

# Talon FAST™ Reference Architecture Storage Consolidation into AWS





## Contents

Overview.....	3
Operating environment summary .....	4
Centralized data store in AWS.....	4
Talon FAST™ AWS VM(s).....	4
Talon FAST™ Edge VM(s).....	5
Network Connectivity .....	5
Configuration Guidelines .....	6
Hybrid Cloud .....	6
User experience.....	7
SUMMARY .....	8

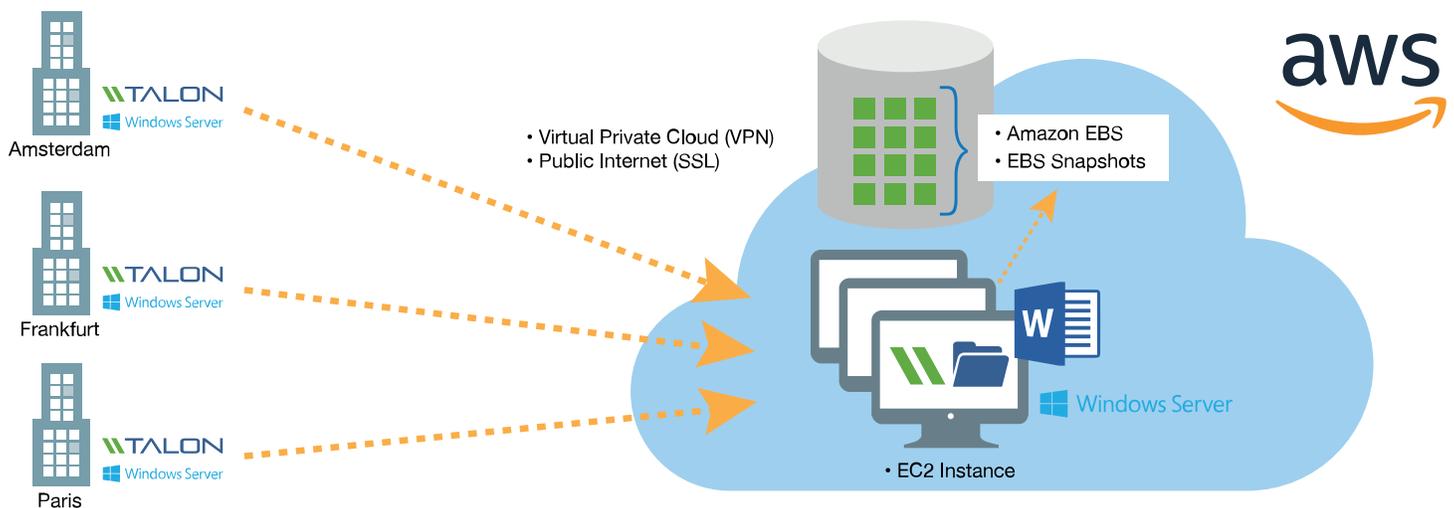
# Introduction

Talon FAST™ allows businesses to centralize data into a single footprint, leveraging a customer’s existing traditional datacenter or cloud storage infrastructure by consolidating distributed storage and IT assets. The software enables enterprises to transparently extend this centralized data to users globally, with real-time global file sharing and collaboration, without changing their workflow or experience.

The Talon FAST™ topologies include ‘hub-and-spoke’ (the most common), in which there are numerous edge instances addressing a single core instance; or mesh configurations, in which there are multiple cores (each front-ending a storage

pool) with edges addressing specific cores (and potentially other cores through their primary). This application note presents a reference architecture for leveraging the Talon FAST™ solution with a centralized storage repository housed in AWS.

Figure 1.





## Operating environment summary

The topology referenced herein is a 'hub and spoke' model, whereby the network of distributed offices/locations are all accessing one common set of data in the AWS cloud. The key points of the reference architecture herein are:

1. Centralized data store: EBS file shares attached to centralized ECS instance
2. Extension of the central data store to the distributed locations
  - a. Talon FAST™ Core running on ECS instance
    - i. Talon FAST™ Core virtual instance
    - ii. Talon FAST™ Licensing Management Server virtual instance
  - b. Talon FAST™ Edge VMs running in each location
3. Network configuration
  - a. Virtual Private Network connectivity
    - i. Point-to-site, or
    - ii. discreet Site-to-site
  - b. (optional) SSL connectivity
4. Integration with customer's Active Directory domain
5. DFS-Namespaces for the use of a global namespace (recommended)

## Centralized data store in AWS

The main repository for the unstructured data is a share (or number of shares) configured on the EC2 instance, with the storage provided by attached EBS resource. With this central repository, all locations of the enterprise, regardless of their locations in the world, can use this 'central file server instance' as if they were local. The result is a single storage footprint, versus a distributed storage architecture that requires local management, backup, security, footprint, etc. in each location. EBS storage is a storage service attached to EC2 compute instances offering the ability to create one or more file shares for the Talon Core instance, supporting the standard Server Message Block (SMB) Protocol. Both SMB 2.1 and SMB 3.0 are supported.

Further, applications running in other AWS virtual machines or cloud services can mount the file share in the cloud, just as a desktop application mounts a typical SMB share. Any number

of application components can then mount and access the file storage share simultaneously. With this approach, a logical SMB share can be created centrally and with Talon FAST™, that share can be accessed by users around the world. This allows customers to use the most scalable, cost-effective storage resource while using the native protocol to users and applications (SMB), extended through the Talon FAST™.

## Talon FAST™ AWS VM(s)

Running on the EC2 instance to which the EBS resource is attached is a Talon FAST™ core instance. The EC2 instance selected should conform to the Talon FAST™ Initial Sizing Guidelines provided by your Talon representative or Talon authorized partner; however, the minimum recommended VM for the Talon FAST™ core is a m4.xlarge instance. The Talon FAST™ core will extend the EBS-resident shares to the network of distributed locations using the Talon FAST™ distributed network file system. Through the use of the Talon FAST™ core, AD ACLs and NTFS file semantics are fully supported to ensure data coherency and integrity; in addition, global file locking is ensured regardless of where in the world users are physically located when they open and work on files.

In addition to the Talon FAST™ core, there will be a Talon FAST™ Licensing Management Server (LMS) instance running in the EC2 instance to ensure proper licensing, billing, and accounting. The VM size selected should take this into account, again conforming to the Talon FAST™ Initial Sizing Guidelines provided by your Talon representative or Talon authorized partner.

## Talon FAST™ Edge VM(s)

Each remote office will run a virtual instance (VM instance) of Talon FAST™ configured as an 'edge' instance; the edge VMs provide the critical performance-enhancing functions such as file caching, file-level differencing, and local service to users. The edge VMs can run on any available Windows Server 2012 R2 or Windows Server 2016 environment in the remote location, and will be configured with a local cache partition which will use algorithms to retain the most often used data for that location to ensure local-like performance for common tasks.

## Network Connectivity

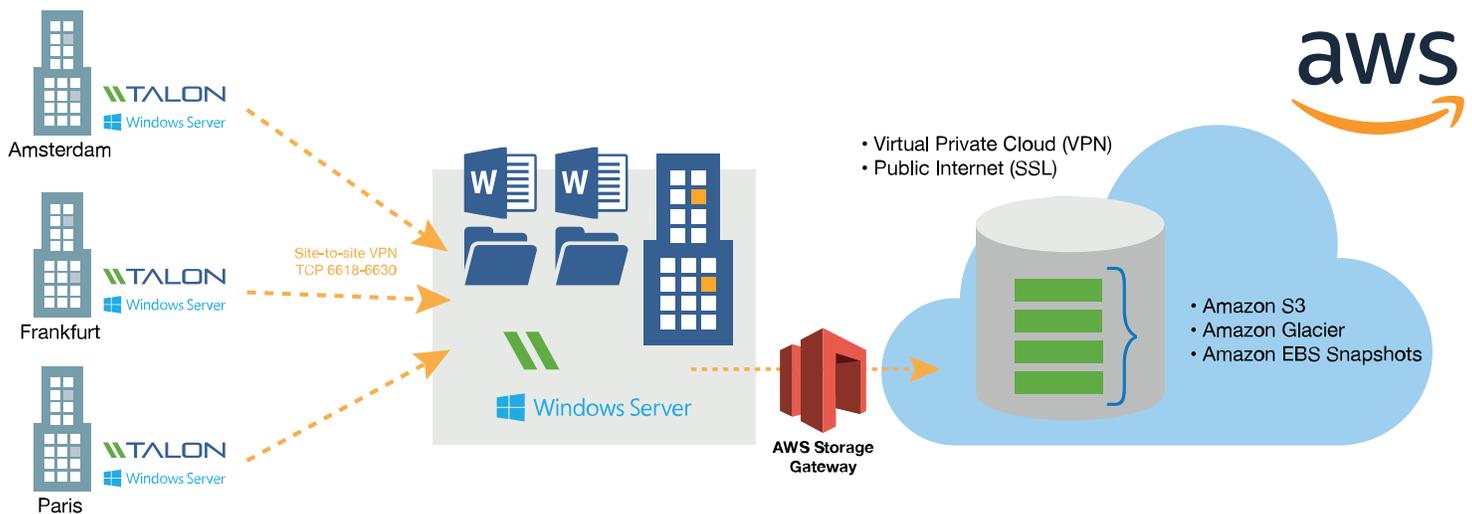
Connectivity is provided by a secure site-to-site or point-to-site Virtual Private Network (VPN) connection between each location needing access to the centralized data, and the AWS network. The VPN must be capable of carrying bidirectional traffic on TCP ports 6618 – 6622 between a Talon FAST™ core and the corresponding edges. A topological representation of the overall solution using a site-to-site VPN is presented in **Figure 2**.

## Configuration guidelines

Please review the Talon FAST™ Requirements and Talon FAST™ Initial Sizing Guidelines documents to ensure a successful configuration. In addition, for certain applications (such as CAD applications), it would be helpful to review any applicable 'best practices' documents available from Talon to ensure optimum configuration and performance parameters.

It is strongly recommended, in addition, that the customer employs DFS-Namespace to create a unified namespace to reduce complexity and simplify management of the distributed storage network.

Figure 2.





## Hybrid Cloud

Another topology is a hybrid approach using the AWS Storage Gateway. AWS Storage Gateway is a hybrid storage service that enables on-premise applications to seamlessly use AWS cloud storage. Customers can use the service for numerous use cases, including storage tiering, and migration.

In this use case, the main repository for the unstructured data is a share (or number of shares) configured on the on-premise Microsoft Windows Server 2012 R2 or Windows Server 2016 VM running the Talon FAST™ Core; this VM then connects to the on-premise AWS Storage Gateway (configured in Volume Gateway configuration mode); the Volume Gateway configuration connects to the on-premise Talon Core instance as a local disk using iSCSI, and the Talon Core VM creates a file system over the presented iSCSI volumes.

The AWS Storage Gateway then connects to AWS storage services such as Amazon S3, and data in these volumes is transferred into Amazon S3 cloud storage and accessed through the Volume Gateway as needed by the FAST™ Core. The topology will blend latency and scale by storing frequently-accessed data locally with “cooler” data in the cloud (with snapshots and clones for protection as well). The service includes a highly-optimized data transfer mechanism, with bandwidth management, automated network resilience, and efficient data transfer, along with a local cache for low-latency on-premises access to the most active data used by the Talon FAST™ fabric.

## User experience

By using the Talon FAST™ solution to extend AWS-resident file shares to globally distributed users, these users have access to these shares via either a UNC path or a DFS Namespace. The user experience for properly configured systems is analogous to the experience of having a local file server; i.e. users or applications can navigate to a directory structure, select shares/folders, and work with files. The complete range of file operations (open, save, copy, paste, etc.) are available to the user so that there is no change in workflow.

- When a user requests use of a file in the central repository, Active Directory will authenticate that user's access rights
- After a successful authentication, the file is opened centrally from the EBS share by the Talon FAST™ core, and a lock is applied (centrally) to that file
- If the file has not ever been used by that particular location, the file is served over the network using the proprietary streaming and compression technologies inherent in the Talon FAST™ solution to improve performance
  - If, however, the file has been used by (any user in) that location before, it is likely resident in cache; in this case the file will be served out of the local cache without incurring network transfer operations, thereby providing a high performance experience
  - If the file version in local cache is not the most up-to-date version that is in the AWS repository, any differences (and only the differences) will be sent to the local cache and merged with the cached version of the file upon open; this maximizes performance and minimizes network resource use
  - **note:** the file remains locked at the central repository, and was only served after authentication and lock were performed
- User operations continue as normal, and any updates/changes/writes will be cached locally
- Upon save/exit, any changes to the file will be 'differenced' back to the authoritative central copy
- Upon exit, after the saves are completed centrally the file is closed and the lock will be released and available to other users





## SUMMARY

This combination of AWS and Talon FAST™ allows enterprises to migrate their unstructured data 'to the cloud', taking advantage of the flexibility, availability, and economics of the cloud computing model for one of their largest use cases, that of unstructured data. The resulting effect(s) on the business are significant:

- a reduced footprint both at remote locations in the datacenter
- significant cost savings in the hardware and management components of providing file services to users
- increased flexibility and agility through enhanced global collaboration
- enhanced security and compliance by removing the risk of data loss/leak at remote locations through error, disaster, and intrusion

For more information, please visit [www.talonstorage.com](http://www.talonstorage.com)

