



Trends & Challenges in Enterprise Data Management

January 2021

Contents

1. High Level IT Trends
2. Data Management Trends
3. Common Use Cases of Data Portability
4. Challenges for Enterprise Data Portability
5. How Parsec Makes Data Portability Simple
6. Summary and Conclusion

High Level IT Trends

A high-level review of current trends in IT is helpful to set context for discussion of Enterprise Data Management. The following are some of the key issues facing the industry today:

Security and Compliance - *IDC reports that global spending on security continues to grow at a compound annual rate of almost 10% and will top \$133 billion by 2022.*

Data security concerns due to government regulation, industry specific requirements and threats such as hacking, malware and ransomware continue to occupy the minds of enterprise IT professionals. The ongoing investments in solutions required for Security, Risk Management and Compliance are consuming an ever-growing share of the overall IT budget... *As a result, other IT priorities are forced to do more with less.*

Data Explosion - *90% of data in the world today has been created in the past two years and by 2025 humans alone will be adding an astonishing 463 Exabytes per day! IoT devices will add even more.*

Primary storage array capacities will continue to grow geometrically. This in turn creates the need for tiering and retention strategies to make sure data is on the most economical tier based on where and how long the data is to be stored. Current data management solutions limit where and when data can be moved so technology lags economic need.

Data Retention - *The average requirement for Enterprise data retention time is now six years.*

Generally, retention periods are even longer for personnel, health, and financial records, and a bit less for operational data. Regardless, this means that as the primary storage grows, archival storage will continue to mushroom as well.

Emerging Storage Technologies - *According to Yahoo Finance, the global NVMe market size is projected to grow from USD 44.6 billion in 2020 to USD 163.5 billion by 2025; a CAGR of 29.7%.*

These protocols unlock even higher performance due to a quantum leap in number and depth of queues compared to legacy SATA and SAS. Only 3 years ago NVMe and NVMe-oF were unheard of. They are the latest example of the rapid innovation in storage that nimble IT teams are looking to exploit.

Cloud Adoption – *According to IDC, Enterprise spending on cloud (30% of total) surpassed spending for on-prem solutions (24%) for the first time in 2020. And there is no end in sight.*

AWS, Azure and Google Cloud all reported growth rates of 30-48% in Q3 2020. The pressure on IT to reduce CAPEX, lower management costs and provide better accessibility is driving a tidal wave of moving applications to the cloud – public and private. Implicit in this trend is the need to easily and securely migrate data from on-premise to the Cloud.

Data Management Trends

The explosive growth of and increased retention requirements for data, combined with the budget squeeze caused by never-ending security concerns, is causing many IT leaders to rethink their approach to managing data. As a result, several trends are emerging:

Enterprise Data Portability - *According to the Institute of Electrical and Electronics Engineers (IEEE), cloud interoperability and data portability pose the major challenges to enterprise adoption of cloud computing services.*

Whether archiving “cold” data to lower cost storage tiers, moving “hot” data to higher performance storage, or migrating data to new platforms on demand, IT leaders increasingly expect to be able to put their data where they want it, when they want it.

Cloud Migration - As stated above, the public cloud is more often than not the destination for data being moved as part of a data portability strategy. At the same time, the costs to access (or retrieve) data stored in the public cloud is giving rise to many private cloud providers. Tools that can support both, as well as on-premise object stores, are gaining in popularity.

Independence from Storage Vendors – *Gartner states “The industry doesn’t have a Storage Management problem. It has a Data Management problem.”*

Historically, IT teams viewed data management challenges through the prism of their storage vendor. In other words, as a storage management problem. This thinking has begun to change however. Vendor specific tools that only work with their own hardware, and are often even limited to specific models within their brand, have resulted in customers having to learn and maintain a multitude of tools. For many IT leaders, this “tool sprawl” – and the cost and complexity that comes along with it -- is no longer acceptable.

Innovative companies such as Parsec have emerged to solve this problem and thought leaders in IT are increasingly looking to these data management vendors as valued partners.

Common Use Cases of Enterprise Data Portability

Migrations are probably the most common use case for data portability in any Enterprise. They are triggered by several common events:

Tech Refresh - Customers who acquire an array with new features and technology desire to place the platform into production as soon as possible, so that they can realize immediate benefits and maximize the return on their capital investments. To achieve these goals, Enterprise IT managers need a platform that supports multiple protocols and data transports which are storage brand agnostic.

Storage Consolidation - Another classic data portability use case. A new storage array may include migration tools which assist moving from legacy storage, but in most cases these are limited to arrays within a specific brand of storage or family of storage within a brand. But if the legacy storage is on disparate arrays, the migration tool will fall short and the administrator will need to augment the migration with multiple tools and external subject matter experts to achieve heterogeneous support.

Array Rebalancing - Numerous legacy arrays do not support this as a native function so users must look for additional software tools to meet the need. For arrays that have disks added to them, restriping and rebalancing may be necessary to gain maximum advantage of all the drives in the storage pool.

Data Center Moves - Massive amounts of data, disparate arrays and multiple protocols are typically in play. Having to use multiple tools and consuming cycles of numerous stakeholders renders the data center move to extreme complexity. These activities also generally entail hiring outside experts to plan and deliver the project. This adds costs and introduces risks that must be budgeted for and addressed.

Replication is another common use case for data portability. Unlike migrations, which are often “one-time events,” replicating data to a remote DR site is an on-going use case. In the event of an outage at the primary data center, users can be diverted to the backup site and continue working with minimal disruption. Unfortunately, as with migration tools, vendor specific replication tools are limited to their own hardware forcing customers to purchase 2x of everything!

Challenges for Enterprise Data Portability

Project Complexity - *Tech Data* observes that migration project expenses are on average 200% of the array acquisition cost. With an average of four years useful life, the annual operating expenses associated to migration represent over 50% of the acquisition cost.

Complexity is a significant barrier to Enterprise Data Migrations. Enterprise storage migration costs can exceed \$15,000 per terabyte. On average, migration planning consumes 70% of cycles while execution is relegated to 30%. Complexity is identified as the culprit behind these startling numbers and vendors who can simplify data migration are welcome.

Disparate Data Storage - Virtually all enterprise data centers contain storage from multiple vendors, numerous storage protocols and data transports. This impedes data portability because vendor specific toolsets are confined to their own brand while third-party tools are often limited to specific protocols and transports.

Point Solution Limitations - Third-Party solutions may be used across array brands but are limited to specific protocols and do not address the broad issue of general data portability. At the enterprise level this results in complexity - choosing and learning a series of point solutions which lengthens project planning, requires coordinated procurement, tool training cycles and ultimately, protracted execution timeline.

Cloud Cost - While moving to the cloud can offer significant cost savings, the reality is that it's more complicated than that. Cloud vendors typically advertise low fees to store data, but charge additional fees to access or retrieve data (egress fees). These costs can add up, and enterprises wanting to move to the cloud have to invest more time to fully understand the life cycle costs. Worse, some move ahead without realizing the cost implications only to face an unpleasant surprise later.

Security and WAN Optimization - Migrating or replicating data across a WAN requires both encryption in motion and network optimization. Often this means multiple tools, which adds further complexity and friction.

User Connection Complexity - An often overlooked part of a successful replication implementation is making it easy for users to access the remote data store in the event of an outage in the primary data center. Most replication vendors focus on the "moving the data" part and leave it to the customer to sort out the "user cutover" part. As they say, the devil's in the details...

How Parsec Makes Enterprise Data Portability Simple


Parsec employs a unique approach in many ways. Parsec connects at the network level so does not have to be concerned with the complexity of application and database rights and security. Once Parsec is zoned into the network, LUN's and Files are ready to migrate. This radically simplifies planning and procurement while reducing project complexity by a degree of magnitude. Parsec competitors spend voluminous cycles planning, meeting and coordinating across the enterprise for issues and roadblocks that do not exist with a Parsec solution.

Impact on Project Planning and Execution

Due to the superior architecture and feature inclusions of the Parsec appliance, the traditional planning and meeting cycles associated with data migrations are flattened. The single appliance approach with Block, File and Object support simplifies tool selection and operation. The ability to support data migration between disparate arrays without gateways removes another category of preparation and planning. Parsec makes long haul migration or replication from one site to another fast and secure, transferring data using our native WAN accelerator and encryption.

When replicating to a secondary DR data center, Parsec can, on command, retrieve the file to object map (including directory structures, metadata and permissions) and reconstruct entire filers, mount points or sub directories at the recovery site, without requiring the same model of storage at both ends.

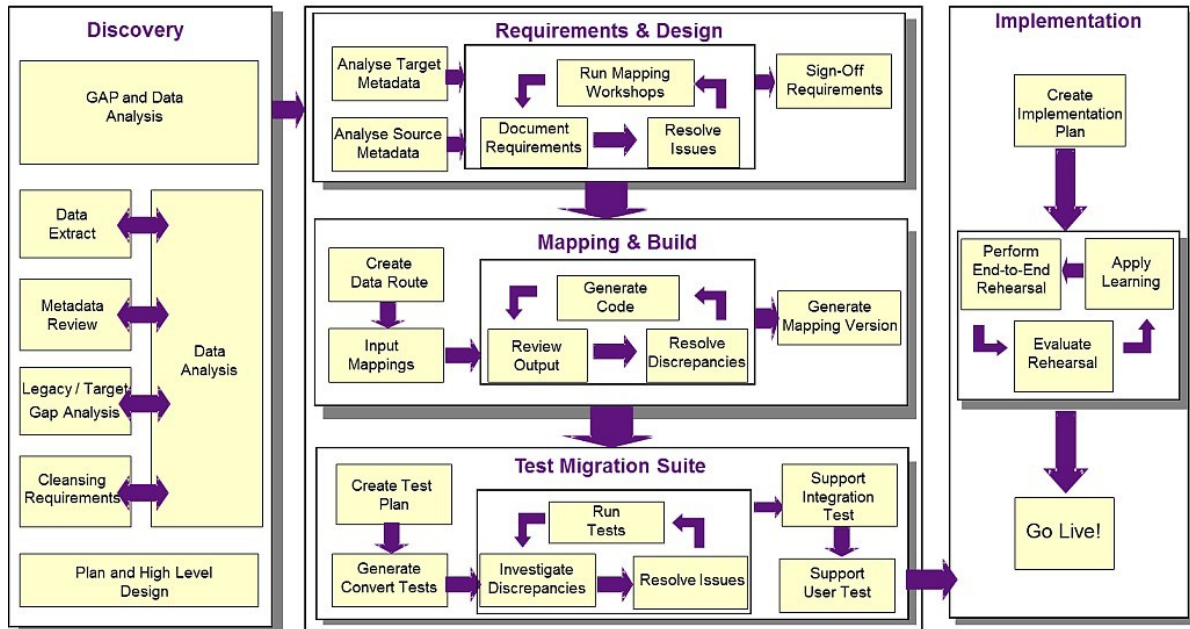
Steps for Parsec Planned Migration




1. Parsec is attached to the array via the network (IP, FC, iSCSI, NVME) infrastructure and begins to discover data that can be migrated.
2. The Customer can then choose what data they want to move and where they want it to reside.
3. Parsec begins migrating data to the new array while applications and users are still accessing the old array

4. Once migration is complete and validated, clients are pointed to the new array
5. Parsec is ready for the next migration

Traditional Process Workflow for Data Migrations:



Simplified Parsec Process Workflow for Data Migrations:



The Parsec appliance approach simplifies SAN/NAS connection. Our unique approach needs less people to operate and requires more basic skills instead of advanced skillsets and associated specialized training. In aggregate, these Parsec advantages shorten the migration process from weeks and months to days and hours, while mitigating risk and removing stress. Furthermore, Parsec licenses its solution on a per appliance basis with unlimited data transfers. There are no hidden data tariffs as with others who vend their tool and then add a “Per TB migrated” surcharge for using the solution. With Parsec, data portability projects do not disrupt applications or impact users, while significantly reducing both CAPEX and OPEX.

Summary and Conclusion

The need and desire for data portability has never been greater in the enterprise. Yet the challenges of project complexity, security, and limitations of point solutions are significant. Legacy, storage-centric solutions are not enough to overcome these challenges.

Innovative companies such as Parsec have emerged to provide vendor agnostic solutions to these challenges. Supporting File, Block, and Object data, Parsec scales to billions of files and provides for multiple concurrent migrations and/or replications. To find out more about Parsec Labs and our Data Management Appliances, please visit our website at www.parseclabs.com