Solving Top SharePoint Online Challenges

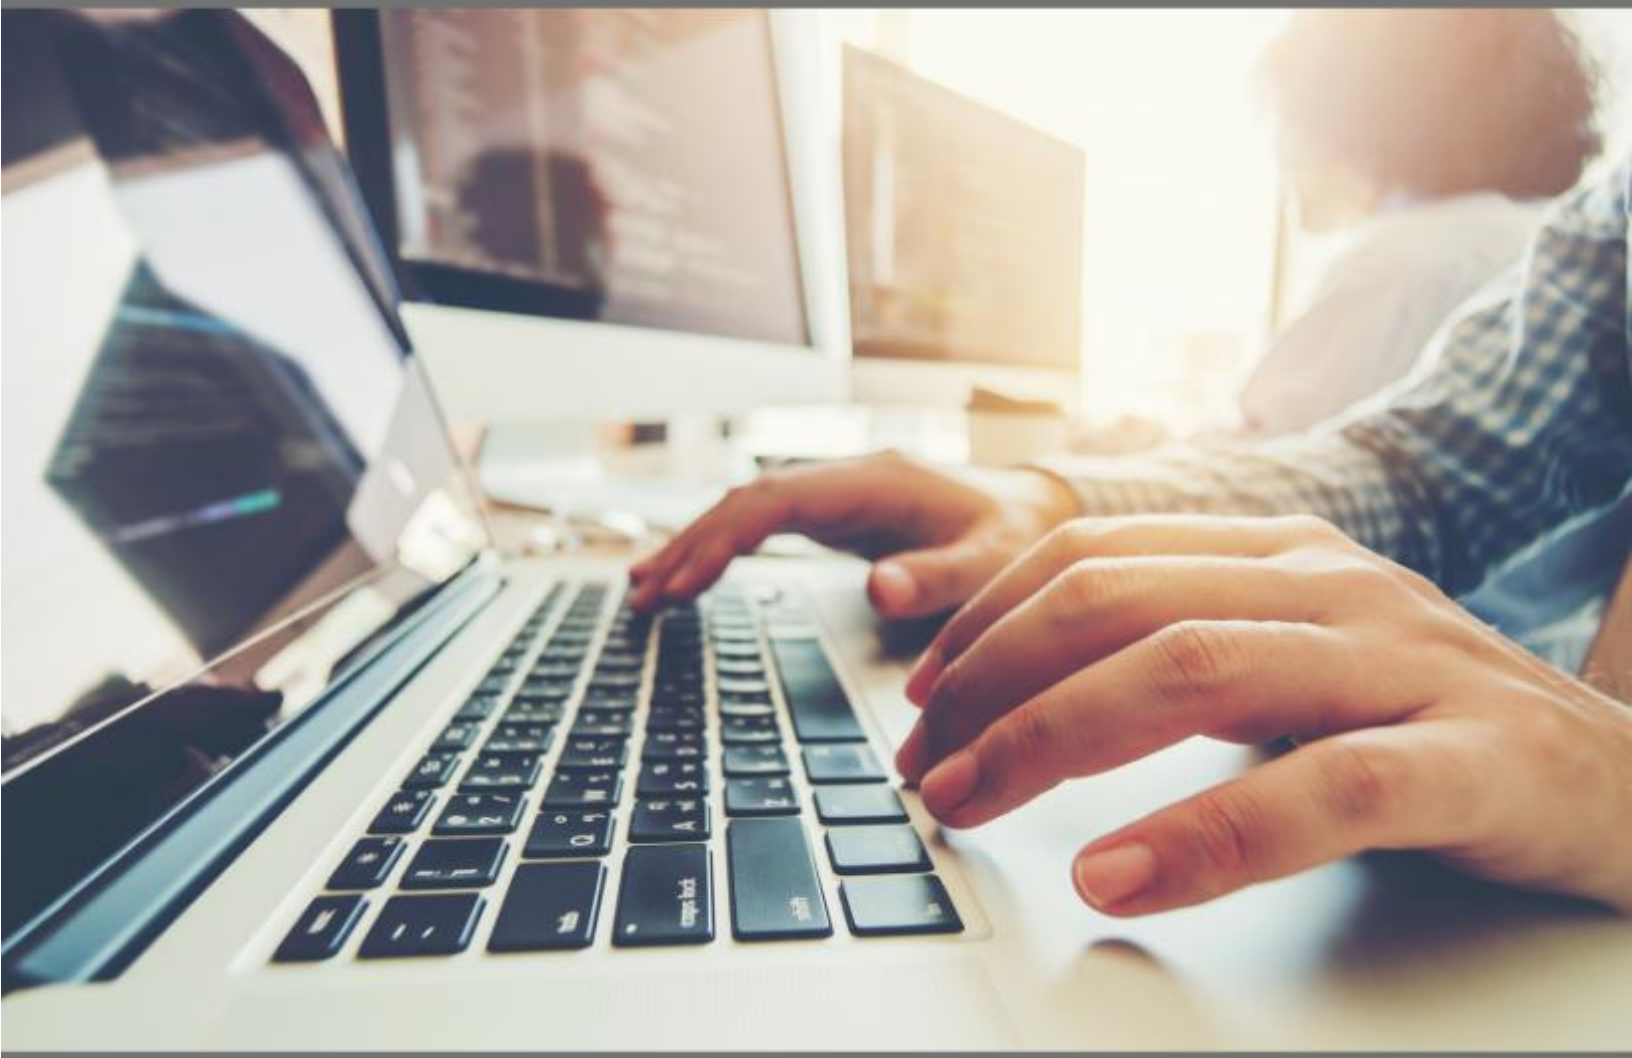

ENZO

Contents

Introduction

Microsoft SharePoint Online, a product of the Office 365 suite, delivers unparalleled flexibility and simplicity, which explains the high adoption rate of the platform. However organizations looking to adopt or migrate to SharePoint Online for improved operational efficiency need to better understand how inherent limitations of an online service may affect their experience. This white paper outlines the top limitations that customers can expect from SharePoint Online, how these limitations can affect their operations, and the available workarounds. In addition, this white paper introduces Enzo Unified, a new platform that directly addresses these limitations.

SharePoint Online Limitations

5,000 Item Limit

SharePoint lists, including document lists, can hold a large number of items. The current limit is 30 million documents per library. However, the List View threshold is 5,000. This means that while a SharePoint list can hold a large number of items, it can only display 5,000 at a time. As you can see in Figure 1, the SharePoint Online site cannot display list items when the view contains more than 5,000

items. As we will see later, this limitation impacts Microsoft Flow, ODBC drivers, reporting systems and integration platforms.

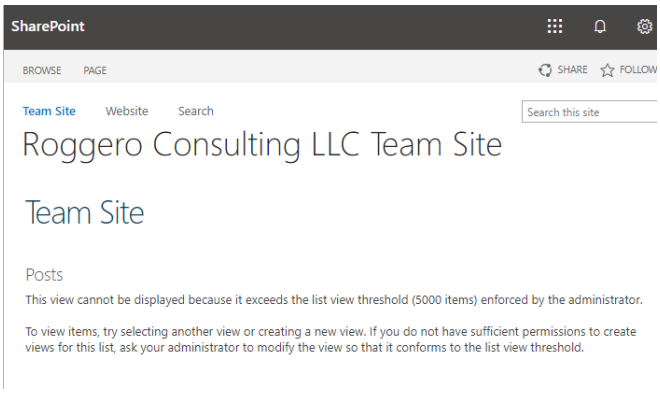

Figure 1 - 5,000 Item Limitation in SharePoint Online

It is worth nothing that several workarounds are available, including creating indexes and views with filters that leverage those indexes. Nevertheless, these workarounds can only partially answer this limitation since many scenarios require access to the entire list.

In order to help deal with this limitation, the SharePoint API provides a paging mechanism that allows developers to build programs with the necessary logic to progressively consume 5,000 items, until all the items have been read. This paging mechanism is predominantly useful for custom application development, but can be slow and is not very productive for most reporting and integration systems.

Rate Limiting

Since SharePoint Online is a shared service hosted by Microsoft, your service may be rate limited under certain scenarios when you attempt to extract too many records, too quickly. While Microsoft does not officially publish rate limiting information, rate limiting falls into two usage patterns as seen in Figure 2: sporadic and heavy.

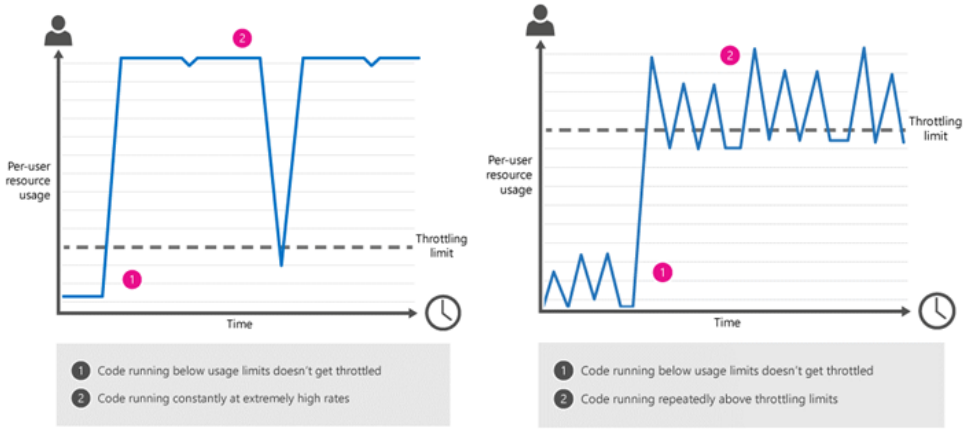

Figure 2 - Rate Limiting in SharePoint Online (source: Microsoft)

The sporadic usage pattern takes place when too many requests are made over a short period of time, hence increasing the amount of resources being used over time; the heavy usage pattern takes place when few requests are being made, but a lot of resources are being consumed per request.

The only recovery mechanism available when rate limiting is taking place is to wait. Developers should be warry of rate limiting and apply a back-off mechanism when making API calls to SharePoint Online. The best way to avoid rate limiting however is to offload the data into a database and point the offending workload to the database instead of SharePoint.

Far Connection and Shared Service

As a hosted service, Microsoft offers SharePoint Online in various data centers across the globe and will select the most appropriate data center based on your location when you sign up for the service. For the United States, SharePoint Online is currently hosted in Chicago, IL, Des Moines, IA, Quincy, WA and San Antonio, TX as reported by Microsoft's Data Center page. As a result, you can expect your data to be stored relatively far away from your office location, which can impact performance for integration and reporting scenarios that are dealing with larger lists.

In addition to being far, your SharePoint Online service also shares resources with other tenants, and as a result, your workload is prone to resource limitations. Either Microsoft implements a form of resource leveling (so that no one tenant can consume too many resources at one time), or your workloads compete with other tenants' workloads. Or both.

Combined with rate limiting and the 5,000 item limitation, consumers can expect to experience an impact on workload performance, affecting integration and reporting scenarios.

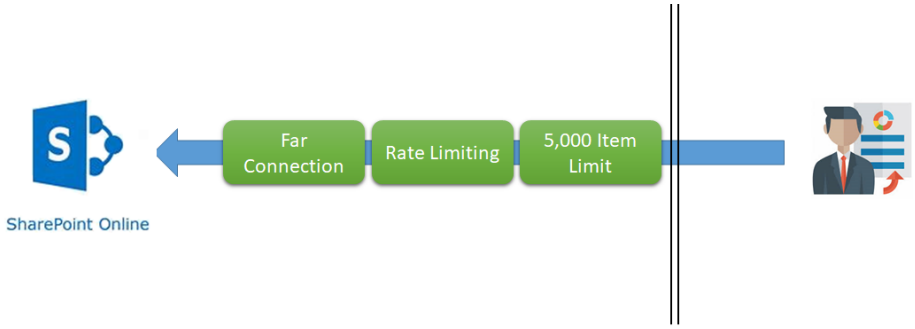

Figure 3 - Consumer Impact of SharePoint Online

No Direct SQL Access

Organizations that already use SharePoint Server may be accessing directly the SQL Server database on which SharePoint runs to extract information and enable specific integration scenarios. However with SharePoint Online, access to the underlying database is not possible. While ODBC drivers can be used to circumvent this specific limitation, the previously discussed limitations cannot be solved by ODBC drivers; this means that accessing SharePoint Online data using SQL commands can become challenging, especially for large lists.

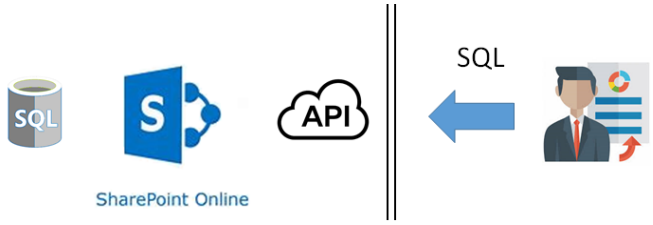

Figure 4 - No SQL Access to SharePoint Online

Impact Analysis

Reporting & Power Users

Power Users generally use SQL as the language of choice due to its simplicity; however SharePoint Online can make it very difficult for power users to remain effective even when using ODBC drivers. That's because ODBC drivers are subject to the 5,000 item limitation, rate limiting and far connection limitations. This also impacts the ability to write reports against large lists. Most systems, such as Excel, SSRS and other reporting tools will not be able to provide access to large lists in real-time.

ODBC drivers are subject to SharePoint Online limitations.

Limitation	Solved by ODBC Driver
Access SharePoint Lists using SQL	Yes
5,000 Item Limitation	No
Rate Limiting	No
Far Connection Performance	No

Far connections also imply a reduction in overall performance, which can translate into minutes of wait time. In our tests, we experienced a significant delay in reading data from large lists stored in SharePoint Online compared to obtaining the same records from a SharePoint database locally. As shown in Figure 5 we observed a throughput of roughly 5 seconds per 1,000 records when reading from a SharePoint Online list, retrieving all available fields; accessing the database directly in a SQL Server database usually allows throughput of 200ms per 1,000 items or better.

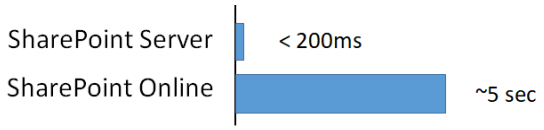

Figure 5 - SharePoint Server vs. SharePoint Online Performance

Development & Integration

From a development and integration standpoint, tools that deal with large lists face the same limitations outlined previously. For example, Microsoft Flow will experience the same 5,000 item limitation as ODBC drivers, and will also be relatively slow since it must go through the API layer of SharePoint Online.

Other integration tools providing additional flexibility, such SSIS, while subject to the same limitations, will provide a more comprehensive design surface so that integration engineers can begin to design

workflows that leverage the paging capability of SharePoint Online. This means that organizations may need to build and maintain complex workflows, increasing time to market and support costs.

Developers working with the SharePoint Server API already know how complex it can be to build solutions. Introducing rate limiting, paging and far connections creates a significant design and development burden that wasn't a factor until now.

Solutions

To address the above limitations and reduce the impact to your organization when adopting SharePoint Online, you should be aware of available workarounds that can help you avoid some of the pitfalls, as well as how a new SharePoint proxy, Enzo Unified, can help eliminate these limitations.

Workarounds

Let's start by discussing available workarounds, with their pros and cons.

- **Create Smaller Views**
  This approach has to do with the 5,000 item limitation; in theory you should be able to create views that return fewer than 5,000 items, so that access to those views will not hit the threshold. For example if your list contains sales records for the US, creating a view by Zip Code per year may solve the problem. You should however make sure to build the views before you reach the first 5,000 items in the list, otherwise creating the view will fail.

- **Index your Views**
  This approach allows you to improve performance of your list data by allowing SharePoint to index on fields that are queried the most. If the index contains fewer than 5,000 items, it can also be used to create Filtered Views on larger lists.

- **Use Modern**
  The modern experience allows SharePoint Online users to navigate large lists on the SharePoint Online portal. While this is a good solution for end-user experience, it will not help with the use cases described in this white paper.

For more information about how to best manage large lists in SharePoint Online see Microsoft's reference documentation on this topic.

Enzo Solutions

Enzo Unified was built to answer the many challenges that hosted platforms pose, including paging, rate limiting, performance and integration. In general, Enzo offers two different solutions to these limitations: real-time (synchronous) and near-time (asynchronous).

Enzo removes the 5,000 item limitation of SharePoint Online using SQL or REST commands.

Real-Time Access using SQL or REST

Enzo offers real-time access to SharePoint Online list items through both native SQL commands and REST commands. As depicted in Figure 6 this capability allows developers, integration tools and reporting engines to communicate to SharePoint Online in real-time without

worrying about the 5,000 item limitation. Unlike an ODBC driver, Enzo was built as a server platform, and as a result can provide advanced capabilities and optimizations out of the box. [Click here to learn more about differences between Enzo and ODBC Drivers](#).

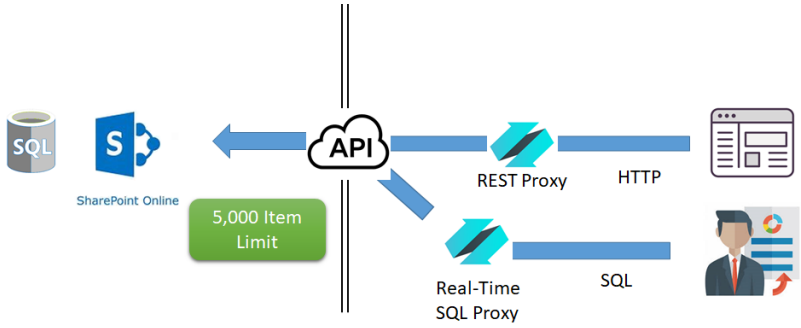

Figure 6 - Solving the 5,000 Item Limitation with Enzo

Because Enzo is a server product, it does not require client-side driver installation. The SQL proxy of Enzo offers wire-compatible SQL Server access, and as a result it is possible to query SharePoint Online directly from SQL Server triggers/procedures/views, and from SQL Server Management Studio.

In addition to its SQL proxy, Enzo also offers a REST compatible interface, which eliminates complex client-side SDKs, and enables simple webhook integration with SharePoint Online without having to worry about paging.

Near-Time using Edge Caching

In addition to real-time access to SharePoint Online lists, Enzo offers an edge caching mechanism that allows consumers to access SharePoint data in near-time very quickly, using both SQL and REST commands. As depicted in Figure 7 the same cache can be used by both SQL and REST calls. The Edge Caching option improves performance that is closer to running SharePoint Server on premises, and eliminates the rate limiting and 5,000 item limitation of SharePoint Online.

Enzo Edge Caching offers up to 10x performance improvement compared to accessing data from SharePoint Online directly.

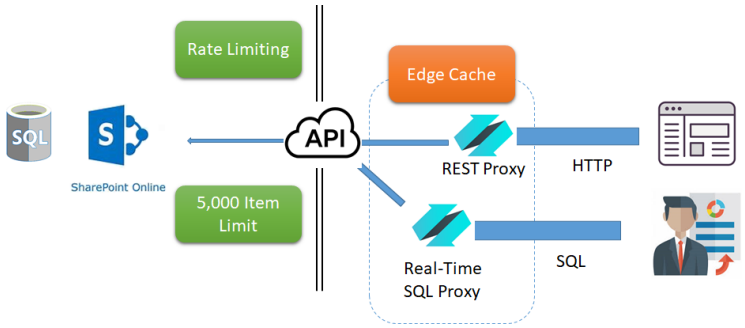

Figure 7 - Edge Caching with Enzo

Because Enzo behaves like a database engine, and stores its cache locally, it can respond to REST and SQL requests very quickly. The overall performance of Enzo matches that of a database engine, and has shown to offer up to 10x performance improvement compared to access of the same data against SharePoint Online.

The same cache can be used by both SQL and REST calls.

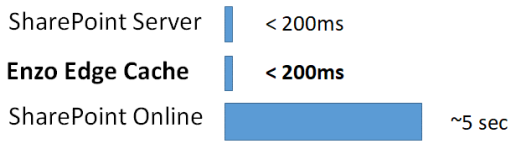

Figure 8 - Edge Cache Performance

Near-Time using Replication

Enzo Pipelines is a feature of Enzo that provides an advanced replication engine that offers both one time synchronization, and ongoing Change Data Capture (CDC). The CDC engine forwards changes detected in SharePoint to any number of systems using a Publish-Subscribe (Pub/Sub) pattern. In its simplest implementation as shown in Figure 9, Enzo Pipelines copy a SharePoint list to a SQL Server database table, and keeps the table synchronized by replicating changes made to the source list.

This capability offers an alternative to Enzo Edge Caching and solves all the limitations identified in this paper, including the 5,000 item limit, rate limiting, and the performance issues. However it does not provide the ability to service the data through REST commands out of the box since the data is stored in a SQL Server database.

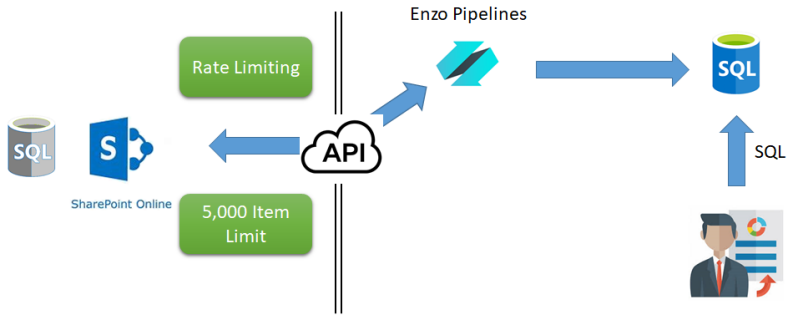

Figure 9 - Replicating SharePoint Lists with Enzo Pipelines

The primary advantage of this approach is to offer a full replication of SharePoint Online lists as if they were stored on the local network, offering complete backward compatibility with existing tools and reporting systems. Visit this page for a more in-depth overview of Enzo Pipelines.

Enzo Offerings

Two Enzo offerings are available: Enzo Server and Enzo Online. Enzo Server is a product that can be installed in your infrastructure, while Enzo Online is a Platform as a Service. The following matrix can help you decide which offering provides which capability.

Capability	Enzo Server	Enzo Online
Native SQL Access	Yes	No
REST Access	Yes	Yes
No 5,000 Item Limitation	Yes	Yes
Edge Caching	Yes	No
Replication (Pipeline)	Yes	Yes
Learn More	Learn More Download	Sign Up

Summary

This white paper provides an overview of the fundamental challenges organizations face when adopting or migrating to SharePoint Online, along with workarounds recommended by Microsoft. However the limitations of SharePoint Online impose various boundaries that neither ODBC drivers nor the workarounds suggested can solve for most implementations, such as the 5,000 item limit, far connections, rate limiting and performance.

Enzo was built to address hosted platform limitations and offers both real-time and near-time solutions, including native SQL and REST access, edge caching, and replication capabilities helping organizations beat the inherent limitations of SharePoint Online.

For more information, visit https://www.enzounified.com, contact us at info@enzounified.com or call 1 800-610-2521 to speak with a representative.