCSI interface. [Learn More](#).

Capacity: TBs (Max: 32 TBs)

iSCSI Target Name: iqn.1997-05.com.amazon:

Based on Snapshot ID:

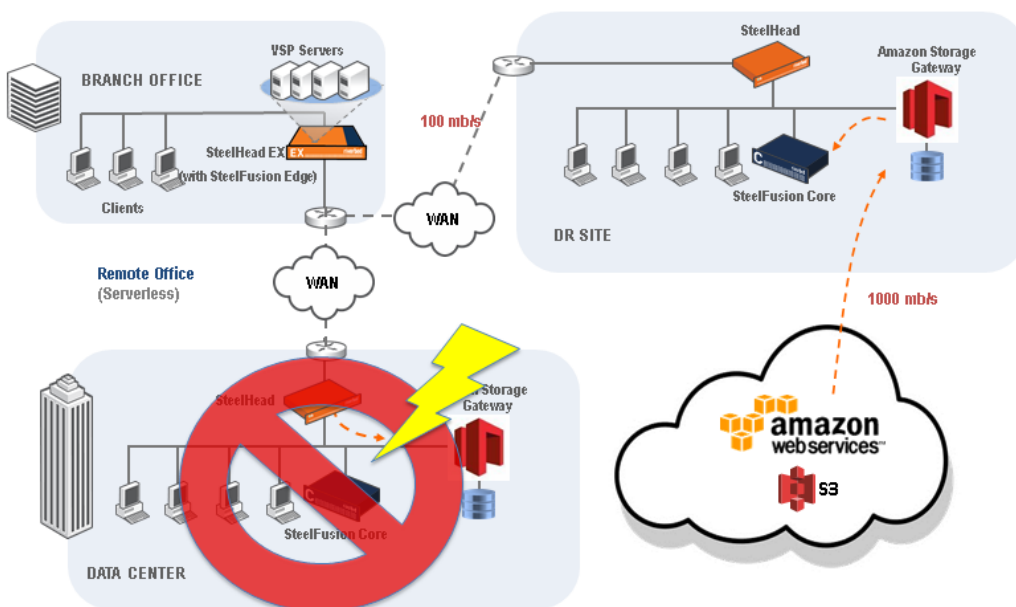
Host IP:

Port:

- After recovering the AWS Storage Gateway, you will need to deploy a new SteelFusion Core, as outlined in the previous chapter above. When adding the iSCSI target and LUN, you will point to the new iSCSI target and attach to the new LUN delivered by the DR AWS Storage Gateway, and provide initiator access to the correct host(s) at the branch which will need access (for example, SteelFusion Edge).

Note: If you want to recover data from the LUN directly in the data center, you may mount the LUN directly to a Windows or ESXi host as described in the sections Error! Reference source not found. and Error! Reference source not found.. SteelFusion is not required in this case and data can be directly recovered via Windows or VMware iSCSI connections to the LUN.

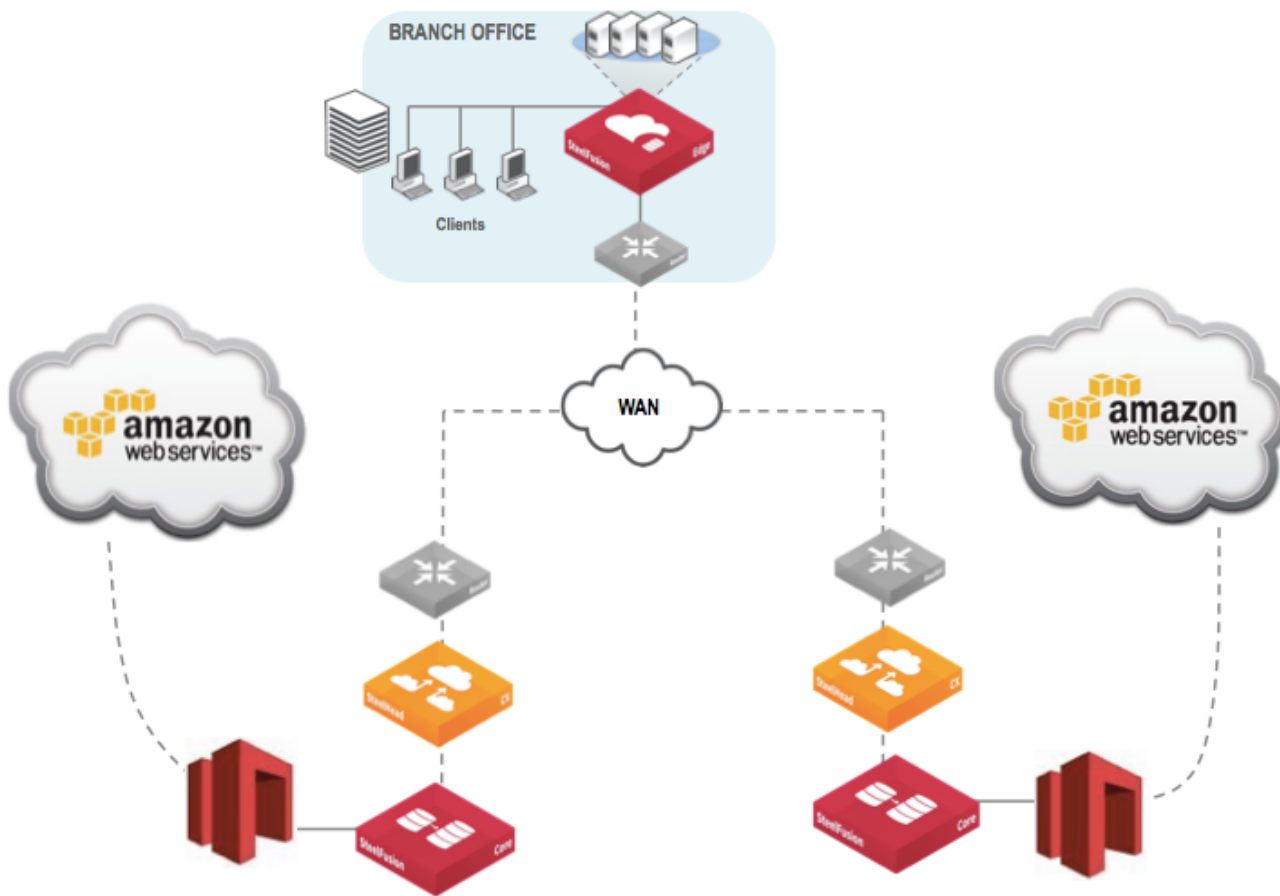
- When data access occurs (such as from an ESXi server at the branch launching a VM delivered via SteelFusion), SteelFusion will make requests for the data from the AWS Storage Gateway, and deliver that received data from the recovered snapshot to SteelFusion across the WAN to the branch SteelFusion Edge.
- The data recovery time will depend on your WAN speed between Amazon S3 and your SteelFusion Core, and the WAN speed between your SteelFusion Core and your SteelFusion edge. For example, if your SteelFusion Core is on a 1 gb/s Direct Connect link to Amazon S3, your recovery time will most likely depend on the WAN speed at which SteelFusion traffic can pass from the DR site to the branch. In the below example, a 100mb/s WAN to the branch could yield a Windows VM boot time of roughly 20-30 minutes. Alternatively, your DR procedures may dictate to recover the branch environment at the DR site, rather than at the branch, which could reduce the amount of time needed to initially recover services and bring them online for users.



Data center Disaster Recovery with FusionSync

A single data center is susceptible to large-scale failures (power loss, natural disasters, hardware failures) that can bring down your network infrastructure. To mitigate such scenarios, SteelFusion Replication enables you to connect branch offices to data centers across geographic boundaries and replicate data between them. FusionSync enables Cores in two data centers to remain in synchronization and enables the Storage Edges to switch to another Core in case of disaster and prevent data loss and downtime. This protects from data loss in case of a data center failure and from network downtime that can affect many branch offices at the same time.

For more information on how to set this up, please refer to SteelFusion Design Guide [here](#).



About Riverbed

Riverbed, at more than \$1 billion in annual revenue, is the leader in application performance infrastructure, delivering the most complete platform for the hybrid enterprise to ensure applications perform as expected, data is always available when needed, and performance issues can be proactively detected and resolved before impacting business performance. Riverbed enables hybrid enterprises to transform application performance into a competitive advantage by maximizing employee productivity and leveraging IT to create new forms of operational agility. Riverbed's 26,000+ customers include 97% of the *Fortune* 100 and 98% of the *Forbes* Global 100.

Learn more at riverbed.com

riverbed

©2016 Riverbed Technology. All rights reserved. Riverbed and any Riverbed product or service name or logo used here are trademarks of Riverbed Technology. All other trademarks used herein belong to their respective owners. The trademarks and logos displayed herein may not be used without the prior written consent of Riverbed Technology or their respective owners.