



Web Services for Metadata Archive

6-69021-01, Rev. A

Quantum[®]

StorNext Web Services for Metadata Archive, 6-69021-01, Rev. A, November 2021, Product of USA.

Quantum Corporation provides this publication “as is” without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Quantum Corporation may revise this publication from time to time without notice.

COPYRIGHT STATEMENT

© 2021 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. ActiveScale, DXi, DXi Accent, FlexSync, FlexTier, iLayer, Lattus, Quantum, The Quantum Logo, QXS, Scalar, StorNext, SuperLoader, Vision, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Quantum specifications are subject to change.



Contents

Preface	iv
Audience	iv
Chapter 1: Introduction	1
Introduction	1
Chapter 2: Basics	2
Basics	2
Character Encoding	2
Security	2
Non-realtime Nature of Data	3
Service Location	3
Chapter 3: API Calls	4
API Calls	4
Change Tracking	4



Preface

This document describes the commands supported by the StorNext Web Services for Metadata Archive and contains the following chapters:

- [Chapter 1: Introduction](#)
- [Chapter 2: Basics](#)
- [Chapter 3: API Calls](#)

Audience

This manual is written for StorNext operators, system administrators, and field service engineers.


Notational Conventions

This manual uses the following conventions:

Convention	Example
User input is shown in bold monospace font.	./DARTinstall
Computer output and command line examples are shown in monospace font.	./DARTinstall
User input variables are enclosed in angle brackets.	http://<ip_address>/cgi-bin/stats

Convention	Example
For UNIX and Linux commands, the command prompt is implied.	<code>./DARTinstall</code> is the same as <code># ./DARTinstall</code>
File and directory names, menu commands, button names, and window names are shown in bold font.	/data/upload
Menu names separated by arrows indicate a sequence of menus to be navigated.	Utilities > Firmware

The following formats indicate important information:

 **Note:** Note emphasizes important information related to the main topic.

 **Caution:** Caution indicates potential hazards to equipment or data.


 **WARNING:** Warning indicates potential hazards to personal safety.


- Right side of the system - Refers to the right side as you face the component being described.
- Left side of the system - Refers to the left side as you face the component being described.
- Data sizes are reported in base 10 (decimal) rather than base 2¹⁰ (binary). For example:
10,995, 116,277,769 Bytes are reported as 11.0 TB (decimal/1000). In binary, this value is 10 TiB (binary/1024).


Product Safety Statements









Quantum will not be held liable for damage arising from unauthorized use of the product. The user assumes all risk in this aspect.

This unit is engineered and manufactured to meet all safety and regulatory requirements. Be aware that improper use may result in bodily injury, damage to the equipment, or interference with other equipment.

 **WARNING:** Before operating this product, read all instructions and warnings in this document and in the *Quantum Products System, Safety, and Regulatory Information Guide*.

 **ADVARSEL:** Læs alle instruktioner og advarsler i dette dokument og i *Informationsvejledning vedrørende system-, sikkerheds- og lovbestemmelser for Quantum produkter, før produktet betjenes*.

 **AVERTISSEMENT :** Avant d'utiliser ce produit, lisez toutes les instructions et les avertissements de ce document et du *Guide d'informations sur le système, la sécurité et la réglementation de Quantum*.

-
-  **WARNUNG:** Lesen Sie vor der Inbetriebnahme dieses Produkts alle Anleitungen und Warnungen in diesem Dokument und im *System-, Sicherheits- und Betriebsbestimmungen-Handbuch für Quantum-Produkte*.
-
-  **ADVERTENCIA:** Antes de hacer funcionar este producto, lea todas las instrucciones y advertencias de este documento y de la *Guía de información normativa, del sistema y de seguridad de los productos de Quantum*.
-
-  **VARNING:** Läs igenom alla instruktioner och varningar i detta dokument och i *Quantums produktsystem, säkerhet och reglerande informationsguide* innan denna produkt används.
-
-  **ВНИМАНИЕ!** Перед началом эксплуатации данного изделия прочтите все инструкции и предупреждения, приведенные в настоящем документе и в *Руководстве по системе, технике безопасности и действующим нормативам компании Quantum*.
-
-  警告：本製品を使用される前に、本書と『*Quantum製品システム、安全、規制情報ガイド*』に記載されているすべての説明と警告をお読みください。
-
-  경고: 본 제품을 작동하기 전에 본 문서와 *Quantum 제품 시스템, 안전 및 규제 정보 설명서*에 있는 모든 지침과 경고를 참조합니다.
-
-  警告：在操作本产品之前，请阅读本文档和 *Quantum 产品系统、安全和法规信息指南*中的所有说明和警告。
-
-  警告：操作此產品前，請閱讀本檔案及 *Quantum 產品系統、安全與法規資訊指南*中的指示與警告說明。

 **אזהרה:** לפני ההפעלה של מוצר זה, קרא את כל ההוראות והאזהרות הכלולות במסמך זה וכן במדריך המידע בנושא מערכת, בטיחות ותקינה עבור מוצרי Quantum.

For the most up to date information on StorNext, see:

<https://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

Contacts

For information about contacting Quantum, including Quantum office locations, go to:

<https://www.quantum.com/aboutus/contactus/index.aspx>

Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

doc-comments@quantum.com

Getting More Information or Help

Quantum University, Quantum’s comprehensive service approach, leverages advanced data access and diagnostics technologies with cross-environment, multi-vendor expertise to resolve backup issues faster and at lower cost.

Accelerate service issue resolution with these exclusive Quantum University services:

- **Service and Support Website** - Register products, license software, browse Quantum Learning courses, check backup software and operating system support, and locate manuals, FAQs, firmware downloads, product updates and more in one convenient location. Get started at:

<https://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

- **eSupport** - Submit online service requests, update contact information, add attachments, and receive status updates via email. Online Service accounts are free from Quantum. That account can also be used to access Quantum’s Knowledge Base, a comprehensive repository of product support information. Get started at:

<https://www.quantum.com/customercenter/>

For further assistance, or for training opportunities, contact the Quantum Customer Support Center:

Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

<https://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

Worldwide End-User Product Warranty

For more information on the Quantum Worldwide End-User Standard Limited Product Warranty:

<https://www.quantum.com/serviceandsupport/warrantyinformation/index.aspx>



Chapter 1: Introduction

This chapter contains the following topics:

Introduction	1
--------------------	---

Introduction

You can enable the **mdarchive** service and web services on a StorNext file system to support a number of basic queries about content. The web services do not represent a comprehensive query interface to all content in the system, rather a means to avoid scanning a file system to locate content which matches specific criteria.

In application use, you are expected to combine the web services with an external database to serve a specific query domain, or some data movement / processing service which acts on the results.



Chapter 2: Basics

This chapter contains the following topics:

Basics	2
--------------	---

Basics

Character Encoding

- All queries use JSON encoding for content; no other data encoding is supported.
- File names are UTF-8 encoded. Both the macOS and Windows StorNext clients restrict file names to UTF-8. The Linux StorNext client does not; it can create file names which are not legal UTF-8 encoding and will fail to present in JSON.

Security

Access to the web service is controlled using a shared secret and access is using HTTPS. The shared secret is automatically generated by the file system and can be obtained in one of a number of ways to be covered later, depending on how the service is configured. This shared secret limits access to the query interface, it is not used to limit the returned information based on permissions in the file system. Access to the secret provides access to any content in the file system independent of the ownership of that data in the file system.

In order to protect against scenarios where an end user is allowed access to results from the interface, an extra level of security is needed to limit results to results the end user has file system permission to see.

A certificate can be installed into the system, but by default the web service providing the information uses a self-signed certificate which would be rejected by many web clients without special handling.

Non-realtime Nature of Data

The query results from **mdarchive** are based on a different copy of metadata than the actual file system data. This state is not maintained synchronously with the state of the actual file system and, as a result, may lag behind the actual file system contents. The architecture is designed to minimize this lag, but system load factors do mean that the returned information may report things which have since been changed in the file system. For example, the existence of a file may be reported, but the file may subsequently have been renamed or moved and this has to be handled by the caller.

All requests make use of a sequence number, this always increases with time including across service restarts or system crashes. This sequence number is not guaranteed to be monotonically increasing, it may skip values. It is referred to as a marker, or transaction ID (tid). A request for file system changes would ask for content starting at a specific marker and be returned changes in order. A new sequence number is returned with the response which can be used to get subsequent changes.

Service Location

The web services to access **mdarchive** are provided by the file system **fsm** process. The port number is assigned by the system on service startup (it does pay attention to the StorNext cvports file). The host is the host the **fsm** is running on.

The `restAccess` configuration parameter controls the existence of the services and how their locations are found. Allowed values are:

- **Disabled** : No REST services.
- **Privileged**: REST enabled, only a privileged user can access the location information. In this mode, the extended attributes **system.com.quantum.sn_rest_uri** and **system.com.quantum.sn_rest_challenge** contain the HTTPS endpoint and the shared secret for accessing services.
- **Enabled**: REST enabled and any user may access the location information. In this mode, the **user.com.quantum.sn_rest_uri** and **user.com.quantum.sn_rest_challenge** versions can also be used.

These are the Linux extended attribute names; on macOS there is no system or user prefix. There is no support for accessing these on Windows except using a **cvapi** call as extended attributes are not exposed to a user space.

All endpoints in this document consist of the **postfix** which goes after the endpoint address.



Chapter 3: API Calls

This chapter contains the following topics:

API Calls	4
-----------------	---

API Calls

Change Tracking

Operating system event notification interfaces do not work with StorNext as they are implemented above the actual file system and can only track operations on a single host. You can use a services such as Linux **inotify**; however, if a directory is changed on another host, a notification is not generated.

The primary interfaces offered by **mdarchive** are around reporting changes in the file system; the interface does cover metadata changes from all hosts accessing the file system.

There are four basic calls for change tracking that you can access as **GET** or **POST** calls.

- For a **GET** request, an argument is expected on the URI.
- For a **POST** request, an argument is expected in the request body as a JSON object.

```
/api/metadb/v1/marker
```

```
/api/metadb/v1/change_tree
```

```
/api/metadb/v1/change_list
```

```
/api/metadb/v1/change_handles
```

The typical call sequence is to use the **marker** call initially followed by repeated use of one of the other calls. The **marker** call returns the range of markers in the file system:

```
{  
  "marker": 28054589,  
  "first_marker": 108,  
  "uuid": "eb6f980e-ad77-491a-b854-e1bce856e6f1",  
  "time": 1554663837  
}
```

Above, the **first_marker** represents the oldest marker recorded by **mdarchive** since it started capturing state, **marker** is the last change known about. The **uuid** is a unique identifier for the **mdarchive** instance, if the **mdarchive** is rebuilt, the **uuid** changes. Finally, **time** is the seconds since the epoch associated with the last marker update.

There are three possible parameters to a request:

Parameter	Description
min_age	The minimum age in seconds to report on as an integer. The marker value returned is close to this amount of time old. If content is to be looked at once it is at least a certain age, then this value can be used. As time advances, more recent changes are returned.
sync	A boolean, using sync causes the mdarchive to flush more recent changes through from the file system. Quantum recommends you not use this in normal circumstances as it adds significant cost to the query.

Parameter	Description
-----------	-------------

details A boolean, using this causes a verbose report which provides a mapping between marker values and timestamps. An array of records sorted by time is returned:

```
{
  "marker": 26724648,
  "first_marker": 108,
  "uuid": "eb6f980e-ad77-491a-b854-e1bce856e6f1",
  "time": 1554518791,
  "details": [
    {
      "date": "2019-04-02T22:23:48Z",
      "time": 1554243828,
      "marker": 23391569
    },
    {
      "date": "2019-04-03T15:11:44Z",
      "time": 1554304304,
      "marker": 23392018,
      "operations": 449
    },
    {
      "date": "2019-04-03T15:12:25Z",
      "time": 1554304345,
      "marker": 23396350,
      "operations": 4332
    },
  ],
}
```

A set of marker values are returned along with the epoch time, a GMT time and a count of the changes between the two markers. In general these details will not be useful beyond tracking the rate of change in the file system.

After a marker has been obtained, the file system can be scanned. At the end of this process you know that the scan has seen all content as it existed at the marker and possibly some which came into existence after the marker. This corresponds to the set of content which existed before starting to monitor the file system.

Subsequent to a one time scan, one of the other three calls can be used to locate content changed since the point in time represented by the marker obtained here.

The different requests return information about changed content in different formats. They each take a marker parameter and return changes which occurred after that marker. Several other modifiers are supported and discussed later.

Request	Description
/api/metadb/v1/change_tree	Returns a hierarchical set of directories as a nested json object which indicate where in the file system hierarchy there were changes. Just the directory names are returned, the user is expected to use this to perform a selective scan of the file system limited to only the identified directories, no other parts of the file system name space contain changes in the time range covered. When used in combination with another database containing a copy of the metadata, this allows identifying where content has been removed from the system as well as where it has been added.
/api/metadb/v1/change_list	Returns a list of path names in the file system of content which has been changed or added since the specified marker.
/api/metadb/v1/change_handles	Returns a list of file handles into the file system for content which has been changed or added since the marker. These are special file names which can be used to access the file without using an absolute path to get to them. Using these in the file system is a privileged operation, reading a file in this manner will also not trigger storage manager activity on it – but will return zeros if it is truncated.

Request arguments for these calls are about filtering the result set in various ways, and controlling the number of results.

Argument	Purpose	Change_list	Change_handle	Change_tree
marker	Starting marker value for returns. Required integer .	Yes	Yes	Yes
uuid	The uuid value from marker call. If the mdarchive has been rebuilt and the value no longer matches, the call will return an error. Optional string .	Yes	Yes	Yes
max_results	Maximum results to return, optional integer .	Yes	Yes	Yes
min_age	Minimum age in seconds of results to return, optional integer .	Yes	Yes	Yes
end_marker	A maximum marker to report results for. This can be used instead of min_age , but requires the caller to know the value. Will be overridden by the min_age parameter. Optional integer .	Yes	Yes	Yes
sync	Force mdarchive up to date with file system. Optional boolean	Yes	Yes	Yes

Argument	Purpose	Change_list	Change_handle	Change_tree
path	A directory path to limit results to. Only children of this directory will be included in results. Optional string . If the path cannot be found an error is returned.	Yes	Yes	Yes
relative	Only meaningful when path used. Instead of returning paths based on the root of the filesystem, a query returns paths relative to the specified path argument. Optional boolean .	Yes	No	No
affinity	Only report files or directories with the specified StorNext affinity. Optional string .	Yes	Yes	Yes
array	Return results as a JSON array of objects instead of an object containing objects. Contents are sorted by age – oldest first. Json objects do not actually imply an order in iterating their keys, arrays do. The original API used objects, but this can cause problems for some JSON library implementations accessing the contents in sequence. Optional boolean .	Yes	Yes	No
max_size min_size	Limit reported results to files greater than min_size or less than max_size bytes. One or both can be specified. Optional integer .	Yes	Yes	Yes
file_only	Only report files in results. Optional boolean .	Yes	Yes	No
dir_only	Only report directories in results. Optional boolean .	Yes	Yes	No
max_total_blocks	Stop reporting results after files containing this number aggregate file system blocks have been reported. Optional integer .	Yes	Yes	No

Argument	Purpose	Change_list	Change_handle	Change_tree
details	<p>When returning objects the object name is the file path, the object is null. When returning an array, the array is just strings. The details argument changes this. Each file reported on has an object with information about the file, the contents of which are controlled by the details string argument contents.</p> <ul style="list-style-type: none"> • s: File size reported. • n: File name reported. • a: Time since last change in seconds. • A: File affinity. • b: File blocks. • t: File marker (or tid) reported. 	Yes	Yes	No
match	Only return files which match a regular expression. Does not filter directory names.	Yes	Yes	Yes
exclude	Filter out files which match a regular expression.	Yes	Yes	Yes



Quantum[®]

Quantum technology, software, and services provide the solutions that today's organizations need to make video and other unstructured data smarter – so their data works for them and not the other way around. With over 40 years of innovation, Quantum's end-to-end platform is uniquely equipped to orchestrate, protect, and enrich data across its lifecycle, providing enhanced intelligence and actionable insights. Leading organizations in cloud services, entertainment, government, research, education, transportation, and enterprise IT trust Quantum to bring their data to life, because data makes life better, safer, and smarter. Quantum is listed on Nasdaq (QMCO) and the Russell 2000[®] Index. For more information visit www.quantum.com.

www.quantum.com | 800-677-6268