



UM Storage Services

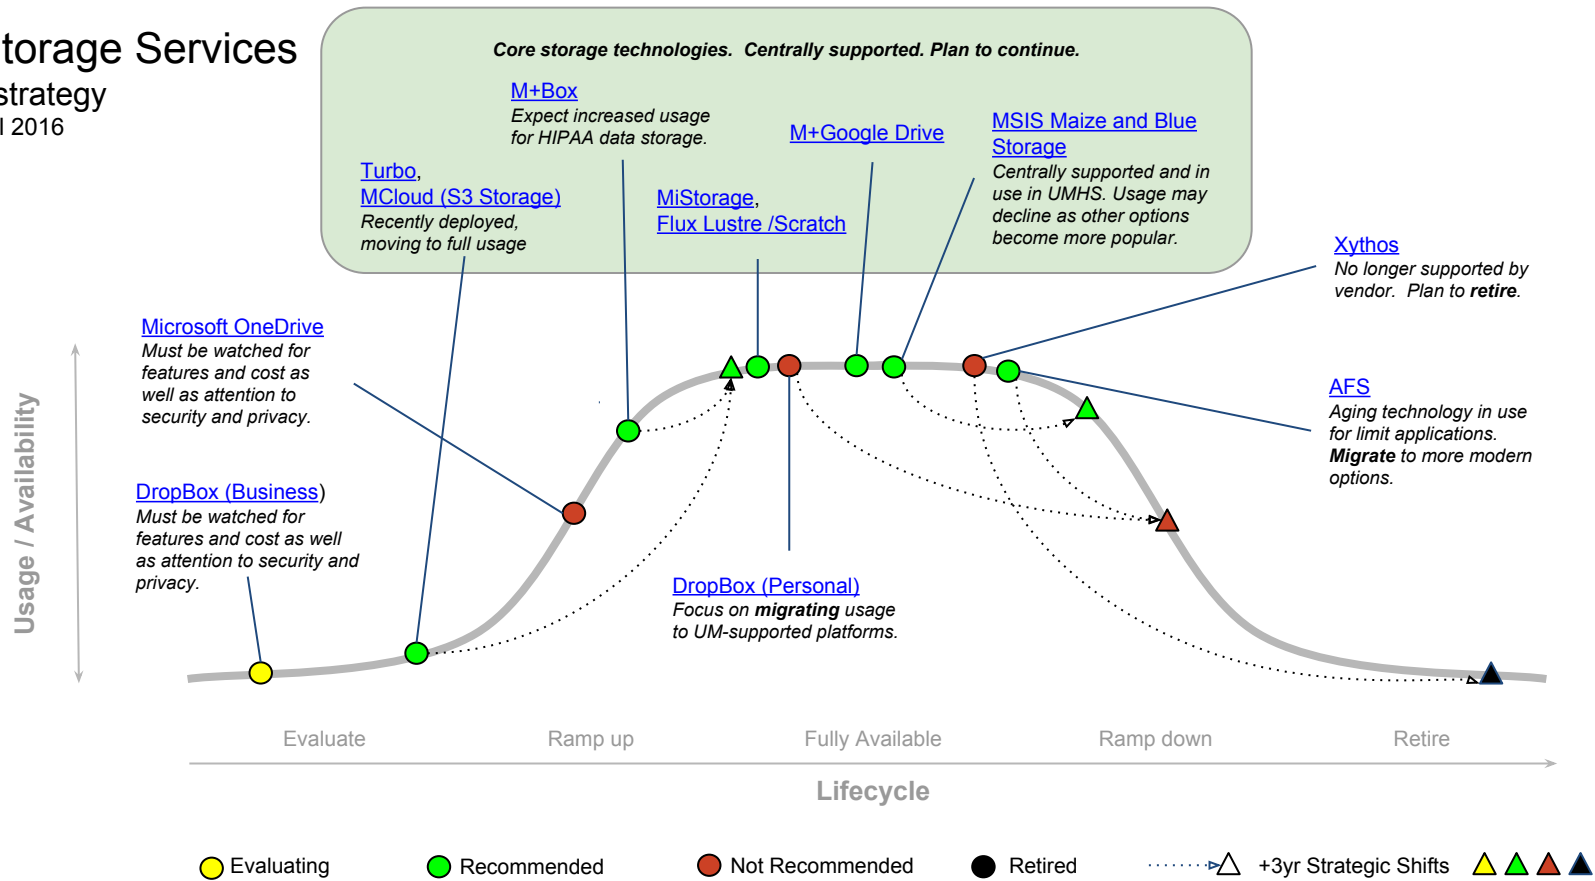
ITS, MSIS, ARC-TS

As of April 2016

UM Storage Services

3 year strategy

As of April 2016



Scope

This MESA covers storage services provided by central IT organizations for all of the University of Michigan - Ann Arbor campus. These organizations include [ITS](#), [MSIS](#), and [ARC-TS](#). The services cover the storage of data for administrative, academic, and research uses, including uses where patient health information is involved. The technologies used to implement these services include both those hosted on-premises and cloud-hosted ones. This rapidly evolving market must be closely watched as competing services continue to mature. Factors to watch include security, user experience, user acceptance, features, and cost. For assistance in choosing the appropriate storage options, please see the [Storage Solutions page](#).

Goals

- Provide cost-efficient, secure storage options that are either hosted on-premises or are cloud-hosted and protected by a contract between UM and the vendor
- Improve collaboration by promoting a small number of feature-rich storage platforms for general file storage and sharing. Provide the means for users to migrate off legacy storage to the common platforms.
- Provide sufficient breadth of storage options to meet the differing requirements for research, teaching, learning, and administrative work.



Initiatives

- Promote strategic storage and collaboration platforms
 - Encourage use of modern storage solutions such as MiStorage, Box, Google, etc. where applicable.
 - Provide education and incentives to move off of consumer cloud storage to university contracted solutions. In particular, provide education and tools to move content from Dropbox and OneDrive.
 - Transition CTools sites used for projects to appropriate cloud storage.
- Integrate Box and Google with Canvas.
- Look for alternatives to AFS to meet specific use cases.
- Retire Xythos
 - Stop using Xythos for any new file sharing and storage needs.
 - Migrate all users and applications using Xythos to the solution most appropriate from both a functional and data classification. For data classified as sensitive or ePHI the Box platform should be used because U-M has a BAA with Box.
 - Retire Xythos before having to extend the annual license (Current expires in 2017)



Google Drive

Google Drive is part of Google Apps for Education (GAFE), a suite of free productivity tools for classroom collaboration. “Google Drive is a file storage and synchronization service created by Google. It allows users to store files in the cloud, share files, and edit documents, spreadsheets, and presentations with collaborators. Google Drive encompasses Google Docs, Sheets, and Slides, an office suite that permits collaborative editing of documents, spreadsheets, presentations, drawings, forms, and more.”¹

Google Drive is **recommended** for cloud storage because it provides unlimited storage with real-time collaborative editing capabilities for a variety of document types.

U-M entered into a 10-year contractual agreement for GAFE in 2012. GAFE also includes email, calendaring, and other services. ITS provides GAFE at no cost to campus, the ITS Service Center is trained in Google Apps, and further support is provided by a knowledgeable Tier 3 team with access to Google support for escalation. The Tier 3 team has access to a console providing management capabilities for U-M users and access to logs. ITS provides training on request.

The 2015 Magic Quadrant for Enterprise File Synchronization and Sharing names Google as a Challenger. Google is focused on the consumer market and on the K-12 environment. As Google Drive continues to evolve, their attention to the needs of Higher Education must be watched.

GAFE was adopted by U-M after a process that engaged campus at all levels. GAFE was originally chosen for its email and calendaring platforms, with the goal of decreasing the large number of distributed email and calendaring systems on campus. As Google Drive improved in functionality and usability, its close integration with mail and calendar drove adoption.

Contact: collaborate@umich.edu

¹ http://en.wikipedia.org/wiki/Google_Drive

Box

“Box (formerly Box.net), based in Redwood City, California, is an online file sharing and content management service for businesses. The company uses a freemium business model to provide cloud storage and file hosting for personal accounts and businesses. Official clients and apps are available for Windows, Mac OSX, and several mobile platforms. Box was founded in 2005.”¹

Box is **recommended** for cloud storage because it offers unlimited cloud storage for and collaboration on any type of file. Because our contract with Box includes a Business Associates Agreement (BAA), Box is approved for HIPAA data and utilized by the medical campus to store content that may include patient data, which may not be stored in Google Drive.

The ITS Service Center is trained in Box and is further supported by a knowledgeable Tier 3 team with access to Box support for escalation. The Tier 3 team has access to a console providing management capabilities for U-M users and access to reports. ITS provides training on request.

The 2015 Magic Quadrant for Enterprise Files Synchronization and Sharing names Box as a Leader.

U-M entered into a 2-year contractual agreement for Box in 2012 (through Internet2 Net+) and renewed for 3 additional years (the longest renewal available) in 2014. The 2014 renewal added the BAA to protect HIPAA data. The contract was updated to unlimited storage in 2015. U-M participates on an Internet2/Box advisory board providing enhanced access and influence to the company direction. Box is focused on the the corporate environment and has Education and Healthcare & Life Sciences as target industries. As Box continues to evolve, their attention to the needs of Higher Education must be watched.

Contact: collaborate@umich.edu

¹ http://en.wikipedia.org/wiki/Box_%28company%29

UMHS Maize and Blue Storage

[MSIS Maize or MSIS Blue Storage](#) is an on-premises file storage system provided through UMMS File Storage Services. It is suitable for storing all data types, including ePHI. The system is scalable, secure and backed up for disaster recovery.

UMHS Maize and Blue Storage is **recommended** for staff and researchers within UMHS who have NAS storage needs for their data. It is available in both NFS and CIFS, and is also available in a flavor that can be made accessible to Campus IP ranges (Blue), or restricted solely within the UMHS firewall (Maize).

The services is currently supported by MSIS and MCIT. MCIT handles the infrastructure maintenance, and MSIS currently handles the customer facing interactions and service management.

Contact: msishelp@umich.edu

MiStorage

[MiStorage](#) provides scalable network-attached storage to U-M students, faculty, and staff, accessible to Mac, Linux, and Windows users. It is intended for administrative and departmental data.

MiStorage is **recommended** for storage because it is supported 7x24x365 by ITS. It is available to all members of the University of Michigan community. It is maintained and monitored as a production service, including security updates, on-call, and help desk support.

The MiStorage service offers a Silver and Gold tier. MiStorage Silver provides low-cost mass storage. MiStorage Gold is designed for applications that require faster performing disk than what is available at the Silver tier. MiStorage offers the option of CIFS or NFSv3 protocol. MiStorage with the CIFS protocol is approved for use with sensitive data. MiStorage using the NFSv3 protocol is not approved for use with sensitive data.

MiStorage replaces previously offered Mainstream and Value Storage FY16.

Contact: ifs-support@umich.edu

Turbo

[Turbo Research Storage](#) offers scalable storage to U-M researchers. The service provides NFSv3 and NFSv4 access. Turbo offers two security levels, one for some types of sensitive data, and one for non-sensitive data. Performance is intended to be sufficient for both IO operations and build file access. The service is designed to easily connect with Flux, the shared U-M computing cluster, as well as off-campus computing systems and collaborators. CIFS protocol and multi-protocol (simultaneous NFS and CIFS) volume capability will be added in FY16.

Turbo is **recommended** for faculty, scientists, clinicians, and students working on any campus (Ann Arbor, Flint, Dearborn, Medical) or off site with needs to store and retain research data.

User requests for support for Turbo are processed through [ARC-TS](#) by contacting hpc-support@umich.edu. Support is available during normal business hours (M-F, 8am-5pm). Details of Service Expectations may be found on the [Turbo Research Storage website](#).

Turbo is a new storage service targeting research data at the University of Michigan. Turbo can only be used for research data. Turbo along with Globus sharing should work well for sharing and hosting data for external collaborators and institutes.

Contact: hpc-support@umich.edu

M Cloud (S3)

M Cloud is the U-M offering of public cloud services to the University of Michigan (U-M) community. The service enables the U-M community to more easily consume public cloud services by simplifying billing and providing account management and other support services.

[MCloud](#) is a **recommended** service because it supports the IT strategic principle of “Cloud first”. Some of the Services under this umbrella also provide features that cannot be achieved by on premises compute services.

MCloud currently includes access to Amazon Web Services (AWS) under a University of Michigan enterprise agreement. The agreement allows the U-M community to opt-in to use AWS’s cloud-based IT infrastructure services under a U-M Master Account with greater protections than the standard AWS click-through license terms.

Amazon Simple Storage Service ([Amazon S3](#)), provides developers and IT teams with secure, durable, highly-scalable cloud storage. Amazon S3 is easy to use object storage, with a simple web service interface to store and retrieve any amount of data from anywhere on the web. With Amazon S3, you pay only for the storage you actually use.

Contact: Service Manager: Mark Personett, Service Owner: Chris Wood

Flux Lustre /Scratch Storage

[Flux Lustre/Scratch storage](#) is storage included in the Flux Service.

Flux Lustre/Scratch Storage is **recommended** for in situ storage where performance is important for computation. It is connected directly to the Flux cluster through a high speed, low latency infiniband network. There are no backups to the Lustre storage, and it is periodically purged of data not accessed for over 90 days, so it is not designed for any sort of long term storage. Data here should be copied from this service to another storage service when analysis is complete

This storage is supported by ARC admins as a part of the Flux cluster. There are infrastructure upgrades planned both in amount of storage as the Flux cluster grows, and in as newer interconnects reach mainstream acceptance (100 GBit Infiniband).

Contact: hpc-support@umich.edu

AFS

“The Andrew File System ([AFS](#)) is a distributed file system which uses a set of trusted servers to present a homogeneous, location-transparent file name space to all the client workstations. It was developed by Carnegie Mellon University as part of the Andrew Project. It is named after Andrew Carnegie and Andrew Mellon. Its primary use is in distributed computing.”¹

AFS is **recommended** for storage because it is supported by ITS. It is available to all members of the University of Michigan community, maintained and monitored as a production service, including security updates, on-call, and help desk support.

The AFS service lifecycle is under review for retirement as part of the Project Horizon project. AFS has very limited support for file sharing and collaboration. Web access is limited to [MFILE](#). Users are advised to use more modern storage options such as Box, Google, or MiServer. AFS is no longer automatically provisioned as part of the standard computing package. Faculty, staff, and students that require AFS can still obtain the service using the self provisioning tool.

Contact: ifs-support@umich.edu

¹ http://en.wikipedia.org/wiki/Andrew_File_System

Xythos

“Xythos Software, a Blackboard company, has long focused on providing standards-based technology, using WebDAV to provide flexible access to content. The Xythos Enterprise Document Management Suite (EDMS) is a Java-based alternative to SharePoint with proven ability to scale and a good track record of adoption in the higher education sector. Xythos EDMS has been expanded to include Web 2.0 and offline capabilities.”^[1] It has integrated wiki support and Really Simple Syndication (RSS) feeds. Xythos supports external collaboration by a feature it calls a "ticket." Essentially, a ticket lets documents be shared with external users (non-named users) via a secure URL.

Xythos is **not recommended** for any uses either by individuals or by applications which rely on its file storage capabilities.

Minimal support is currently provided by the vendor. Recently, MSIS contacted the vendor and the customer service representative was not even aware of the product.

Since the acquisition by Blackboard, Gartner has seen the majority of client interest coming from the higher education area. This is consistent with Blackboard's overall business focus. Gartner has not seen the same level of investment or broader strategy to address buyers in other markets. Gartner has removed Xythos Software from its research, because they have seen no client interest in its content management technology, nor have they seen any promotional efforts, since 2010.

Xythos is used extensively by the Medical School including individual users and many software applications that use it for file storage and file delivery. MSIS understands the user base and will need to plan to migrate from this platform.

[1] Gartner Magic Quadrant for Enterprise Content Management 2008, G00160668

Contact: David Glaser (dsglaser@umich.edu)

Microsoft OneDrive

“OneDrive (previously SkyDrive, Windows Live SkyDrive and Windows Live Folders) is a file hosting service that allows users to upload and sync files to a cloud storage and then access them from a web browser or their local device. It is part of the suite of online services formerly known as Windows Live and allows users to keep the files private, share them with contacts, or make the files public. Publicly shared files do not require a Microsoft account to access.”¹

Microsoft OneDrive is **not recommended** for cloud storage because it is a consumer based service with no U-M contract. With no contract with Microsoft for OneDrive, we have no ability to manage, and no protection for files stored in OneDrive accounts, even though they use the @umich.edu address.

ITS does not support Microsoft OneDrive.

The 2015 Magic Quadrant for Enterprise File Synchronization and Sharing names OneDrive as a Challenger.

In October, 2014, Microsoft began to include OneDrive in the Office 365, making it available to campus users through our Microsoft campus agreement. While this agreement covers use of Office 365 software applications, use of OneDrive storage requires that the user accept an End User License Agreement (EULA) and is not covered by the campus agreement. For that reason, no sensitive data is allowed and it is not recommended for any university data. However, the Office 365 software suite makes it natural to store files there. Since the OneDrive space is available because of our software agreement, users may believe that their data is also covered by university contract. The widespread use of Microsoft Office software paired with the ease of storing data in OneDrive makes OneDrive a solution to watch. Users wishing to use the Office 365 applications may store files in Box, which provides an Office 365 integration.

Contact: collaborate@umich.edu

¹ <http://en.wikipedia.org/wiki/OneDrive>

Dropbox (personal)

“Dropbox is a [file hosting service](#) operated by Dropbox, Inc., headquartered in [San Francisco, California](#), that offers [cloud storage](#), [file synchronization](#), [personal cloud](#), and [client](#) software. Dropbox allows users to create a special folder on their computers, which Dropbox then synchronizes so that it appears to be the same folder (with the same contents) regardless of which computer is used to view it. Files placed in this folder are also accessible via the Dropbox website and mobile apps. Dropbox provides client software for [Microsoft Windows](#), [Mac OS X](#), [Linux](#), [Android](#), [iOS](#), [BlackBerry OS](#), [Windows Phone](#) and [web browsers](#), as well as unofficial ports to [Symbian](#) and [MeeGo](#).”¹

Dropbox (personal) is **not recommended** for cloud storage because it is a consumer based service with no U-M contract. Dropbox has had a number of incidents showing lax privacy policies and attention to security².

U-M has not entered into a contractual agreement with Dropbox, therefore we have no ability to manage Dropbox accounts, even though they use the @umich.edu address. U-M users of Dropbox do so through a personal EULA. ITS does not support Dropbox.

The 2015 Magic Quadrant for Enterprise File Synchronization and Sharing names Dropbox as a Challenger.

Dropbox was founded in 2007. While Dropbox has never been an official U-M service, it has been used extensively by U-M users, including for university business. Users find it easy to use and are reluctant to change to official U-M services. A report sent to the CIO in February, 2015 from Dropbox stated that there are 42,268 Dropbox accounts registered under umich.edu.

Contact: collaborate@umich.edu

¹ http://en.wikipedia.org/wiki/Dropbox_%28service%29

² http://en.wikipedia.org/wiki/Dropbox_%28service%29#Privacy_concerns



Dropbox (business)

“Dropbox is a [file hosting service](#) operated by Dropbox, Inc., headquartered in [San Francisco, California](#), that offers [cloud storage](#), [file synchronization](#), [personal cloud](#), and [client software](#). Dropbox allows users to create a special folder on their computers, which Dropbox then synchronizes so that it appears to be the same folder (with the same contents) regardless of which computer is used to view it. Files placed in this folder are also accessible via the Dropbox website and mobile apps. Dropbox provides client software for [Microsoft Windows](#), [Mac OS X](#), [Linux](#), [Android](#), [iOS](#), [BlackBerry OS](#), [Windows Phone](#) and [web browsers](#), as well as unofficial ports to [Symbian](#) and [MeeGo](#).”¹

Dropbox for Business (which includes Dropbox for Education) is **being evaluated** for cloud storage because To date, U-M has not signed a contract for this service. Dropbox has had a number of incidents showing lax privacy policies and attention to security² However, Dropbox claims to have improved privacy and security in their Business offering. The popularity of Dropbox personal with faculty and staff would lend to easy adoption.

Dropbox is not supported by ITS. Users storing university data on Dropbox must rely on vendor support.

The 2015 Magic Quadrant for Enterprise File Synchronization and Sharing names Dropbox as a Challenger.

Dropbox was founded in 2007. When U-M and other universities were looking for an enterprise File Sync and Share service, Dropbox was approached and was not interested. As a result, U-M and other universities went on to work with Box in developing the Internet2 Box service. Given the popularity of the personal service, this must be watched for its features and cost as well as its attention to security and privacy.

Contact: collaborate@umich.edu

¹ http://en.wikipedia.org/wiki/Dropbox_%28service%29

² http://en.wikipedia.org/wiki/Dropbox_%28service%29#Privacy_concerns



Appendix

Tolerate, Invest, Migrate, Eliminate

Cloud Storage

