Windows File Server Migrations to Microsoft Azure Cloud

Brien M. Posey

Executive Summary

Migrating an organization's file data to Azure cloud can potentially address a number of pain points that are commonly associated with storing and managing large quantities of unstructured data. However, enterprise class organizations have traditionally shied away from performing large scale file server migrations to the cloud, citing concerns over logistics and latency. In this white paper we review the reasons why an organization should consider file migration to Azure, and take an in-depth look at three solutions—Azure Files, Azure NetApp Files, and Nasuni with Azure Blob storage—to understand their most appropriate use cases.

Trends Driving Migration to Azure

Today, trends ranging from Windows Server 2008 end-of-life to the need for remote file access are motivating enterprise class organizations to migrate their file data to Microsoft Azure cloud.

Windows Server 2008 EOL

Windows Server 2008 and 2008 R2 reached their collective end-of-life date in January 2020. Although these versions of Windows Server continue to function, Microsoft no longer provides regular security updates. This leaves organizations who continue to operate Windows Server 2008 and 2008 R2 servers vulnerable to attack.

Microsoft understands that some customers have situations requiring them to continue using these aging operating systems and has provided two options for those who need ongoing security updates.

The first option is to purchase extended security updates. While this option provides organizations with the security updates that they need, many consider this option to be cost prohibitive. The extended security update pricing is 75% (annually) of the Enterprise Agreement or Server and Cloud Enrollment license cost for the latest Windows Server version. Furthermore, organizations will be required to pay for the first full year up front. Those who decide to purchase the service mid-year will be required to pay for the full year.

The other, less costly option is to migrate the server to Microsoft Azure. Microsoft has committed to providing organizations that move their Windows Server 2008 and 2008 R2 workloads to Azure with three years of extended support at no additional charge (beyond what it costs to host the workload in Azure). Organizations who decide to rehost their Windows Server 2008 and 2008 R2 workloads in Azure also have the option of upgrading to the current version of Windows Server whenever they are ready to do so.

Microsoft understands that some customers have situations requiring them to continue using these aging operating systems and has provided two options for those who need ongoing security updates.

Migration Drivers for New Operating Systems

Although the extended support costs are driving many organizations to move their Windows Server 2008 and 2008 R2 workloads to Azure, there are a number of additional drivers that are causing organizations to consider an Azure migration even for workloads running on newer Windows operating systems.

Remote Worker Support

In 2020, before the onset of the COVID-19 pandemic, approximately 3.4% (7 million people) of the US population was working remotely, with approximately 43% working remotely at least some of the time. Since the beginning of the pandemic however, 88% of organizations have either encouraged or required their employees to work from home. According to a Gallup poll, 60% of all employees are now working from home. While this number might seem low, it is worth remembering that the poll spanned a variety of industries, including those supporting essential services whose employees did not have the luxury of working remotely. Regardless, the statistics point to a massive increase in the number of people who are working remotely. Interestingly, Gallop also determined that 59% of US workers want to continue working remotely after the pandemic is over. As such, it seems unlikely that everyone will go back to working on site as so many people did before COVID-19. The remote work trend is clearly here to stay.

It seems unlikely that everyone will go back to working on site as so many people did before COVID-19. The remote work trend is clearly here to stay.

Because the mandatory quarantines happened so quickly, enterprise IT had to scramble to put systems in place that could handle all of the users who were suddenly working remotely. In many cases, this meant migrating unstructured data and critical business workloads to Microsoft Azure or other public clouds over the course of several days or a few weeks.

Avoiding the Pitfalls of Consumer Oriented Cloud Services

Another trend that is driving the transition to Azure-based file services is the need to prevent users from using consumer-grade services such as Dropbox to store or transfer corporate data. Although such services are undoubtedly popular, their use presents a number of logistical challenges. For instance, data residing in Dropbox or a similar service is siloed and is outside the IT department's direct control. The IT department lacks the ability to back up the data or to prevent users from sharing sensitive data with people outside of the organization.

The use of consumer-grade services for the storing and transferring of data may also undermine an organization's compliance initiatives. When a user resorts to using a consumer-grade service to store or transfer corporate data, they are effectively undermining the logging, access control, and other security mechanisms that the organization has put into place. If a compliance audit reveals that such services are being used to store or transfer regulated data, then the organization may be punished with substantial fines or other legal actions.

Migrating file data to Azure cloud can help to end the practice of using consumer-grade services to store or transfer corporate data. File data that is hosted in Azure cloud is accessible from anywhere. Making data universally accessible may remove the need for users to resort to taking matters into their own hands by leveraging consumer-grade services.

Data Growth

One of the biggest reasons why an organization may opt to migrate its file data to Azure cloud is because Azure is an ideal solution to the problem of exponential data growth. While it is true that an organization can accommodate data growth by adding hardware to its own datacenter, this approach requires the organization to purchase storage hardware before it is actually needed. This means that an organization that decides to store its file data on-premises will have to make accurate data growth projections in order to avoid over or under-purchasing storage hardware.

If an organization underestimates its data growth, it will be left scrambling to purchase even more hardware to accommodate the unanticipated data growth. These unplanned hardware purchases will inevitably force the organization to shift its priorities, since a significant portion of the IT budget is being spent on an unplanned expense. Additionally, there is a ripple effect since adding storage capacity to accommodate data growth also requires organizations to invest in increased backup and DR capacity.

Capacity planning is of course, a normal part of the storage management process. Even so, administrators are often surprised by the sheer volume of data that exists within the organization. This data might not all exist in one place, but if all of the organization's data is examined collectively, its total size is likely larger than anyone realizes. Bringing all of this data together in one place (Azure cloud) can help an organization to optimize its cost management strategy, because having all of the data in one place makes it easy to see what the organization is actually paying for.

Microsoft Azure uses pay-as-you-go pricing as opposed to requiring an organization to make a large upfront investment in storage hardware. As such, an organization will never have to worry about purchasing storage that ultimately goes unused. Likewise, underestimating the organization's data growth also ceases to be a problem because Azure storage is available on-demand without the preplanning that is required when storing data on-premises.

Need for Cost Control

It's worth considering that public cloud providers have been sharply criticized in recent years for no longer being the inexpensive alternative to on-premises operations that they once were. As such, there is no guarantee that simply migrating data to the cloud is going to reduce an organization's cost. However, migrating file data to Azure has the potential to greatly reduce both operational costs and administrative effort, especially for large data sets and/or multiple locations.

Cloud migration allows organizations to transition from a CapEx model to an OpEx model. OpEx cost structures give organizations flexibility that cannot be achieved with a CapEx pricing structure. This is especially important today, with so many organizations seeing dramatic swings with regard to the size of their workforce and to the locations within which employees are working.

Transitioning to an OpEx model is not the only cost control benefit to transitioning to Azure. Large organizations with multiple locations will likely benefit from cost savings associated with infrastructure consolidation. This consolidation also reduces the complexity of management and may free up IT staff to focus less on repetitive management tasks and more on IT projects that play a direct role in helping the organization to achieve its business objectives.

Bringing all of this data together in one place (Azure cloud) can help an organization to optimize its cost management strategy.

Removal of Data Silos

One more reason why so many organizations have begun migrating their file data to Azure is to remove data silos. Organic growth within the enterprise almost always leads to the creation of data silos. It is relatively common, for example, for individual teams and working groups to deploy data storage solutions that exist separately from the organization's primary file storage. The data within these repositories is siloed from the rest of the organization's data, making it difficult to back up, or to apply retention policies and other safeguards.

Migrating an organization's file data to Azure can help to eliminate these data silos, making it far easier for enterprise IT to manage the data, and for end-users to access it.

Three Microsoft Solutions

Those who are considering migrating their Windows File Servers to Azure cloud have at least three different options for doing so. It is important to note that these three methods differ in terms of both architecture, scalability, and cost. Perhaps more importantly, there are several different types of storage available in the Azure cloud. The type of storage used when storing file data in Azure will dramatically affect the price. As such, it is important to carefully consider the type of storage that each solution uses. For example, Azure Files for instance is billed at a higher rate than Azure Blob (object) storage.

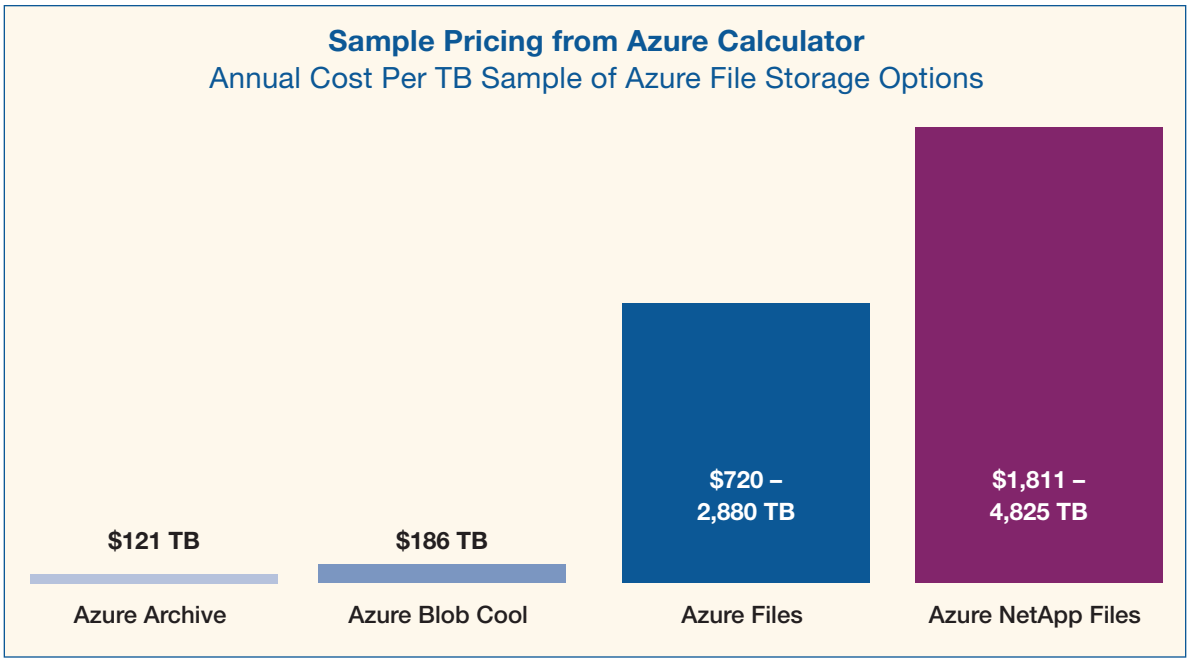

Not all Azure storage is priced the same; here are some pricing examples from time of publication relevant to Azure for file storage. For updates, visit the Azure Pricing Calculator. Annual cost per TB, based on 100 TB per month.

Azure Files

Azure Files is the first solution we'll look at for hosting file data in Azure cloud. Microsoft provides two different options for organizations who wish to use Azure Files.

The first option is to map a network drive directly to an Azure-based SMB share. The advantage to this approach is that it is simple to implement. Network endpoints connect to an Azure-based file share in exactly the same way that they would connect to an on-premises file share.

Of course, the disadvantage to using this approach is that file read and write performance is less than optimal. File data must be accessed across a WAN link. In an enterprise environment, there are commonly numerous users accessing file data at any given moment. The sheer volume of traffic stemming from users directly accessing Azure-based SMB shares can make this approach completely impractical.

This is where the second option comes in – Azure File Sync. This add-on to Azure Files allows organizations to create storage tiers in which hot data is kept on-premises, while cooler data resides in Azure Files. Rather than connecting directly to an Azure-based file share, network endpoints connect to a Windows Server that is running Azure File Sync. This server acts as a proxy to the Azure Files share. Because the hot data is cached within the on-premises server, the volume of file access-related WAN traffic is reduced, and users see vastly improved performance.

Although Azure File Sync is effective at improving performance over direct SMB access to Azure Files in the cloud, its use undermines some of the benefits that are driving organizations to go to the cloud in the first place.

One such lost benefit for example, is that of reducing costs by eliminating on-premises infrastructure. Azure File Sync requires an on-premises Windows Server to run the synchronization service and host the caching data. Like any other Windows Server, the synchronization server must be properly licensed, and kept up to date with the latest patches. It is also worth noting that because the synchronization server runs a Windows Server operating system, the server's OS will eventually reach its endoflife data and have to be upgraded.

Although Azure File Sync is effective at improving performance over direct SMB access to Azure Files in the cloud, its use undermines some of the benefits that are driving organizations to go to the cloud in the first place.

Another consideration that affects cost is that Azure Files require organizations to pay extra for snapshots (maximum of 255) and for the added-cost Azure Backup service that uses these snapshots for recoveries.

Azure NetApp Files (ANF)

A second option for moving file server data to Azure cloud is to use Azure NetApp Files (https://azure.microsoft.com/en-us/services/netapp/). Azure NetApp Files is a premium solution that is designed specifically for use with high-performance database workloads.

Although Azure NetApp Files has its place, it is not cost-effective for general SMB/CIFS workloads, regardless of tier. It is also important to note that Azure NetApp Files is not designed for multi-site file synchronization and sharing.

Azure NetApp Files is essentially a NetApp FAS or AFF system hosted in Azure and offered as a monthly subscription service. This means that the interface will be very familiar to those who use NetApp on-premises on a regular basis. It is worth noting however, that Azure NetApp Files is not as scalable as NetApp's on-premises solution in terms of overall capacity. Volumes can only accommodate 368 TB (slightly more than 1/3 PB) of data.

Azure NetApp Files also requires a separate backup service for long-term data protection, so that is another cost that must be factored in.

If an organization requires on-premises file syncing capabilities similar to the Azure File Sync option with Azure Files, they will need to add NetApp Global File Cache (formerly known as Talon). This increases both the operational overhead and cost, and like Azure Files Sync, requires a dedicated on-site Windows File Server.

It is worth noting however, that Azure NetApp Files is not as scalable as NetApp's on-premises solution in terms of overall capacity.

Azure NetApp Files also requires a separate backup service for long-term data protection, so that is another cost that must be factored in.

Nasuni and Azure Blob Storage

The third option for organizations wishing to migrate their file data to Azure is Nasuni with Azure Blob object storage. Nasuni is a Microsoft Gold ISV Partner, which means their solution is a Microsoft-approved migration path to Azure. More importantly however, Nasuni's solution helps organizations to overcome the challenges of storing file data in Azure cloud, without sacrificing the benefits that initially attracted the organizations to Azure. See Figure 1: Features Comparison Chart.

Nasuni works by creating a global file system that lives in cloud object storage such as Azure Blob. This becomes the authoritative source of all file data and metadata. Nasuni then caches copies of the metadata and active (hot) files on virtual appliances that can be deployed on-premises or in the cloud, giving users fast access to file shares wherever they are needed. Nasuni's approach creates several distinct advantages over competing solutions.

First is that Nasuni is designed to use object storage like Azure Blob as its back-end repository. Note that organizations will still get fast file access even with object storage as Nasuni's back-end because almost all file reads and writes are serviced by the SSDs behind each caching VM.

Next is that Nasuni is designed to work in both hybrid cloud or in pure Azure environments. If an organization wants on-premises access to file shares, they deploy a Nasuni caching VM in the office. If an organization wants files shares to be hosted in the cloud, they deploy a Nasuni caching VM in Azure. Additionally, Nasuni makes it easy to transition between the two. If an organization starts with a hybrid cloud and then later decides to host all file shares in Azure, Nasuni can easily adapt by turning off the VMs on-premises and turning on more VMs in the cloud.

The fact that Nasuni supports hybrid cloud and cloud-only deployments may be especially compelling for organizations that use both. Some organizations for instance rely primarily on hybrid cloud environments, but may have individual departments that are operating purely within Azure.

Regardless of how it is deployed, Nasuni costs much less per terabyte than the other two solutions. One reason for this is that Nasuni uses Azure Blob storage, which is far less expensive than the storage used by Azure Files or Azure NetApp Files, as shown in the example above. Another reason is that Nasuni has backup and disaster recovery built-in, so organizations don't have to pay extra for these the way they would with Azure Files and Azure NetApp Files. See Figure 2: Cost Comparison Chart.

The Elimination of Data Silos

One of the primary advantages to Nasuni's approach is that it eliminates data silos. All of the organization's file data is brought together under the Nasuni Unified File System (UniFS®). UniFS is able to span all on-premises and cloud locations, providing a consolidated view and a unified namespace that is readily accessible, regardless of where users are located.

Nasuni uses Azure Blob storage, which is far less expensive than the storage used by Azure Files or Azure NetApp Files.

Multi-Site File Sharing

With so much of the world's workforce continuing to operate remotely, it has become critically important to make sure that remote users are able to securely access file data, just as they would if they were working on site. As simple as this requirement may sound, organizations are repeatedly finding that there are challenges associated with giving users a consistent experience.

Consider for a moment what happens when a large enterprise replicates file data across multiple sites. Users within a given location typically all connect to the same site, thereby ensuring that everyone in that location is looking at the same data. When users connect from remote locations however, there is no guarantee that users will be accessing data within the same site that they would be accessing if they were in the office. As such, two users who normally work in close proximity to one another might be accessing data located in two different sites. This could potentially result in the two users having an inconsistent view of the data. If one user were to make an update, a user who is accessing a replica of the data located in another site won't see the change until the next replication cycle occurs.

When used with public cloud storage like Azure Blob, Nasuni uses high-speed internet links to securely propagate just the changes to active files from Nasuni Edge Appliances to cloud storage, and then to other edge appliances. Nasuni Global Volume Manager® aligns all changes from all locations by sequencing the file deltas in cloud storage, creating an immutable version history of all files that can be retrieved at any time. With every appliance kept consistently in sync, users globally will think they're working on one big, fast local file server.

Cloud-Based Data Protection

As users create and modify file data, Nasuni stores files as immutable files in WORM format within Azure Blob. Nasuni allows for nearly unlimited changes to each file to be retained, in fact it can snapshot versions as frequently as every 5 minutes. If an organization is impacted by a ransomware infection, it can easily mitigate the attack by simply recovering the most recently uncorrupted files to within 5 minutes of the attack. Nasuni's ability to restore files and volumes in minutes makes it a much faster recovery platform than traditional backup or cloud backup systems.

Nasuni protects data using unlimited snapshots in Azure Blob and eliminates the need for a physical backup infrastructure.

In the event that a more extensive DR operation is required, Nasuni can bring an entire site back online in about fifteen minutes. The reason why it is possible to perform such rapid recovery operations is because Nasuni handles data access differently than competing solutions like Azure File Sync or NetApp ONTAP. Those solutions are based around the use of storage tiering. Hot data is stored on-premises, while cooler data is stored in Azure cloud. Conversely, Nasuni stores all data in the cloud via frequent snapshots and uses virtual machines for caching copies of active files and metadata.

Having the ability to perform rapid recovery operations is not the only data protection benefit provided by Nasuni. Because of the way that Nasuni protects data using unlimited snapshots in Azure Blob, it eliminates the need for a physical backup infrastructure. This means that organizations can free themselves from the costs associated with their legacy backup system, such as the costs tied to purchasing backup media, purchasing backup software licenses, and maintaining backup hardware.

Migration Considerations

It is important to consider your actual migration path to the cloud. Neither Azure Files nor Azure NetApp Files can absorb data directly from on-premises Windows File Servers. Both require a multi-step migration process. If you are migrating from on-premises storage, you will want to look at how their process actually works. Is it site by site? Or can data be migrated from several locations simultaneously? Does it require that the WFS systems be offline for a significant period? Or is the vendor able to migrate then switch your endusers to the new system immediately with a clean sync-up?

Share Data with Azure Analytics Applications

Another particularly compelling benefit afforded by Nasuni's approach is that an organization's file data can be connected to any third-party cloud service. This means that organizations can leverage technologies such as AI, data analytics, and search as tools for finding additional business insights hidden within their file data. Best of all, Nasuni supports multi-cloud environments, so organizations can run any cloud service against their data, even if the data is all stored in Azure Blob.

Conclusion

Moving file data to Azure can be both practical and cost effective. It is worth noting however, that there are significant differences between the available solutions. Azure Files is useful for small businesses that may have a few file servers they want to move to the cloud, but it lacks the scalability needed by large enterprises, the sync servers become hard to manage when files need to be shares across more than a few locations, and it is expensive because of it does not leverage object storage and backup is not included. Similarly, Azure NetApp Files works well for high-performance workloads, but it is even more expensive than Azure Files, making it cost-prohibitive for general purpose file storage.

Nasuni with Azure Blob gives enterprise-class organizations a practical migration option for their Windows File Servers that is fully supported by Microsoft. Additionally, Nasuni overcomes the most troubling logistical challenges that have historically been associated with migrating file servers to the cloud, such as latency, cost, and migration simplicity. Nasuni's use of Azure Blob storage as its back-end and its continuous snapshot technology, which eliminates the need for backup makes it much less expensive than Azure Files or Azure NetApp Files. Because it scales to support multi-site file sharing and intrinsically provides data protection, it's ideal for remote work environments. Additionally, it scales to support rapid ransomware recovery or recovery from multiple site outages. Overall Nasuni is an ideal cloud migration solution for the current cost-sensitive times.

Figure 1: Features Comparison Chart

Category	Azure Files w/ Sync	Azure NetApp Files	Nasuni
Capacity per volume	100 TB	368 TB	Unlimited capacity Unlimited file size
# Sites	Less than 100 Windows server required at each site	Single site standalone Multi-site with addition of Global File Cache (GFC)-windows server required at each site	Unlimited[1] Free VM license for each Site[2]
File sharing, synchronization, global file locking	Requires Azure Files with Sync[4] No file locking	Multi-site synchronization and file sharing requires addition of GFC Remote file locking requires addition of GFC	Included
End-User File Access Performance	Not optimized for multi-site file sharing performance	Optimized for single site access	LAN speed access to files regardless of distance/ source location
Migration from WFS Environment	Via Azure Files Sync or Robocopy	Robocopy or multi-site via Global File Cache	Cloud Migrator or Robocopy Migrate up to 100 TB in 30 days or less End-to-end migration support
Azure Storage Category	Azure Files[3]	Azure NetApp Files (NetApp arrays running in Azure)	Azure Blob
Data Protection Strategy	Not included; Separate purchase of Azure Backup Service required.	Not included; Snapshot to cloud available separately	Backup included; DR included
Range of Workloads	General purpose file storage Application file access Multi-site file access	High-performance database file access	General purpose file storage Application file access Multi-site file access, sharing, locking

[1] *The average Nasuni customer has 3+ sites and over 30 TB under management*

[2] *Nasuni VM's are stateless*

[3] *You cannot create Azure file shares from Blob storage accounts or premium general purpose (GPv1 or GPv2) storage accounts.*

[4] *Azure File Sync supports syncing only with an Azure file share that's in the same region as the Storage Sync Service.*

[5] *Nasuni also supports WFS migration to AWS S3 and GCP*
https://cloud.netapp.com/global-file-cache-faq#technical-questions
https://docs.microsoft.com/en-us/azure/storage/files/storage-files-faq

Figure 2: Sample Cost Comparison Chart (at time of publication—9/2020)

Category	Azure Files	Azure NetApp Files	Azure Blob and Nasuni
Price Per Year, Per TB @ 100 TB	Standard $864*	Standard $1812	Starting at $650/TB
	Premium $3812	Premium $3615 Ultra $4826	
Backup	Azure Backup (added cost)	NetApp Cloud Backup (added cost)	Nasuni Continuous File Versioning (included)
On-Premises File Sync	Azure File Sync and Windows File Servers (added cost)	NetApp Global File Cache and Windows Servers (added cost)	Nasuni Edge Appliance VMs (included)

Includes capacity required to retain the maximum number of snapshots allowed (200) at a rate of .1% change per snapshot. For the latest pricing, check the Azure Pricing Calculator.

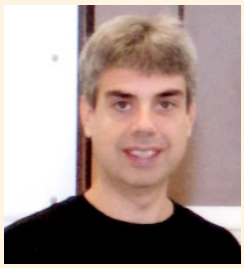

About Brien M. Posey

As an internationally best selling technology author and 19 time Microsoft MVP, Brien Posey has written or contributed to dozens of books, and created numerous full-length video training courses on a huge variety of IT and space related topics. In addition, Brien has published over 7000 technical articles and white papers for various Web sites and periodicals. In addition to his writing, Brien routinely records webcasts and speaks internationally at various live events IT on subjects ranging from information technology to astronautics.

Prior to going freelance, Brien was CIO for a national chain of hospitals and healthcare facilities. He has also served as the lead Network Engineer for the United States Department of Defense at Fort Knox, and has worked as a Network Administrator for some of the nation's largest insurance companies. Brien also previously served in a volunteer capacity as the Association of Spaceflight Professionals' Technology and Engineering Group lead.

In addition to his ongoing work in information technology, Brien is a Commercial Scientist-Astronaut Candidate. Over the last several years, Brien has been training extensively in preparation for a mission to study polar mesospheric clouds from space.

One Marina Park Drive
Boston, USA
www.nasuni.com

SAL-0128 09/20