

**Oracle® ZFS Storage Appliance Analytics
Guide, Release OS8.7.0**

ORACLE®

Part No: E78913-01
March 2017

Oracle ZFS Storage Appliance Analytics Guide, Release OS8.7.0

Part No: E78913-01

Copyright © 2014, 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Référence: E78913-01

Copyright © 2014, 2017, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

Working with Analytics	11
▼ Setting a Retention Policy (BUI)	12
▼ Setting a Retention Policy (CLI)	12
▼ Setting a Hostname Lookup Policy (BUI)	13
▼ Setting a Hostname Lookup Policy (CLI)	14
Managing Worksheets	15
▼ Creating a Worksheet (BUI)	15
▼ Creating a Worksheet (CLI)	15
▼ Closing a Worksheet (BUI)	16
▼ Saving a Worksheet (BUI)	16
▼ Renaming a Worksheet (CLI)	17
▼ Destroying a Worksheet (BUI)	18
▼ Destroying a Worksheet (CLI)	18
▼ Cloning a Worksheet (BUI)	19
▼ Removing a Dataset from a Worksheet (BUI)	19
▼ Removing a Dataset from a Worksheet (CLI)	20
▼ Viewing When a Worksheet was Last Modified (CLI)	21
▼ Showing Graph Hierarchy	21
▼ Viewing Available Datasets (BUI)	23
▼ Viewing Available Datasets (CLI)	24
▼ Reading Datasets (CLI)	25
▼ Suspending and Resuming a Dataset (BUI)	26
▼ Suspending and Resuming a Dataset (CLI)	27
▼ Suspending and Resuming All Datasets (CLI)	28
▼ Discarding Data in a Dataset (BUI)	29
▼ Discarding Data in a Dataset (CLI)	29
▼ Identifying CPU Performance Issues (BUI)	32
▼ Identifying CPU Performance Issues (CLI)	32
▼ Identifying Network Performance Issues (BUI)	34

▼ Identifying Network Performance Issues (CLI)	35
▼ Identifying Memory Performance Issues (BUI)	36
▼ Identifying Memory Performance Issues (CLI)	37
▼ When to Add the First Read Cache Device (BUI)	39
▼ When to Add the First Read Cache Device (CLI)	39
▼ When to Add More Read Cache Devices (BUI)	42
▼ When to Add More Read Cache Devices (CLI)	43
▼ When to Add the First Write Log Device (BUI)	44
▼ When to Add the First Write Log Device (CLI)	45
▼ When to Add More Write Log Devices (BUI)	48
▼ When to Add More Write Log Devices (CLI)	48
▼ When to Add More Disks (BUI)	50
▼ When to Add More Disks (CLI)	50
▼ Configuring a Threshold Alert (BUI)	52
▼ Configuring a Threshold Alert (CLI)	52
▼ Exporting a Worksheet (BUI)	53
▼ Exporting a Worksheet (CLI)	54
▼ Downloading a Dataset to a CSV File (BUI)	55
▼ Viewing a Dataset in a CSV Format (CLI)	55
Analytics Data Retention Policies	57
Analytics Data Retention Policies	58
Data Retention Properties	59
Understanding Analytics Worksheets	61
Worksheet Graphs and Plots	61
Adjusting Graphs	62
Adjusting Quantize Plots	63
Background Patterns	64
Toolbar Reference	64
Worksheet Tips	66
Saved Worksheet Properties	66
BUI Icon Reference	66
About Analytics Datasets	69

Understanding Analytics Statistics	71
Storage Performance Impact	72
Execution Performance Impact	75
Statistics Actions	77
Default Statistics	78
CPU: Percent Utilization	79
Cache: ARC Accesses	81
Cache: L2ARC I/O Bytes	84
Cache: L2ARC Accesses	84
Capacity: Capacity Bytes Used (BUI)	85
Capacity: Capacity Bytes Used (CLI)	87
Capacity: Capacity Percent Used (BUI)	89
Capacity: Capacity Percent Used (CLI)	91
Capacity: Meta Device Capacity Bytes Used (BUI)	93
Capacity: Meta Device Capacity Percent Used (BUI)	94
Capacity: System Pool Bytes Used	94
Capacity: System Pool Percent Used	95
Data Movement: Shadow Migration Bytes	96
Data Movement: Shadow Migration Ops	97
Data Movement: Shadow Migration Requests	97
Data Movement: NDMP Bytes Statistics	98
Data Movement: NDMP Operations Statistics	99
Data Movement: Replication Bytes	99
Data Movement: Replication Operations	100
Disk: Disks	101
Disk: I/O Bytes	102
Disk: I/O Operations	104
Network: Device Bytes	106
Network: Interface Bytes	107
Protocol: SMB Operations	108
Protocol: Fibre Channel Bytes	109
Protocol: Fibre Channel Operations	110
Protocol: FTP Bytes	112
Protocol: HTTP/WebDAV Requests	113
Protocol: iSCSI Bytes	114
Protocol: iSCSI Operations	115
Protocol: NFSv[2-4] Bytes	116

Protocol: NFSv[2-4] Operations	117
Protocol: OISP Bytes	119
Protocol: OISP Operations	120
Protocol: SFTP Bytes	122
Protocol: SMB/SMB2 Bytes	123
Protocol: SRP Bytes	124
Protocol: SRP Operations	124
Using Advanced Analytics Statistics	127
CPU: CPUs	128
CPU: Kernel Spins	129
Cache: ARC Adaptive Parameter	129
Cache: ARC Evicted Bytes	130
Cache: ARC Size	130
Cache: ARC Target Size	131
Cache: DNLC Accesses	132
Cache: DNLC Entries	132
Cache: L2ARC Errors	133
Cache: L2ARC Size	134
Data Movement: NDMP Bytes Transferred to/from Disk	134
Data Movement: NDMP Bytes Transferred to/from Tape	135
Data Movement: NDMP File System Operations	135
Data Movement: NDMP Jobs	136
Data Movement: Replication Latencies	136
Data Movement: Replication Send/Receive Bytes	137
Disk: Average Number of I/O Operations	138
Disk: Percent Utilization	138
Disk: ZFS DMU Operations	139
Disk: ZFS Logical I/O Bytes	140
Disk: ZFS Logical I/O Operations	140
Memory: Dynamic Memory Usage	141
Memory: Kernel Memory	141
Memory: Kernel Memory in Use	142
Memory: Kernel Memory Allocated, But Not in Use	142
Network: Datalink Bytes	143
Network: IP Bytes	144
Network: IP Packets	144

Network: TCP Bytes	145
Network: TCP Packets	145
Network: TCP Retransmissions	146
System: NSCD Backend Requests	146
System: NSCD Operations	147

Working with Analytics

The Oracle ZFS Storage Appliance is equipped with an advanced DTrace-based facility for server analytics so you can examine the details of different layers of the storage stack. Analytics provides real-time graphs of various statistics, which you can save for later viewing. It is designed for both long-term monitoring and short-term analysis.

To manage and monitor analytics, use these tasks:

- Setting a Retention Policy - [BUI](#), [CLI](#)
- Setting a Hostname Lookup Policy - [BUI](#), [CLI](#)
- Creating a Worksheet - [BUI](#), [CLI](#)
- Closing a Worksheet - [BUI](#)
- Saving a Worksheet - [BUI](#)
- Cloning a Worksheet - [BUI](#)
- Removing a Dataset from a Worksheet - [BUI](#), [CLI](#)
- Viewing When a Worksheet was Last Modified - [CLI](#)
- Showing Graph Hierarchy - [BUI](#)
- Viewing Available Datasets - [BUI](#), [CLI](#)
- Reading Datasets - [CLI](#)
- Suspending and Resuming a Dataset - [BUI](#), [CLI](#)
- Suspending and Resuming All Datasets - [CLI](#)
- Discarding Data in a Dataset - [BUI](#), [CLI](#)
- Identifying CPU Performance Issues - [BUI](#), [CLI](#)
- Identifying Network Performance Issues - [BUI](#), [CLI](#)
- Identifying Memory Performance Issues - [BUI](#), [CLI](#)
- When to Add the First Read Cache Device - [BUI](#), [CLI](#)
- When to Add More Read Cache Devices - [BUI](#), [CLI](#)
- When to Add the First Write Log Device - [BUI](#), [CLI](#)
- When to Add More Write Log Devices - [BUI](#), [CLI](#)
- When to Add More Disks - [BUI](#), [CLI](#)
- Configuring a Threshold Alert - [BUI](#), [CLI](#)

- Exporting a Worksheet - [BUI](#), [CLI](#)
- Downloading a Dataset to a CSV File - [BUI](#)
- Viewing a Dataset in a CSV Format - [CLI](#)

For information about Analytics worksheets, datasets, and statistics, see these topics:

- “Analytics Data Retention Policies” on page 57
- “Understanding Analytics Worksheets” on page 61
- “About Analytics Datasets” on page 69
- “Understanding Analytics Statistics” on page 71
- “Using Advanced Analytics Statistics” on page 127

▼ Setting a Retention Policy (BUI)

Use the following task to set a retention policy, which limits the amount of data collected over a retention period. It is strongly recommended that you set retention policies to meet your minimum business requirements, in order to preserve disk space. The maximum retention period is two years.

1. **Go to Analytics > Settings.**
2. **Type an integer value into the text box.**
3. **From the drop-down menu, select one of the following retention periods: hours, days, weeks, months.**
4. **Click APPLY to save the retention settings.**

▼ Setting a Retention Policy (CLI)

Use the following task to set a retention policy, which limits the amount of data collected over a retention period. It is strongly recommended that you set data retention policies to meet your minimum business requirements, in order to preserve disk space. The maximum retention period is two years, or 17520 hours.

1. **GO to analytics settings.**

```
hostname:> analytics settings
```
2. **Enter show to view a list of data retention properties.**

```
hostname:analytics settings> show
Properties:
    retain_second_data = all
    retain_minute_data = all
    retain_hour_data = all
    hostname_lookup = true
```

3. Set the data retention interval and define a retention policy.

The retention policy is measured in hours. If you want to set your policy to a certain number of days, weeks, or months, you must first calculate the number of hours in that period. In the following example, set `retain_second_data=72` retains data recorded at a per-second interval for 72 hours, or 3 days.

```
hostname:analytics settings> set retain_second_data=72
    retain_second_data = 3 days (uncommitted)
```

4. Enter `commit`.

```
hostname:analytics settings> commit
```

▼ Setting a Hostname Lookup Policy (BUI)

Use the following task to enable or disable a hostname lookup policy for statistics that are broken down by client or by hostname. Enabling hostname lookup performs hostname resolution for each IP address in the breakdown, and then stores and displays the data by hostname. This is the default setting.

Disabling hostname lookup saves all client breakdowns by their IP addresses, which reduces overhead and improves performance.

Changing the hostname lookup policy results in a breakdown that contains both hostnames and IP addresses. For example, if you disable hostname lookup and then enable it, the old breakdowns display as IP addresses and the new breakdowns display as hostnames.

1. **Go to Analytics > Settings.**
2. **Perform one of the following actions:**
 - **To save client breakdowns by hostname: Select "Enable Hostname Lookup."**
 - **To save client breakdowns by IP address: Deselect "Enable Hostname Lookup."**

3. Click **APPLY**.

▼ Setting a Hostname Lookup Policy (CLI)

Use the following task to enable or disable a hostname lookup policy for statistics that are broken down by client or by hostname. Enabling hostname lookup performs hostname resolution for each IP address in the breakdown, and then stores and displays the data by hostname. This is the default setting.

Disabling hostname lookup saves all client breakdowns by their IP addresses, which reduces overhead and improves performance.

Changing the hostname lookup policy results in a breakdown that contains both hostnames and IP addresses. For example, if you disable hostname lookup and then enable it, the old breakdowns display as IP addresses and the new breakdowns display as hostnames.

1. **Go to analytics settings.**

```
hostname:> analytics settings
```

2. **Enter show.**

```
hostname:analytics settings> show
Properties:
    retain_second_data = 1 weeks
    retain_minute_data = 2 weeks
    retain_hour_data = 730 days
    hostname_lookup = true
```

3. **Perform one of the following actions:**

- **To save client breakdowns by hostname: Ensure the policy is set to true. If not, enter set hostname_lookup=true and then enter commit.**

```
hostname:analytics settings> set hostname_lookup=true
    hostname_lookup = true (uncommitted)
hostname:analytics settings> commit
```

- **To save client breakdowns by IP address: Ensure the policy is set to false. If not, enter set hostname_lookup=false and then enter commit.**

```
hostname:analytics settings> set hostname_lookup=false
    hostname_lookup = false (uncommitted)
hostname:analytics settings> commit
```

Managing Worksheets

Worksheets are the main interface for Analytics. To work with worksheets, use the following tasks:

- Creating a Worksheet - [BUI](#), [CLI](#)
- Closing a Worksheet - [BUI](#)
- Saving a Worksheet - [BUI](#)
- Renaming a Worksheet - [CLI](#)
- Destroying a Worksheet - [BUI](#), [CLI](#)
- Cloning a Worksheet - [BUI](#)
- Removing a Dataset from a Worksheet - [BUI](#), [CLI](#)
- Viewing When a Worksheet was Last Modified - [CLI](#)
- Showing Graph Hierarchy - [BUI](#)
- Viewing Available Datasets - [BUI](#), [CLI](#)
- Reading Datasets - [CLI](#)
- Suspending and Resuming All Datasets - [BUI](#), [CLI](#)
- Discarding Data in a Dataset - [BUI](#), [CLI](#)

▼ Creating a Worksheet (BUI)

Use the following procedure to create a worksheet using the BUI. To save the worksheet after you have created it, see [“Saving a Worksheet \(BUI\)”](#) on page 16.

1. **Go to Analytics > Open Worksheets > New.**
2. **Click Untitled worksheet, click in the field, and type a name for the worksheet.**
3. **Click the add icon  and select a statistic to add to the worksheet.**

▼ Creating a Worksheet (CLI)

Use the following procedure to create a worksheet using the CLI. To add statistics to the worksheet after you have created it, use the tasks in [“Working with Analytics”](#) on page 11. The worksheet is automatically saved after you create it.

1. **Go to analytics worksheets.**

```
hostname:> analytics worksheets
```

2. **Enter create and type the new name of your worksheet.**

```
hostname:analytics worksheets> create example_1
```

3. **Enter show to view a list of open worksheets, including the one that you created.**

```
hostname:analytics worksheets> show
```

```
Worksheets:
```

WORKSHEET	OWNER	NAME
worksheet-000	root	example_1

▼ Closing a Worksheet (BUI)

Use the following procedure to close an open worksheet and consequently discard all statistics in the worksheet.

In the CLI, only saved worksheets can be viewed; therefore, there is no CLI equivalent procedure to close an open worksheet.

1. **Go to Analytics > Open Worksheets.**
2. **Click Worksheets to view a list of open worksheets.**
3. **Select the worksheet that you want to close.**
4. **Click Close.**

▼ Saving a Worksheet (BUI)

Worksheets can be saved for later viewing. In consequence, all visible statistics will be archived, meaning that they will continue to save new data after the saved worksheet has been closed.

In the CLI, worksheets are automatically saved after they are created.

1. **Go to Analytics > Open Worksheets.**

2. **Click Worksheets to view a list of open worksheets.**
3. **Select the worksheet that you want to save.**
4. **Click Save.**

Note - When you create a worksheet on a standalone or clustered system, the worksheet statistics are not permanently saved on the controller until you click Save.

▼ Renaming a Worksheet (CLI)

Use the following procedure to rename a saved worksheet.

1. **Go to analytics worksheets.**

```
hostname:> analytics worksheets
```

2. **Enter show to view a list of saved worksheets.**

```
hostname:analytics worksheets> show
Worksheets:
```

WORKSHEET	OWNER	NAME
worksheet-000	root	worksheet
...		

3. **Select the worksheet you want to rename, and then list its properties.**

```
analytics worksheets> select worksheet-000
```

```
analytics worksheet-000> ls
```

Properties:

```

        uuid = a442e761-4048-4738-b95f-be0824d7ed09
        name = worksheet
        owner = root
        ctime = 2016-12-14 03:58:28
        mtime = 2016-12-14 03:58:28
```

4. **Enter set name= and the new name for the worksheet.**

```
analytics worksheet-000> set name=test
```

```
name = test (uncommitted)hostname:
```

5. **Commit the change and list the properties to confirm the new worksheet name.**

```
analytics worksheet-000> commit
analytics worksheet-000> ls
Properties:
    uuid = a442e761-4048-4738-b95f-be0824d7ed09
    name = test
    owner = root
    ctime = 2016-12-14 03:58:28
    mtime = 2016-12-14 03:58:28
```

▼ Destroying a Worksheet (BUI)

Use the following procedure to destroy a saved worksheet.

Datasets in saved worksheets are archived. As such, data is not discarded when a worksheet is destroyed. To discard data in a dataset, see [“Discarding Data in a Dataset \(BUI\)” on page 29](#).

1. **Go to Analytics > Saved Worksheets.**
2. **Hover over a worksheet and click the discard icon .**
3. **Click OK to confirm your action.**

▼ Destroying a Worksheet (CLI)

Use the following procedure to destroy a saved worksheet.

Datasets in saved worksheets are archived. As such, data is not discarded when a worksheet is destroyed. To discard data in a dataset, see [“Discarding Data in a Dataset \(CLI\)” on page 29](#).

1. **Go to analytics worksheets.**

```
hostname:> analytics worksheets
```

2. **Enter show to view a list of saved worksheets.**

```
hostname:analytics worksheets> show
Worksheets:
```

```
WORKSHEET      OWNER  NAME
worksheet-000  root   example_1
```

3. Enter `destroy` and the worksheet that you want to destroy.

```
hostname:analytics worksheets> destroy worksheet-000
```

4. Enter `Y` to confirm your action.

```
This will destroy "worksheet-000". Are you sure? (Y/N) Y
hostname:analytics worksheets>
```

▼ Cloning a Worksheet (BUI)

Use the following procedure to clone, or create a copy, of a worksheet. To save the cloned worksheet after you have created it, see [“Saving a Worksheet \(BUI\)” on page 16](#).

This procedure cannot be performed in the CLI.

- 1. Go to Analytics > Open Worksheets.**
- 2. Click Worksheets to view a list of open worksheets.**
- 3. Select the worksheet that you want to clone.**
- 4. Click Clone.**
- 5. To name the cloned worksheet, click on the worksheet name and type a new name for the worksheet.**

Related Topics

- [“Closing a Worksheet \(BUI\)” on page 16](#)
- [“Removing a Dataset from a Worksheet \(BUI\)” on page 19](#)

▼ Removing a Dataset from a Worksheet (BUI)

Datasets in saved worksheets are archived. As such, data is not discarded when a dataset is removed from a worksheet. To discard data in a dataset, see [“Discarding Data in a Dataset \(BUI\)” on page 29](#).

1. **Go to Analytics > Saved Worksheets.**
2. **Click on the worksheet from which you want to remove a dataset.**
3. **Click the exit icon  on the upper right-hand corner of the statistic to remove it from the worksheet.**

▼ Removing a Dataset from a Worksheet (CLI)

Datasets in saved worksheets are archived. As such, data is not discarded when a dataset is removed from a worksheet. To discard data in a dataset, see [“Discarding Data in a Dataset \(CLI\)” on page 29](#).

1. **Go to analytics worksheets.**

```
hostname:> analytics worksheets
```

2. **Enter show to view a list of open worksheets.**

```
hostname:analytics worksheets> show
Worksheets:
```

WORKSHEET	OWNER	NAME
worksheet-000	root	example_1
worksheet-001	root	example_2

3. **Enter select and the worksheet from which you want to remove a dataset.**

```
hostname:analytics worksheets> select worksheet-000
```

4. **Enter show to view a list of datasets in the worksheet.**

```
hostname:analytics worksheet-000> show
```

Properties:

```
uuid = e268333b-c1f0-401b-97e9-ff7f8ee8dc9b
name = 830 MB/s NFSv3 disk
owner = root
ctime = 2009-9-4 20:04:28
mtime = 2009-9-4 20:07:24
```

Datasets:

DATASET	DATE	SECONDS	NAME
dataset-000	2009-9-4	60	nic.kilobytes[device]

```
dataset-001  2009-9-4      60  io.bytes[op]
```

5. Enter **remove** and the dataset that you want to remove.

```
hostname:analytics worksheet-000> remove dataset-000
This will remove "dataset-000". Are you sure? (Y/N)
```

6. Enter **Y** to confirm your action.

The modified worksheet is automatically saved.

▼ Viewing When a Worksheet was Last Modified (CLI)

Use this procedure to view the time that a saved worksheet was last modified. This time can be useful for auditing purposes, and for determining the time that a dataset within the worksheet was modified. For example, it is helpful to know when a dataset was added to a worksheet when troubleshooting a problem. The worksheet last modified time can also be compared to such events as a dataset suspension. This procedure is not supported in the BUI.

1. Go to **analytics worksheets**.

```
hostname:> analytics worksheets
```

2. Enter **select** and the name of the saved worksheet.

```
hostname:analytics worksheets> select worksheet-001
```

3. Enter **get mtime**.

```
hostname:analytics worksheet-001> get mtime
      mtime = 2015-6-1 13:09:01
```

▼ Showing Graph Hierarchy

Graphs that are broken down by filename have a special feature that allows you to view the hierarchical breakdown of traced filenames. As with graphs, the left panel will show components based on the statistic breakdown. If the filenames become too long to fit in the left panel, you can expand it by clicking and dragging the divider between the panel and the graph.

1. Create a worksheet as described in [“Creating a Worksheet \(BUI\)” on page 15](#).

2. Click the add icon  and select a statistic broken down by file name.

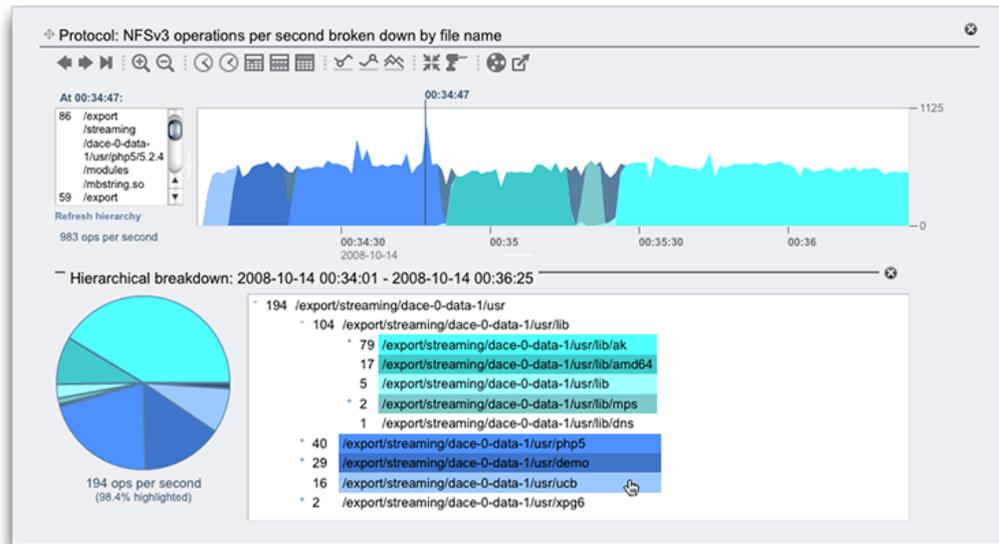
Not all statistics can be broken down by file name. The following statistics can be broken down by file name:

- CACHE > ARC accesses
- CACHE > L2ARC accesses
- DATA MOVEMENT > Shadow migration bytes
- DATA MOVEMENT > Shadow migration ops
- DATA MOVEMENT > Shadow migration requests
- PROTOCOL > FTP bytes
- PROTOCOL > HTTP/WebDAV requests
- PROTOCOL > NDMP bytes
- PROTOCOL > NDMP operations
- PROTOCOL > NFSv2 bytes
- PROTOCOL > NFSv2 operations
- PROTOCOL > NFSv3 bytes
- PROTOCOL > NFSv3 operations
- PROTOCOL > NFSv4 bytes
- PROTOCOL > NFSv4 operations
- PROTOCOL > NFSv4.1 bytes
- PROTOCOL > NFSv4.1 operations
- PROTOCOL > SFTP bytes
- PROTOCOL > SMB operations
- PROTOCOL > SMB2 operations

3. Click Show hierarchy on the left side of the statistic.

4. Click Refresh hierarchy to update the pie-chart and tree-view.

5. Click the exit icon , in the upper-right corner, to close the hierarchical breakdown.



Related Topics

- [“Adjusting Graphs” on page 62](#)
- [“Adjusting Quantize Plots” on page 63](#)
- [“Background Patterns” on page 64](#)
- [“Toolbar Reference” on page 64](#)

▼ Viewing Available Datasets (BUI)

Statistics that are being viewed in an open worksheet are temporary datasets that disappear when the worksheet is closed. You can view these temporary datasets, as well as statistics that are archived to disk, as a list on one page.

On the BUI, you can also view since when the analytics dataset was created and when the data was last accessed. The last access time of an analytics dataset is updated when a dataset is created, read, saved, resumed, or used to plot graph in the BUI.

- **Go to Analytics > Datasets.**

▼ Viewing Available Datasets (CLI)

Use the following procedure to view available datasets using the CLI.

1. **Go to analytics datasets.**
2. **Enter `show` to view a list of active and suspended datasets.**

In the example below, `dataset-007` is a temporary statistic because ONDISK size is zero. All other statistics are archived.

Note - The names of the statistics are abbreviated versions of what is visible in the BUI. For example, `dnlc.accesses` is short for Cache: DNLC accesses per second.

```
hostname:analytics datasets> show
Datasets:

DATASET   STATE   INCORE ONDISK NAME
dataset-000 active   674K   35.7K arc.accesses[hit/miss]
dataset-001 active   227K   31.1K arc.l2_accesses[hit/miss]
dataset-002 active   227K   31.1K arc.l2_size
dataset-003 active   227K   31.1K arc.size
dataset-004 active   806K   35.7K arc.size[component]
dataset-005 active   227K   31.1K cpu.utilization
dataset-006 active   451K   35.6K cpu.utilization[mode]
dataset-007 active   57.7K    0 dnlc.accesses
dataset-008 active   490K   35.6K dnlc.accesses[hit/miss]
dataset-009 active   227K   31.1K http.reqs
dataset-010 active   227K   31.1K io.bytes
dataset-011 active   268K   31.1K io.bytes[op]
dataset-012 active   227K   31.1K io.ops
...
```

3. **To view specific dataset properties, enter `select` and the name of the dataset.**

```
hostname:analytics datasets> select dataset-007
```

4. **Enter `show` to list the properties for the selected dataset. The properties include when the dataset was created and last accessed. The last access date and time is updated when a dataset is created, read, saved or resumed.**

```
hostname:analytics dataset-007> show
Properties:
name = dnlc.accesses
grouping = Cache
explanation = DNLC accesses per second
incore = 65.5K
size = 0
suspended = false
since = 2017-1-2 08:30:11
last_access = 2017-10-3 01:35:47
```

▼ Reading Datasets (CLI)

Use the following procedure to read datasets.

In the BUI, the same information is presented in a graph. For more information, see [“Worksheet Graphs and Plots” on page 61](#).

1. **Go to analytics datasets.**
2. **Enter show to view a list of available datasets.**
3. **Enter select and the name of the dataset that you want to read.**

```
hostname:analytics datasets> select dataset-007
```

4. **Enter read and the number of previous seconds to display.**

```
hostname:analytics dataset-007> read 10
DATE/TIME           /SEC      /SEC BREAKDOWN
2015-10-14 21:25:19    137      - -
2015-10-14 21:25:20    215      - -
2015-10-14 21:25:21    156      - -
2015-10-14 21:25:22    171      - -
2015-10-14 21:25:23   2722      - -
2015-10-14 21:25:24    190      - -
2015-10-14 21:25:25    156      - -
2015-10-14 21:25:26    166      - -
2015-10-14 21:25:27    118      - -
2015-10-14 21:25:28   1354      - -
```

Breakdowns will also be listed, if available. The following example shows CPU utilization broken down by CPU mode (user/kernel), which corresponds to dataset-006.

In this example, line 21:30:10 shows 14% kernel time and 1% user time, which adds up to 15% total utilization.

```
hostname:analytics datasets> select dataset-006
hostname:analytics dataset-006> read 5
DATE/TIME          %UTIL    %UTIL BREAKDOWN
2015-10-14 21:30:07      7        6 kernel
0 user
2015-10-14 21:30:08      7        7 kernel
0 user
2015-10-14 21:30:09      0        - -
2015-10-14 21:30:10     15       14 kernel
1 user
2015-10-14 21:30:11     25       24 kernel
1 user
```

5. To print comma separated values (CSV) for a number of seconds of data, enter `csv` and the number of seconds.

```
hostname:analytics datasets> select dataset-022
hostname:analytics dataset-022> csv 10
Time (UTC),Operations per second
2015-03-21 18:30:02,0
2015-03-21 18:30:03,0
2015-03-21 18:30:04,0
2015-03-21 18:30:05,0
2015-03-21 18:30:06,0
2015-03-21 18:30:07,0
2015-03-21 18:30:08,0
2015-03-21 18:30:09,0
2015-03-21 18:30:10,0
2015-03-21 18:30:11,0
```

▼ Suspending and Resuming a Dataset (BUI)

Use the following procedure to suspend and resume a single dataset. It is not possible to suspend or resume all datasets with one action using the BUI; you must use the CLI.

To suspend or resume all datasets at once, see [“Suspending and Resuming All Datasets \(CLI\)” on page 28](#).

1. **Go to Analytics > Datasets.**
2. **Hover over a dataset and click the Suspend/Resume icon  to suspend it.**
The green icon indicating that a dataset is actively collecting data will turn gray.
3. **Hover over a suspended dataset and click the Suspend/Resume icon  again to resume it.**
The gray icon indicating that a dataset is suspended will turn green.

▼ Suspending and Resuming a Dataset (CLI)

The CLI has the ability to suspend and resume individual or all datasets. Use the following procedure to suspend or resume an individual dataset. To suspend or resume all datasets at once, see [“Suspending and Resuming All Datasets \(CLI\)” on page 28](#).

1. **Go to analytics datasets.**
2. **Enter `select` and the name of the dataset that you want to suspend.**

```
hostname:analytics datasets> select dataset-043
```

3. **Enter `set suspended=true`.**

```
hostname:analytics dataset-043> set suspended=true
suspended = true (uncommitted)
```

4. **Enter `commit`.**

```
hostname:analytics dataset-043> commit
```

5. **To resume the dataset, enter `set suspended=false`.**

```
hostname:analytics dataset-043> set suspended=false
suspended = false (uncommitted)
```

6. **Enter `commit`.**

```
hostname:analytics dataset-043> commit
```

▼ Suspending and Resuming All Datasets (CLI)

Use the following procedure to suspend or resume all datasets at once. To suspend or resume an individual dataset, see [“Suspending and Resuming a Dataset \(CLI\)” on page 27](#).

In the BUI, it is not possible to suspend or resume all datasets at once. Datasets must be suspended or resumed individually.

1. Go to analytics datasets.

2. Enter `suspend` to suspend all datasets.

```
hostname:analytics datasets> suspend
This will suspend all datasets. Are you sure? (Y/N)y
```

3. Enter `y` to confirm that you want to suspend all datasets.

4. Enter `show` to view a list of suspended datasets.

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE ONDISK NAME
dataset-000 suspend  638K   584K arc.accesses[hit/miss]
dataset-001 suspend  211K   172K arc.l2_accesses[hit/miss]
dataset-002 suspend  211K   133K arc.l2_size
dataset-003 suspend  211K   133K arc.size
...
```

5. Enter `resume` to resume all datasets.

```
hostname:analytics datasets> resume
```

6. Enter `show` to view a list of active datasets.

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE ONDISK NAME
dataset-000 active    642K   588K arc.accesses[hit/miss]
dataset-001 active    215K   174K arc.l2_accesses[hit/miss]
dataset-002 active    215K   134K arc.l2_size
dataset-003 active    215K   134K arc.size
...
```

▼ Discarding Data in a Dataset (BUI)

Use the following procedure to discard an entire archived dataset, or a portion of an archived dataset. Datasets are automatically pruned according to the existing retention policy, but they can also be pruned manually. This process may take several minutes to complete, depending on the size of the dataset and the amount of pruning required. If you discard an entire dataset, the ON DISK size is zero. If you discard only a portion of the dataset, the ON DISK size is reduced. Note that datasets can only be pruned when they are active.

1. **Go to Analytics > Dataset.**
2. **Hover over a dataset and click the discard icon .**
3. **Select one of the following options:**
 - a. **To discard the entire dataset, select Entire dataset and skip to step 5.**
 - b. **To discard only a portion of the dataset, uncheck Entire dataset, select a data granularity of second, minute, or hour, and continue to step 4.**
4. **Select one of the following options:**
 - a. **To discard all data from the selected granularity, select All.**
 - b. **To discard data after a certain period of time, select Older than, type an integer value into the text box, and select one of the following time periods: hours, days, weeks, months.**
5. **Click OK.**

▼ Discarding Data in a Dataset (CLI)

Use the following procedure to discard an entire archived dataset, or a portion of an archived dataset. Datasets are automatically pruned according to the existing retention policy, but they can also be pruned manually. This process may take several minutes to complete, depending on the size of the dataset and the amount of pruning required. Note that datasets can only be pruned when they are active.

1. **Go to analytics datasets.**

2. Enter `show` to view a list of active datasets.

```
hostname:analytics datasets> show
Datasets:

DATASET      STATE   INCORE  ONDISK  NAME
dataset-000  active  1.27M   15.5M   arc.accesses[hit/miss]
dataset-001  active   517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
eligibility]
...
dataset-005  active   290K    7.80M   cpu.utilization
```

3. Select one of the following options:

- a. To discard an entire archived dataset, enter `destroy` and the dataset. Then enter `Y` to confirm your action.

```
hostname:analytics datasets> destroy dataset-005
This will destroy "dataset-005". Are you sure? (Y/N) Y
```

- b. To discard only a portion of an archived dataset, enter `select` and the name of the archived dataset that you want to prune. Then enter `prune`, with any of these options: `date`, `time`, and `granularity`.

`date` A date before which all data is removed. If a date is not specified, then the default is the present date and time. Enter the date using this format: year-month-date.

`time` The time before which all data is removed. If a time is not specified, then the default is 12:00 AM, or 00:00 in 24-hour notation. Enter the time in 24-hour notation, using this format: hour:minute:second.

`granularity` The level of data that is removed. The granularity can be presented as one of the following: second, minute, or hour.

If minute or hour is specified, the lower level of data granularity is also deleted. For example, using the `prune hour` command also deletes the per-second and per-minute data.

Detailed datasets can be pruned on several levels to reduce the quantity of saved data, which allows you to archive only a portion of the dataset. For example, use a series of `prune` commands to keep a day's worth of per-second data, a couple weeks of per-minute data, and six months of per-hour data. See Example 4.

Example 1 Pruning by granularity

In the following example, only the granularity is specified. This example discards all per-second and per-minute data collected before 4:56 PM on April 2, 2012.

```
hostname:analytics datasets> select dataset-001
hostname:analytics dataset-001> prune minute
This will remove per-second and minute data collected prior to 2012-4-02
16:56:52.

Are you sure? (Y/N) Y
```

Example 2 Pruning by date

In the following example, only a date is specified. This example discards all per-second data collected before midnight on December 1, 2015.

```
hostname:analytics dataset-001> prune 2015-12-01 second
This will remove per-second data collected prior to 2015-12-1 00:00.

Are you sure? (Y/N) Y
```

Example 3 Pruning by date and time

In the following example, both a date and time are specified. This example discards per-second data collected at and before 12 PM on June 3, 2015.

```
hostname:analytics dataset-001> prune 2015-06-03 12:00:01 second
This will remove per-second data collected prior to 2015-6-3 12:00:01.

Are you sure? (Y/N) Y
```

Example 4 Pruning detailed datasets

The following prune commands keep a day's worth of per-second data, a couple weeks of per-minute data, and six months of per-hour data older than December 15, 2015.

```
hostname:analytics dataset-001> prune 2015-12-14 second
hostname:analytics dataset-001> prune 2015-12-01 minute
hostname:analytics dataset-001> prune 2015-6-01 hour
```

▼ Identifying CPU Performance Issues (BUI)

Use the following procedure to identify and remedy CPU hardware bottlenecks on the appliance. Based on the results of two analytic datasets, suggested corrective actions are provided to increase data throughput.

1. **Create a worksheet as described in “[Creating a Worksheet \(BUI\)](#)” on page 15.**
2. **Click the add icon  next to Add statistic.**
3. **Go to CPU > Percent utilization > As a raw statistic.**
4. **Click the add icon  again.**
5. **Go to CPU > Percent utilization > Broken down by CPU identifier.**
6. **Wait at least 15 minutes.**

Note - Fifteen minutes is a general guideline. This amount of time may be adjusted if you have frequent, short-duration workloads that are CPU intensive.

7. **Examine the graph of CPUs broken down by percent utilization.**
If the appliance CPUs reach 100% utilization for more than 15 minutes, you should consider adding more CPUs or upgrading to faster CPUs.
8. **Examine the graph of CPU percent utilization broken down by CPU identifier.**
A single CPU core operating at 100% utilization while the others are relatively idle is a likely indication of a single-threaded and/or single-client workload. Consider dividing your workload among multiple clients, or investigate a multi-threaded implementation of your client application to better utilize the many CPU cores offered by other controller models.

▼ Identifying CPU Performance Issues (CLI)

Use the following procedure to identify and remedy CPU hardware bottlenecks on the appliance. Based on the results of two analytic datasets, suggested corrective actions are provided to increase data throughput.

1. **Create a worksheet as described in “[Creating a Worksheet \(CLI\)](#)” on page 15, select that worksheet, and then enter dataset.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Enter `set name=cpu.utilization`, and then enter `commit` to add CPU percent utilization as a raw statistic to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name=cpu.utilization
name = cpu.utilization
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter `dataset`.**

```
hostname:analytics worksheet-000> dataset
```

4. **Enter `set name=cpu.utilization[cpu]`, and then enter `commit` to add CPU percent utilization broken down by CPU identifier to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name=cpu.utilization[cpu]
name = cpu.utilization[cpu]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

5. **Enter `done`, and then enter `done` again to exit the context.**

```
hostname:analytics worksheet-000> done
hostname:analytics worksheets> done
```

6. **Wait at least 15 minutes, and then go to `analytics datasets`.**

Note - Fifteen minutes is a general guideline. This amount of time may be adjusted if you have frequent, short-duration workloads that are CPU intensive.

```
hostname:> analytics datasets
```

7. **Enter `show` to view a list of available datasets.**

```
hostname:analytics datasets> show
Datasets:
```

DATASET	STATE	INCORE	ONDISK	NAME
dataset-000	active	1.27M	15.5M	arc.accesses[hit/miss]
dataset-001	active	517K	9.21M	arc.accesses[hit/miss=metadata hits][L2ARC eligibility]
...				
dataset-005	active	290K	7.80M	cpu.utilization

```
hostname:analytics datasets>
```

8. **Enter `select` and the dataset with the name `cpu.utilization`.**

In this example, dataset name `cpu.utilization` corresponds to `dataset-005`.

```
hostname:analytics datasets> select dataset-005
```

9. **Enter read 900 to read the last 900 seconds, or 15 minutes, of the dataset. When you are finished examining the data, enter done.**

If the appliance CPUs reach 100% utilization for more than 15 minutes, you should consider adding more CPUs or upgrading to faster CPUs.

```
hostname:analytics dataset-005> read 900
...
hostname:analytics dataset-005> done
```

10. **Enter show to view a list of available datasets.**

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active     517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
  eligibility]
...
dataset-006 active     290K    7.80M   cpu.utilization[cpu]
hostname:analytics datasets>
```

11. **Enter select and the dataset with the name cpu.utilization[cpu].**

In this example, dataset name `cpu.utilization[cpu]` corresponds to `dataset-006`.

```
hostname:analytics datasets> select dataset-006
```

12. **Enter read 900 to read the last 900 seconds, or 15 minutes, of the dataset. When you are finished examining the data, enter done.**

```
hostname:analytics dataset-006> read 900
...
hostname:analytics dataset-006> done
```

A single CPU core operating at 100% utilization while the others are relatively idle is a likely indication of a single-threaded and/or single-client workload. Consider dividing your workload among multiple clients, or investigate a multi-threaded implementation of your client application to better utilize the many CPU cores offered by other controller models.

▼ Identifying Network Performance Issues (BUI)

Use the following procedure to identify and remedy network hardware bottlenecks on the appliance. Based on the results of the analytic dataset, suggested corrective actions are provided to increase network throughput.

1. **Create a worksheet as described in “[Creating a Worksheet \(BUI\)](#)” on page 15.**
2. **Click the add icon  next to Add statistic.**
3. **Go to NETWORK > Device bytes > Broken down by device.**
4. **Wait at least 10 minutes.**

Note - Ten minutes is a general guideline. This amount of time may be adjusted if you have shorter-duration workloads that demand maximum available network bandwidth.

5. **Examine the graph.**
If any network device reaches 95% of its maximum throughput for more than 10 minutes, you might need to install additional network devices.

▼ Identifying Network Performance Issues (CLI)

Use the following procedure to identify and remedy network hardware bottlenecks on the appliance. Based on the results of the analytic dataset, suggested corrective actions are provided to increase network throughput.

1. **Create a worksheet as described in “[Creating a Worksheet \(CLI\)](#)” on page 15, select that worksheet, and then enter `dataset`.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Enter `set name=nic.kilobytes[device]`, and then enter `commit` to add Network device bytes broken down by device to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name=nic.kilobytes[device]
name = nic.kilobytes[device]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter `done`, and then enter `done` again to exit the context.**

```
hostname:analytics worksheet-000> done
hostname:analytics worksheets> done
```

4. **Wait at least 10 minutes, and then go to `analytics datasets`.**

Note - Ten minutes is a general guideline. This amount of time may be adjusted if you have shorter-duration workloads that demand maximum available network bandwidth.

```
hostname:> analytics datasets
```

5. Enter show to view a list of available datasets.

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active     517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
eligibility]
...
dataset-032 active    290K    7.80M   nic.kilobytes[device]
hostname:analytics datasets>
```

6. Enter select and the dataset with the name nic.kilobytes[device].

In this example, dataset name `nic.kilobytes[device]` corresponds to `dataset-032`.

```
hostname:analytics datasets> select dataset-032
```

7. Enter read 600 to read the last 600 seconds, or 10 minutes, of the dataset. When you are finished examining the data, enter done.

```
hostname:analytics dataset-032> read 600
...
hostname:analytics dataset-032> done
```

If any network device reaches 95% of its maximum throughput for more than 10 minutes, you might need to install additional network devices.

▼ Identifying Memory Performance Issues (BUI)

Use the following procedure to identify and remedy memory hardware bottlenecks on the appliance. Based on the results of the analytic dataset, suggested corrective actions are provided to increase memory performance by installing more DRAM.

1. **Create a worksheet as described in “[Creating a Worksheet \(BUI\)](#)” on page 15.**
2. **Click the add icon  next to Add statistic.**
3. **Go to CACHE > ARC accesses > Broken down by hit/miss.**
4. **Wait at least 10 minutes.**

Note - Ten minutes is a general guideline. The amount of time may be adjusted if you have shorter-duration workloads that are memory intensive.

5. Examine the graph.

You may want to install more DRAM when all of the conditions in the following table are present.

Condition	Description
ARC access hits for data or metadata are at least 75-97% compared to misses	The ARC is providing a benefit by storing data or metadata that the applications need.
ARC access hits for data or metadata are significantly greater than prefetch hits	The majority of the ARC accesses are for real applications rather than just the prefetch mechanism.
ARC is accessed at least 10,000 times per second	The appliance is hitting DRAM, which is not the typical utilization of an idle system.
Nearly all memory is consumed by the ARC, leaving very little unused memory	The appliance is utilizing all of the DRAM possible for the ARC, not just serving a hot workload out of a small subset of the DRAM that is already present.

▼ Identifying Memory Performance Issues (CLI)

Use the following procedure to identify and remedy memory hardware bottlenecks on the appliance. Based on the results of the analytic dataset, suggested corrective actions are provided to increase memory performance by installing more DRAM.

1. **Create a worksheet as described in “[Creating a Worksheet \(CLI\)](#)” on page 15, select that worksheet, and then enter `dataset`.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Enter `set name=arc.accesses[hit/miss]`, and then enter `commit` to add cache ARC accesses broken down by hit/miss to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name=arc.accesses[hit/miss]
                               name = arc.accesses[hit/miss]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter `done`, and then enter `done` again to exit the context.**

```
hostname:analytics worksheet-000> done
```

```
hostname:analytics worksheets> done
```

4. Wait at least 10 minutes, and then go to analytics datasets.

Note - Ten minutes is a general guideline. The amount of time may be adjusted if you have shorter-duration workloads that are memory intensive.

```
hostname:> analytics datasets
```

5. Enter show to view a list of available datasets.

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active     517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
eligibility]
...
hostname:analytics datasets>
```

6. Enter select and the dataset with the name arc.accesses[hit/miss].

In this example, dataset name arc.accesses[hit/miss] corresponds to dataset-000.

```
hostname:analytics datasets> select dataset-000
```

7. Enter read 600 to read the last 600 seconds, or 10 minutes, of the dataset.

```
hostname:analytics dataset-000> read 600
```

8. Examine the data.

You may want to install more DRAM when all of the conditions in the following table are present.

Condition	Description
ARC access hits for data or metadata are at least 75-97% compared to misses	The ARC is providing a benefit by storing data or metadata that the applications need.
ARC access hits for data or metadata are significantly greater than prefetch hits	The majority of the ARC accesses are for real applications rather than just the prefetch mechanism.
ARC is accessed at least 10,000 times per second	The appliance is hitting DRAM, which is not the typical utilization of an idle system.
Nearly all memory is consumed by the ARC, leaving very little unused memory	The appliance is utilizing all of the DRAM possible for the ARC, not just serving a hot workload out of a small subset of the DRAM that is already present.

▼ When to Add the First Read Cache Device (BUI)

Use the following procedure to determine if you need a first read cache device for the appliance. To determine if you need more than one device, see [“When to Add More Read Cache Devices \(BUI\)” on page 42.](#)

1. **Create a worksheet as described in [“Creating a Worksheet \(BUI\)” on page 15.](#)**
2. **Click the add icon  next to Add statistic.**
3. **Go to CACHE > ARC accesses > Broken down by hit/miss.**
4. **Click metadata hits under the Range average column to show the graph of ARC accesses per second, broken down by hit/miss.**
5. **Click the drilldown icon  and select By L2ARC eligibility.**
6. **Click data hits to show the graph of ARC accesses per second, broken down by hit/miss.**
7. **Click the drilldown icon  and select By L2ARC eligibility.**
8. **Wait several minutes.**

Note - This wait time can be adjusted to better determine your peak I/O. Capturing analytics for 24 hours during normal business operations may reveal the best understanding of your IO patterns.

9. **Examine the graphs.**
Consider adding the first read cache device when all of the following conditions are present:
 - There are at least 1500 L2ARC-eligible ARC access misses for data and/or metadata per second
 - The appliance has an active filesystem or LUN with a ZFS record size of 32k or smaller

▼ When to Add the First Read Cache Device (CLI)

Use the following procedure to determine if you need a first read cache device for the appliance. To determine if you need more than one device, see [“When to Add More Read Cache Devices \(CLI\)” on page 43.](#)

1. **Create a worksheet as described in “Creating a Worksheet (CLI)” on page 15, select that worksheet, and then enter dataset.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Enter set name=arc.accesses[hit/miss], and then enter commit to add cache ARC access broken down by hit/miss to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name=arc.accesses[hit/miss]
name = arc.accesses[hit/miss]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter dataset.**

```
hostname:analytics worksheet-000> dataset
```

4. **Repeat steps 2 and 3 to add the following datasets:**

- Cache ARC accesses broken down by metadata hits and misses ("arc.accesses[hit/miss=metadata hits][L2ARC eligibility]")
- Cache ARC accesses broken down by data hits and misses ("arc.accesses[hit/miss=data hits][L2ARC eligibility]")

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name="arc.accesses[hit/miss=metadata hits][L2ARC eligibility]"
name = arc.accesses[hit/miss=metadata hits][L2ARC eligibility]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
hostname:analytics worksheet-000> dataset
hostname:analytics worksheet-000 dataset (uncommitted)> set name="arc.accesses[hit/miss=data hits][L2ARC eligibility]"
name = arc.accesses[hit/miss=data hits][L2ARC eligibility]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

5. **Enter done, and then enter done again to exit the context.**

```
hostname:analytics worksheet-000> done
hostname:analytics worksheets> done
```

6. **Wait several minutes, and then go to analytics datasets.**

Note - This wait time can be adjusted to better determine your peak I/O. Capturing analytics for 24 hours during normal business operations may reveal the best understanding of your IO patterns.

```
hostname:> analytics datasets
```

7. Enter show to view a list of available datasets.

```
hostname:analytics datasets> show
Datasets:
```

```

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active     517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
  eligibility]
...
hostname:analytics datasets>
```

8. Enter select and the dataset with the name arc.accesses[hit/miss=metadata hits][L2ARC eligibility].

In this example, dataset name arc.accesses[hit/miss=metadata hits][L2ARC eligibility] corresponds to dataset-001.

```
hostname:analytics datasets> select dataset-001
```

9. Enter read 86400 to read the last 86,400 seconds, or 24 hours, of the dataset. When you are finished examining the data, enter done.

```
hostname:analytics dataset-001> read 86400
...
hostname:analytics dataset-001> done
```

Consider adding the first read cache device when all of the following conditions are present:

- There are at least 1500 L2ARC-eligible ARC access misses for metadata per second
- The appliance has an active filesystem or LUN with a ZFS recordsize of 32k or smaller

10. Enter show to view a list of available datasets.

```
hostname:analytics datasets> show
Datasets:
```

```

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active     517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
  eligibility]
dataset-002 active     780K    9.20M   arc.accesses[hit/miss=data hits][L2ARC eligibility]
...
hostname:analytics datasets>
```

11. **Enter `select` and the dataset with the name `arc.accesses[hit/miss=data hits][L2ARC eligibility]`.**

In this example, dataset name `arc.accesses[hit/miss=data hits][L2ARC eligibility]` corresponds to `dataset-002`.

```
hostname:analytics datasets> select dataset-002
```

12. **Enter `read 86400` to read the last 86,400 seconds, or 24 hours, of the dataset. When you are finished examining the data, enter `done`.**

```
hostname:analytics dataset-002> read 86400
...
hostname:analytics dataset-002> done
```

Consider adding the first read cache device when all of the following conditions are present:

- There are at least 1500 L2ARC-eligible ARC access misses for data per second
- The appliance has an active filesystem or LUN with a ZFS recordsize of 32k or smaller

▼ When to Add More Read Cache Devices (BUI)

Use the following procedure to determine if you need more read cache devices for the appliance.

Before You Begin Go to Maintenance > Hardware to identify which chassis and slots contain read cache devices.

1. **Create a worksheet as described in “[Creating a Worksheet \(BUI\)](#)” on page 15.**
2. **Click the add icon  next to Add statistic.**
3. **Go to DISK > I/O operations > Broken down by disk.**
4. **Select a read cache device.**
5. **Click the drilldown icon  and select Percent utilization.**
6. **Repeat steps 4 and 5 for all existing read cache devices.**
7. **Wait at least 10 minutes.**
8. **Examine the graph(s).**

You may want to add more read cache devices when existing devices are 90% utilized.

▼ When to Add More Read Cache Devices (CLI)

Use the following procedure to determine if you need more read cache devices for the appliance.

Before You Begin Go to `maintenance hardware show` to identify which chassis and slots contain read cache devices.

1. **Create a worksheet as described in “[Creating a Worksheet \(CLI\)](#)” on page 15, select that worksheet, and then enter `dataset`.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Add a read cache device to the worksheet, broken down by percent utilization, and then enter `commit`. Repeat for each read cache device used by the appliance.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name="io.utilization
[disk=hostname/HDD 13]"
                               name = io.utilization[disk=hostname/HDD 13]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter `done`, and then enter `done` again to exit the context.**

```
hostname:analytics worksheet-000> done
hostname:analytics worksheets> done
```

4. **Wait at least 10 minutes, and then go to `analytics datasets`.**

```
hostname:> analytics datasets
```

5. **Enter `show` to view a list of available datasets.**

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active     517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
eligibility]
...
dataset-021 active     290K    7.80M   io.utilization[disk=hostname/HDD 13]
hostname:analytics datasets>
```

6. **Enter `select` and the first dataset that you added.**

In this example, dataset name `io.utilization[disk=hostname/HDD 13]` corresponds to `dataset-021`.

```
hostname:analytics datasets> select dataset-021
```

7. **Enter read 600 to read the last 600 seconds, or 10 minutes, of the dataset.**

```
hostname:analytics dataset-021> read 600
```

8. **Repeat steps 6 and 7 for each dataset that you added to the worksheet.**

9. **Examine the data.**

You may want to add more read cache devices when existing devices are 90% utilized.

▼ When to Add the First Write Log Device (BUI)

Use this procedure to determine if you need a first write log device for your appliance. Fibre Channel and iSCSI writes are synchronous and benefit from log devices when the write cache is enabled. Be sure to thoroughly read and understand the documented implications of enabling write caching on LUNs before proceeding. For more information, see [“Space Management for Shares” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#).

To determine if you need to add more than one write log device, see [“When to Add More Write Log Devices \(BUI\)” on page 48](#).

1. **Create a worksheet as described in [“Creating a Worksheet \(BUI\)” on page 15](#).**
2. **Click the add icon  next to Add statistic.**
3. **Go to PROTOCOL > NFSv2 operations > Broken down by type of operation.**
4. **Repeat steps 2 and 3 for each protocol that is configured for your appliance.**
5. **Wait at least 15 minutes.**

Note - Fifteen minutes is a general guideline. This amount of time may be adjusted if you have a performance-sensitive, short-duration synchronous write workload.

6. **Examine the graph(s).**

Consider adding a write log device when one or all of the following conditions are present:

- The sum of iSCSI writes, Fibre Channel writes, and NFS/SMB synchronous operations is at least 1000 per second
- There are at least 100 NFS commits per second

▼ When to Add the First Write Log Device (CLI)

Use this procedure to determine if you need a first write log device for your appliance. Fibre Channel and iSCSI writes are synchronous and benefit from log devices when the write cache is enabled. Be sure to thoroughly read and understand the documented implications of enabling write caching on LUNs before proceeding. For more information, see [“Space Management for Shares” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0.](#)

To determine if you need to add more than one write log device, see [“When to Add More Write Log Devices \(CLI\)” on page 48.](#)

1. **Create a worksheet as described in [“Creating a Worksheet \(CLI\)” on page 15](#), select that worksheet, and then enter `dataset`.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Enter `set name=nfs2.ops[op]`, and then enter `commit` to add NFSv2 operations per second broken down by type of operation to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name=nfs2.ops[op]
                               name = nfs2.ops[op]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter `dataset`.**

4. **Repeat steps 2 and 3 to add the following datasets:**

- NFSv3 operations per second, broken down by type of operation (`nfs3.ops[op]`)
- NFSv4 operations per second, broken down by type of operation (`nfs4.ops[op]`)
- NFSv4.1 operations per second, broken down by type of operation (`nfs4-1.ops[op]`)
- iSCSI operations per second, broken down by type of operation (`iscsi.ops[op]`)
- Fibre Channel operations per second, broken down by type of operation (`fc.ops[op]`)
- SMB operations per second, broken down by type of operation (`smb.ops[op]`)

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name=nfs3.ops[op]
                               name = nfs3.ops[op]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

```
hostname:analytics worksheet-000> dataset
hostname:analytics worksheet-000 dataset (uncommitted)> set name=nfs4.ops[op]
name = nfs4.ops[op]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
hostname:analytics worksheet-000> dataset
hostname:analytics worksheet-000 dataset (uncommitted)> set name=nfs4-1.ops[op]
name = nfs4-1.ops[op]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
hostname:analytics worksheet-000> dataset
hostname:analytics worksheet-000 dataset (uncommitted)> set name=iscsi.ops[op]
name = iscsi.ops[op]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
hostname:analytics worksheet-000> dataset
hostname:analytics worksheet-000 dataset (uncommitted)> set name=fc.ops[op]
name = fc.ops[op]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
hostname:analytics worksheet-000> dataset
hostname:analytics worksheet-000 dataset (uncommitted)> set name=smb.ops[op]
name = smb.ops[op]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

5. Enter done, and then enter done again to exit the context.

```
hostname:analytics worksheet-000> done
hostname:analytics worksheets> done
```

6. Wait at least 15 minutes, and then go to analytics datasets.

Note - Fifteen minutes is a general guideline. This amount of time may be adjusted if you have a performance-sensitive, short-duration synchronous write workload.

```
hostname:> analytics datasets
```

7. Enter show to view a list of available datasets.

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active   1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active    517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
eligibility]
...
dataset-030 active    290K    7.80M   nfs2.ops[op]
hostname:analytics datasets>
```

8. Enter select and the dataset with the name nfs2.ops[op].

In this example, dataset name `nfs2.ops[op]` corresponds to `dataset-030`.

```
hostname:analytics datasets> select dataset-030
```

9. **Enter read 900 to read the last 900 seconds, or 15 minutes, of the dataset, and copy and save the data to you environment for future reference.**

```
hostname:analytics dataset-030> read 900
```

10. **Enter done.**

```
hostname:analytics dataset-030> done
```

11. **Repeat steps 7 through 10 for the following datasets:**

- NFSv3 operations per second, broken down by type of operation (`nfs3.ops[op]`)
- NFSv4 operations per second, broken down by type of operation (`nfs4.ops[op]`)
- NFSv4.1 operations per second, broken down by type of operation (`nfs4-1.ops[op]`)
- iSCSI operations per second, broken down by type of operation (`iscsi.ops[op]`)
- Fibre Channel operations per second, broken down by type of operation (`fc.ops[op]`)
- SMB operations per second, broken down by type of operation (`smb.ops[op]`)

Note - Remember to copy and save the data for each dataset to your environment for future reference.

```
hostname:analytics datasets> show
...
hostname:analytics datasets> select dataset-032
hostname:analytics dataset-032> read 900
...
hostname:analytics dataset-032> done
hostname:analytics datasets> show
...
hostname:analytics datasets> select dataset-034
...
hostname:analytics datasets> select dataset-27
...
hostname:analytics datasets> select dataset-13
...
hostname:analytics datasets> select dataset-07
...
hostname:analytics datasets> select dataset-40
```

12. **Examine the data.**

You may want to add the first write log device when one or all of the following conditions are present:

- The sum of iSCSI writes, Fibre Channel writes, and NFS/SMB synchronous operations is at least 1000 per second
- There are at least 100 NFS commits per second

▼ When to Add More Write Log Devices (BUI)

Use the following procedure to determine if you need more write log devices for the appliance.

Before You Begin Go to Maintenance > Hardware to identify which chassis and slots contain write cache devices.

1. **Create a worksheet as described in “[Creating a Worksheet \(BUI\)](#)” on page 15.**
2. **Click the add icon  next to Add statistic.**
3. **Go to DISK > Disks > Broken down by percent utilization.**
4. **Select a write log device.**
5. **Click the drilldown icon  and select As a raw statistic.**
6. **Repeat steps 4 and 5 for all existing write log devices.**
7. **Wait at least 10 minutes.**

Note - Ten minutes is a general guideline. This amount of time may be adjusted if you have a performance-sensitive, short-duration synchronous write workload.

8. **Examine the graph(s).**

You may want to add more write log devices when existing devices are 90% utilized.

▼ When to Add More Write Log Devices (CLI)

Use the following procedure to determine if you need more write log devices for the appliance.

1. **Create a worksheet as described in “[Creating a Worksheet \(CLI\)](#)” on page 15, select that worksheet, and then enter dataset.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Enter set name="io.disks[utilization=90]", and then enter commit to add disks with utilization of at least 90% to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name="io.disks
[utilization=90]"
                                name = io.disks[utilization=90]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter done, and then enter done again to exit the context.**

```
hostname:analytics worksheet-000> done
hostname:analytics worksheets> done
```

4. **Wait at least 10 minutes, and then go to analytics datasets.**

Note - Ten minutes is a general guideline. This amount of time may be adjusted if you have a performance-sensitive, short-duration synchronous write workload.

```
hostname:> analytics datasets
```

5. **Enter show to view a list of available datasets.**

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active     517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
eligibility]
...
dataset-019 active     290K    7.80M   io.disks[utilization=90]
hostname:analytics datasets>
```

6. **Enter select and the dataset with the name io.disks[utilization=90].**
In this example, dataset name io.disks[utilization=90] corresponds to dataset-019.

```
hostname:analytics datasets> select dataset-019
```

7. **Enter read 600 to read the last 600 seconds, or 10 minutes, of the dataset.**

```
hostname:analytics dataset-019> read 600
```

8. **Examine the data.**

You may want to add more write log devices when existing devices are 90% utilized.

▼ When to Add More Disks (BUI)

Use this procedure to determine if you need to add more disks. Note that disks can be over-utilized by making a poor choice for the RAID profile and/or ZFS record size. In this case, it may be possible to reduce existing disk utilization by moving from RAIDZ to mirrored profiles and/or matching ZFS record sizes to client I/O sizes.

In addition, if the system is configured with no read- or write-optimized flash drives, then all of the I/O operations beyond DRAM are served by disks. For better performance, consider using flash for any workload that includes random reads or synchronous writes.

1. **Create a worksheet as described in “[Creating a Worksheet \(BUI\)](#)” on page 15.**
2. **Click the add icon  next to Add statistic.**
3. **Go to DISK > Disks > Broken down by percent utilization.**
4. **Right-click on the 70% range average and click By disk.**
5. **Wait at least 30 minutes.**

Note - Thirty minutes is a general guideline. If you have short-duration workloads that are bottlenecked on disk utilization, you may want to consider selecting additional disks at less than 30 minutes.

6. **Examine the graphs.**

You may want to use more disks when at least 50% of existing disks are at least 70% utilized.

▼ When to Add More Disks (CLI)

Use this procedure to determine if you need to add more disks. Note that disks can be over-utilized by making a poor choice for the RAID profile and/or ZFS record size. In this case, it may be possible to reduce existing disk utilization by moving from RAIDZ to mirrored profiles and/or matching ZFS record sizes to client I/O sizes.

In addition, if the system is configured with no read- or write-optimized flash drives, then all of the I/O operations beyond DRAM are served by disks. For better performance, consider using flash for any workload that includes random reads or synchronous writes.

1. **Create a worksheet as described in “Creating a Worksheet (CLI)” on page 15, select that worksheet, and then enter dataset.**

```
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheet-000> dataset
```

2. **Enter set name="io.disks[utilization=70]", and then enter commit to add disks with utilization of at least 70% to your worksheet.**

```
hostname:analytics worksheet-000 dataset (uncommitted)> set name="io.disks
[utilization=70]"
                               name = io.disks[utilization=70]
hostname:analytics worksheet-000 dataset (uncommitted)> commit
```

3. **Enter done, and then enter done again to exit the context.**

```
hostname:analytics worksheet-000> done
hostname:analytics worksheets> done
```

4. **Wait at least 30 minutes, and then go to analytics datasets.**

Note - Thirty minutes is a general guideline. If you have short-duration workloads that are bottlenecked on disk utilization, you may want to consider selecting additional disks at less than 30 minutes.

```
hostname:> analytics datasets
```

5. **Enter show to view a list of available datasets.**

```
hostname:analytics datasets> show
Datasets:

DATASET    STATE    INCORE  ONDISK  NAME
dataset-000 active    1.27M   15.5M   arc.accesses[hit/miss]
dataset-001 active    517K    9.21M   arc.accesses[hit/miss=metadata hits][L2ARC
eligibility]
...
dataset-025 active    290K    7.80M   io.disks[utilization=70]
hostname:analytics datasets>
```

6. **Enter select and the dataset with the name io.disks[utilization=70].**

In this example, dataset name `io.disks[utilization=70]` corresponds to `dataset-025`.

```
hostname:analytics datasets> select dataset-025
```

7. **Enter read 1800 to read the last 1800 seconds, or 30 minutes, of the dataset.**

```
hostname:analytics dataset-025> read 1800
```

8. Examine the data.

You may want to use more disks when at least 50% of existing disks are at least 70% utilized.

▼ **Configuring a Threshold Alert (BUI)**

Use this procedure to be automatically notified when a specific dataset exceeds or falls below a set threshold level.

If the automatic suspend policy is enabled, the dataset and the corresponding statistics will be automatically suspended when the maximum idle time you set for the threshold is exceeded. To improve the system performance, you can enable this policy and set a shorter time period for the configurable threshold. Note that the automatic suspend policy does not apply to dashboard datasets and datasets used to collect trendable statistics.

- 1. Go to Configuration > Alerts.**
- 2. Click Threshold alerts, and then click its add icon .**
- 3. For Threshold, select a dataset, when the alert should be generated, and type the percentage value.**
- 4. For Timing, type a value and select an interval.**
- 5. If you want to be notified of events only within a certain time period, or only on certain days, select the appropriate check boxes and values.**
- 6. If you want to be notified of the same event continuously while the condition exists, select the appropriate repost check boxes, enter values, and select intervals.**
- 7. For Alert actions, select an action and complete any accompanying fields.**
- 8. (Optional) Click TEST to test the alert action.**
- 9. Click APPLY.**

▼ **Configuring a Threshold Alert (CLI)**

- 1. Go to configuration alerts thresholds.**

2. **Enter create.**

```
hostname:configuration alerts thresholds> create
```

3. **Enter show to view a list of threshold alert properties.**

4. **Set the properties.**

This example sets a capacity threshold alert for the system pool when it exceeds 80 percent for at least one hour on any day of the week. The alert reposts every day while the condition persists, and an alert is also posted when the condition clears for at least one day.

```
hostname:configuration alerts threshold (uncommitted)> set statname=syscap.percentused
statname = syscap.percentused (uncommitted)
hostname:configuration alerts threshold (uncommitted)> set limit=80
limit = 80 (uncommitted)
hostname:configuration alerts threshold (uncommitted)> set minpost=3600
minpost = 1 hours (uncommitted)
hostname:configuration alerts threshold (uncommitted)> set days=all
days=all (uncommitted)
hostname:configuration alerts threshold (uncommitted)> set frequency=86400
frequency = 1 days (uncommitted)
hostname:configuration alerts threshold (uncommitted)> set minclear=86400
minclear = 1 days (uncommitted)
```

5. **Enter commit to complete the task.**

```
hostname:configuration alerts threshold (uncommitted)> commit
```

▼ Exporting a Worksheet (BUI)

1. Create a worksheet containing the datasets of interest, as described in [“Creating a Worksheet \(BUI\)” on page 15](#).
2. Use this worksheet to collect data during a period of time for which further analysis is needed.
3. Click the Synchronize icon  to synchronize the datasets.
4. Click the Pause icon .
5. Click the Zoom in  and Zoom out  icons to view the period of interest.

Note - Only the data in the view will be included in the Analytics support bundle.

6. **Click Save.**
7. **Go to Analytics > Saved Worksheets.**
8. **Hover over the worksheet that you want to export and click the export icon .**

Note - The appliance must be registered for the Phone Home Service before attempting to upload. For more information, see [“Phone Home Configuration” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0.](#)

9. **Enter a Service Request number, which is provided by Oracle support personnel when requesting a worksheet.**
10. **Click APPLY.**
11. **Record the filename of the support bundle and provide it to Oracle support personnel for retrieval.**

▼ Exporting a Worksheet (CLI)

1. **Go to analytics worksheets.**
2. **Enter show to view a list of available worksheets.**

```
hostname:analytics worksheets> show
Worksheets:
```

WORKSHEET	OWNER	NAME
worksheet-000	root	example_1
worksheet-001	root	example_2

3. **Enter select and the worksheet that you want to export.**

```
hostname:analytics worksheets> select worksheet-000
```

4. **Enter bundle and the service request number.**

```
hostname:analytics worksheet-000> bundle 3-7596250401
A support bundle is being created and sent to Oracle. You will receive an alert
when the bundle has finished uploading. Please save the following filename, as
Oracle support personnel will need it in order to access the bundle:
/upload/issue/3-7596250401/3-7596250401_ak.9a4c3d7b-50c5-6eb9-c2a6-ec9808ae1cd8.tar.gz
```

▼ Downloading a Dataset to a CSV File (BUI)

Use the following information to download a dataset to a CSV file. While it is not possible to download a dataset in the CLI, you can view the data in a CSV format. See [“Viewing a Dataset in a CSV Format \(CLI\)” on page 55](#).

1. **Navigate to the open or saved worksheet with the dataset that you want to download.**
2. **Click the Export data icon , located above the dataset that you want to download.**
3. **Save the file locally.**

▼ Viewing a Dataset in a CSV Format (CLI)

Use the following information to view data from a dataset in a CSV format. To download a dataset to a CSV file, see [“Downloading a Dataset to a CSV File \(BUI\)” on page 55](#).

1. **Go to analytics worksheets.**

```
hostname:> analytics worksheets
```

2. **Enter show to view a list of saved worksheets.**

```
hostname:analytics worksheets> show
Worksheets:
```

WORKSHEET	OWNER	NAME
worksheet-000	root	example_1
worksheet-001	root	example_2

3. **Enter select and the worksheet that contains the dataset for which you want to view the data.**

```
hostname:analytics worksheets> select worksheet-000
```

4. **Enter show to view a list of datasets in the worksheet.**

```
hostname:analytics worksheet-000> show
Properties:
```

```
uuid = 66d05260-8d26-4b69-aae5-f98391048af0
name = example_1
```

```
owner = root
ctime = 2016-12-8 23:55:58
mtime = 2016-12-8 23:56:09
```

Datasets:

DATASET	DATE	SECONDS	NAME
dataset-000	-	60	arc.accesses[L2ARC eligibility]
dataset-001	-	60	cap.bytesused[pool]

5. Enter `select` and the dataset for which you want to view the data.

```
hostname:analytics worksheet-000> select dataset-000
```

6. Enter `csv`.

```
hostname:analytics worksheet-000 dataset-000> csv
Time (UTC),Value per second
2016-12-09 00:02:31,40599
2016-12-09 00:02:32,20134
2016-12-09 00:02:33,22425
2016-12-09 00:02:34,3954
2016-12-09 00:02:35,2185
```

Analytics Data Retention Policies

Default Retention Policy

By default, the appliance retains per-second data for 7 days, per-minute data for 14 days, and per-hour data for 90 days. However, it is strongly recommended that you specify a data retention policy that suits your business needs. Retention policies are especially important if you plan to retain large amounts of historical data for an extended period. The maximum retention period is two years. Software updates to OS8.6.0 or later will cap any previous retention policy settings to a maximum of two years. Modifications to the retention policy are recorded in the audit log. This default retention policy is effective with software version OS8.6.0 and later.

Enabling a Retention Policy

A retention policy limits the minimum amount of data collected by a granularity of per-second, per-minute, or per-hour over a period of time, or retention period. You could set one retention policy per granularity. For example, you could define a retention policy to save a minimum of one day of data at the per-second interval, a second policy to save a minimum of one week of data at the per-minute interval, and a third policy to save a minimum of one month of data at the per-hour interval. It is recommended that you only keep the minimum amount of data according to your business requirements, including compliance needs.

Per-second data is the highest granularity and requires more memory and disk space than per-minute or per-hour data. Likewise, setting a longer retention period corresponds to storing more data. Monitor the size of the datasets in the BUI by navigating to Analytics > Datasets, and in the CLI by using the `analytics datasets` context. Adjust retention policies accordingly to meet business requirements while occupying the least amount of space. Retention policies apply to all active datasets; suspended datasets are not affected.

Note that you must increase the retention time with each increase in granularity. For example, you cannot define a retention period of weeks for per-second data and then define a retention period of days for per-minute data.

When you enable a data retention policy, you should assume that old data is immediately removed. For example, if you set a per-second policy for at least three hours, you should assume that any data over three hours old is deleted. In actuality, the appliance periodically deletes old data and may delay removing old data to avoid causing a performance impact.

You can significantly reduce the space used by Analytics by setting a retention policy that periodically discards the highest data granularity.

To enable a retention policy, you must have Super-User privileges or have configured authorization within the Dataset scope.

Viewing Retained Data

Worksheet graphs are displayed in the highest data granularity available to the appliance. For example, if your retention policies do not collect per-second data but do collect per-minute data, the graphs are rendered using per-minute data.

Related Topics

- See [“Configuring Users” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#) for more information on defining authorization scopes for users.
- To learn more about the properties for analytics policies, see [“Data Retention Properties” on page 59](#).

Analytics Data Retention Policies

Default Retention Policy — By default, the appliance retains per-second data for 7 days, per-minute data for 14 days, and per-hour data for 90 days. However, it is strongly recommended that you specify a data retention policy that suits your business needs. Retention policies are especially important if you plan to retain large amounts of historical data for an extended period. The maximum retention period is two years. Software updates to OS8.6.0 or later will cap any previous retention policy settings to a maximum of two years. Modifications to the retention policy are recorded in the audit log. This default retention policy is effective with software version OS8.6.0 and later.

Enabling a Retention Policy — A retention policy limits the minimum amount of data collected by a granularity of per-second, per-minute, or per-hour over a period of time, or retention period. You could set one retention policy per granularity. For example, you could define a retention policy to save a minimum of one day of data at the per-second interval, a second policy to save a minimum of one week of data at the per-minute interval, and a third policy to save a minimum of one month of data at the per-hour interval. It is recommended that you only keep the minimum amount of data according to your business requirements, including compliance needs.

Per-second data is the highest granularity and requires more memory and disk space than per-minute or per-hour data. Likewise, setting a longer retention period corresponds to storing more data. Monitor the size of the datasets in the BUI by navigating to Analytics > Datasets, and in the CLI by using the `analytics datasets` context. Adjust retention policies accordingly to

meet business requirements while occupying the least amount of space. Retention policies apply to all active datasets; suspended datasets are not affected.

Note that you must increase the retention time with each increase in granularity. For example, you cannot define a retention period of weeks for per-second data and then define a retention period of days for per-minute data.

When you enable a data retention policy, you should assume that old data is immediately removed. For example, if you set a per-second policy for at least three hours, you should assume that any data over three hours old is deleted. In actuality, the appliance periodically deletes old data and may delay removing old data to avoid causing a performance impact. You can significantly reduce the space used by Analytics by setting a retention policy that periodically discards the highest data granularity.

To enable a retention policy, you must have Super-User privileges or have configured authorization within the Dataset scope.

Viewing Retained Data — Worksheet graphs are displayed in the highest data granularity available to the appliance. For example, if your retention policies do not collect per-second data but do collect per-minute data, the graphs are rendered using per-minute data.

Related Topics:

- See [“Configuring Users” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#) for more information on defining authorization scopes for users.
- To learn more about the properties for analytics policies, see [“Data Retention Properties” on page 59](#).

Data Retention Properties

For each of the properties below you either select "All" or "At least". If you select "All" you are not defining a retention policy for the data retention interval and the appliance will not limit active datasets. If you select "At least", enter an integer value into the text box. Then select the period for the retention policy: hours, days, weeks or months. These settings apply to all active datasets and should be set according to business requirements, including compliance needs.

TABLE 1 Setting Properties

Property	Description
Per-second data	Use this setting to define the amount of time to retain data recorded at a per-second interval for active datasets.
Per-minute data	Use this setting to define the amount of time to retain data recorded at a per-minute interval for active datasets.

Data Retention Properties

Property	Description
Per-hour data	Use this setting to define the amount of time to retain data recorded at a per-hour interval for active datasets.

Understanding Analytics Worksheets

A worksheet is the BUI screen on which statistics are graphed. Multiple statistics can be plotted at the same time, and worksheets may be assigned a title and saved for future viewing. The act of saving a worksheet will automatically execute the archive action on all open statistics - meaning whatever statistics were open, will continue to be read and archived forever.

For information on how to use worksheets, see [“Managing Worksheets” on page 15](#).

For more information about worksheets, see the following topics:

- [“Worksheet Graphs and Plots” on page 61](#)
- [“Adjusting Graphs” on page 62](#)
- [“Adjusting Quantize Plots” on page 63](#)
- [“Background Patterns” on page 64](#)
- [“Toolbar Reference” on page 64](#)
- [“Worksheet Tips” on page 66](#)
- [“Saved Worksheet Properties” on page 66](#)
- [“BUI Icon Reference” on page 66](#)

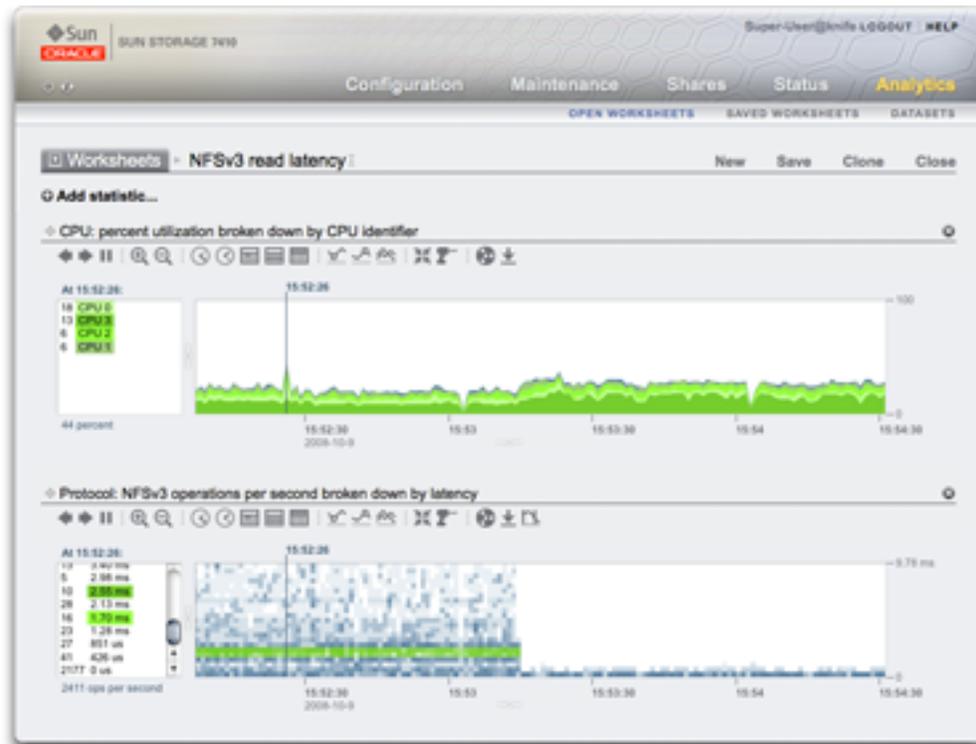
Worksheet Graphs and Plots

Worksheets are the main interface for Analytics. A worksheet is a view where multiple statistics may be graphed. The screenshot that follows shows two statistics:

- CPU: percent utilization broken down by CPU identifier - as a *graph*
- Protocol: NFSv3 operations per second broken down by latency - as a *quantize plot*

Click the screenshot for a larger view. The following sections introduce Analytics features based on that screenshot.

FIGURE 1 Graph showing NFSv3 by Latency



Adjusting Graphs

Graphs provide the following features:

- The left panel lists components of the graph, if available. Since this graph was "... broken down by CPU identifier", the left panel lists CPU identifiers. Only components which had activity in the visible window (or selected time) will be listed on the left.
- Left panel components can be clicked to highlight their data in the main plot window.
- Left panel components can be shift clicked to highlight multiple components at a time (such as in this example, with all four CPU identifiers highlighted).
- Left panel components can be right clicked to show available drilldowns.

- Only ten left panel components will be shown to begin with, followed by "...". You can click the "..." to reveal more. Keep clicking to expand the list completely.
- The graph window on the right can be clicked to highlight a point in time. In the example screenshot, 15:52:26 was selected. Click the pause button followed by the zoom icon to zoom into the selected time. Click the time text to remove the vertical time bar.
- If a point in time is highlighted, the left panel of components will list details for that point in time only. Note that the text above the left box reads "At 15:52:26:", to indicate what the component details are for. If a time wasn't selected, the text would read "Range average:".
- Y-axis auto scales to keep the highest point in the graph (except for utilization statistics, where are fixed at 100%).
- The line graph button  will change this graph to plot just lines without the flood-fill. This may be useful for a couple of reasons: some of the finer detail in line plots can be lost in the flood fill, and so selecting line graphs can improve resolution. This feature can also be used to vertical zoom into component graphs: first, select one or more components on the left, then switch to the line graph.
- To magnify a time range, click where you want the range to begin, shift-click at the end, and click the zoom-in icon .

The following features are common to graphs and quantize plots:

- The height may be expanded. Look for a white line beneath in the middle of the graph, click and drag downwards.
- The width will expand to match the size of your browser.
- Click and drag the move icon  to switch vertical location of the statistics.

Adjusting Quantize Plots

The NFS latency statistic in the screenshot is rendered as a quantize plot. The name refers to the how the data is collected and displayed. For each statistic update, data is quantized into buckets, which are drawn as blocks on the plot. The more events in that bucket for that second, the darker the block will be drawn.

The example screenshot shows NFSv3 operations were spread out to 9 ms and beyond - with latency on the y-axis - until an event kicked in about half way and the latency dropped to less than 1 ms. Other statistics can be plotted to explain the drop in latency (the filesystem cache hit rate showed steady misses go to zero at this point - a workload had been randomly reading from disk (0 to 9+ ms latency), and switched to reading files that were cached in DRAM.)

Quantize plots are used for I/O latency, I/O offset and I/O size, and provide the following features:

- Detailed understanding of data profile (not just the average, maximum or minimum) these visualize all events and promote pattern identification.
- Vertical outlier elimination. Without this, the y-axis would always be compressed to include the highest event. Click the crop outliers icon  to toggle between different percentages of outlier elimination. Hover over this icon to see the current value.
- Vertical zoom: click a low point from the list in the left box, then shift-click a high point. Now click the crop outliers icon to zoom to this range.
- The height may be expanded. Look for a white line beneath in the middle of the plot, click and drag downwards.
- The width will expand to match the size of your browser.
- Click and drag the move icon  to switch vertical location of the statistics.

Background Patterns

Normally graphs are displayed with various colors against a white background. If data is unavailable for any reason the graph will be filled with a pattern to indicate the specific reason for data unavailability:

-  The gray pattern indicates that the given statistic was not being recorded for the time period indicated. This is either because the user had not yet specified the statistic or because data gathering had been explicitly suspended.
-  The red pattern indicates that data gathering was unavailable during that period. This is most commonly seen because the system was down during the time period indicated.
-  The orange pattern indicates an unexpected failure while gathering the given statistic. This can be caused by a number of aberrant conditions. If it is seen persistently or in critical situations, contact your authorized support resource and/or submit a support bundle. For more information on how to submit a support bundle, see [“Working with Support Bundles” in Oracle ZFS Storage Appliance Customer Service Manual](#) .

Toolbar Reference

A toolbar of buttons is shown above graphed statistics. The following is a reference for their function:

TABLE 2 Toolbar Reference

Icon	Click	Shift-Click
	Move backwards in time (moves left)	Move backwards in time (moves left)
	Move forwards in time (moves right)	Move forwards in time (moves right)
	Forward to now	Forward to now
	Pause	Pause
	Zoom out	Zoom out
	Zoom in	Zoom in
	Show one minute	Show two minutes, three, four, ...
	Show one hour	Show two hours, three, four, ...
	Show one day	Show two days, three, four, ...
	Show one week	Show two weeks, three, four, ...
	Show one month	Show two months, three, four, ...
	Show minimum	Show next minimum, next next minimum, ...
	Show maximum	Show next maximum, next next maximum, ...
	Show line graph	Show line graph
	Show mountain graph	Show mountain graph
	Crop outliers	Crop outliers
	Sync worksheet to this statistic	Sync worksheet to this statistic
	Unsync worksheet statistics	Unsync worksheet statistics
	Drilldown	Rainbow highlight
	Save statistical data	Save statistical data
	Export statistical data	Export statistical data

Hover over each button to see a tooltip to describe the click behavior.

Worksheet Tips

- If you want to save a worksheet that displays an interesting event, make sure the statistics are paused first (sync all statistics, then hit pause). Otherwise the graphs will continue to scroll, and when you open the worksheet later the event may no longer be on the screen.
- If you analyze issues after the fact, you will be restricted to the datasets that were already being archived. Visual correlations can be made between them when the time axis is synchronized. If the same pattern is visible in different statistics, there is a good chance that it is related activity.
- Be patient when zooming out to the month view and longer. Analytics is clever about managing long periods of data, however there can still be delays when zooming out to long periods.
- When a worksheet is saved on one node in a clustered system, a copy of the worksheet with the same title is propagated to the peer. To permanently save the worksheet statistics on the peer, click Save.

Saved Worksheet Properties

The following properties are stored for saved worksheets:

TABLE 3 Properties for Saved Worksheets

Field	Description
Name	Configurable name of the saved worksheet. This will be displayed at the top of the Open Worksheets view
Comment	Optional comment (only visible in the BUI)
Owner	User who owns the worksheet
Created	Time the worksheet was created
Modified	Time the worksheet was last modified (only visible in the CLI)

BUI Icon Reference

Hover over saved worksheet entries to expose the following controls:

TABLE 4 BUI Icons

Icon	Description
	Upload this worksheet bundle to Oracle support for analysis. The appliance must be registered for the Phone Home Service (see “Phone Home Configuration” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0) before attempting to upload. You will be prompted to enter a Service Request (SR) number, which is provided by Oracle support personnel when requesting a worksheet.
	Append the datasets saved in this worksheet to the current worksheet in Open Worksheets
	Edit the worksheet to change the name and comment
	Destroy this worksheet

About Analytics Datasets

The term *dataset* refers to the in-memory cached and on-disk saved data for a statistic, and is presented as an entity in Analytics with administration controls. Datasets are automatically created whenever you view statistics in Open Worksheets. A dataset is not saved to disk for future viewing unless you archive it. See [“Statistics Actions” on page 77](#).

The Analytics > Datasets screen in the BUI lists all datasets. These include open statistics that are being viewed in a worksheet (and as such are temporary datasets that disappear when the worksheet is closed) and statistics that are being archived to disk.

The following fields are displayed in the Datasets view for all datasets:

TABLE 5 Field Descriptions of Datasets

Field	Description
Status icon	See Table 6, “BUI Icons,” on page 69 .
Overhead	Low, medium or high performance impact of the dataset; see below table. Also see “Storage Performance Impact” on page 72 and “Execution Performance Impact” on page 75 .
Name	Name of the statistic/dataset
Since	First timestamp in dataset. For open statistics, this is the time the statistic was opened, which may be minutes earlier. For archived statistics, this is the first time in the archived dataset which indicates how far back in the past this dataset goes, which may be days, weeks, months. Sorting this column shows the oldest datasets available.
On Disk	Space this dataset consumes on disk
In Core	Space this dataset consumers in main memory

The following icons are visible in the BUI view; some of these are only visible during hover over of a dataset entry:

TABLE 6 BUI Icons

Icon	Description
	Dataset is actively collecting data

Icon	Description
	Dataset is currently suspended from collecting data
	Dataset has a low overhead
	Dataset has a medium overhead
	Dataset has a high overhead
	Suspend/Resume archived datasets
	Enable archiving of this dataset to disk
	Discard all or some data from this dataset

See [“Statistics Actions” on page 77](#) for descriptions for these dataset actions.

Understanding Analytics Statistics

Analytics statistics provide incredible appliance observability, showing how the appliance is behaving and how clients on the network are using it. While the statistics presented by Analytics may appear straightforward, there may be additional details to be aware of when interpreting their meaning. This is especially true for the purposes of performance analysis, where precise understanding of the statistics is often necessary.

To minimize service downtime in clustered systems, statistics and datasets are not available during failback and takeover operations. Data is not collected, and any attempt to suspend or resume statistics is delayed until failback and takeover operations have completed and data collection automatically resumes.

For information about performance impact, actions, and default statistics, see the following topics:

- [“Storage Performance Impact” on page 72](#)
- [“Execution Performance Impact” on page 75](#)
- [“Statistics Actions” on page 77](#)
- [“Default Statistics” on page 78](#)

For information about the Analytics statistics available for monitoring, see the following topics:

- [“CPU: Percent Utilization” on page 79](#)
- [“Cache: ARC Accesses” on page 81](#)
- [“Cache: L2ARC I/O Bytes” on page 84](#)
- [“Cache: L2ARC Accesses” on page 84](#)
- Capacity: Capacity Bytes Used - [BUI](#), [CLI](#)
- Capacity: Capacity Percent Used - [BUI](#), [CLI](#)
- [“Capacity: Meta Device Capacity Bytes Used \(BUI\)” on page 93](#)
- [“Capacity: Meta Device Capacity Percent Used \(BUI\)” on page 94](#)
- [“Capacity: System Pool Bytes Used” on page 94](#)
- [“Capacity: System Pool Percent Used” on page 95](#)
- [“Data Movement: Shadow Migration Bytes” on page 96](#)

- “Data Movement: Shadow Migration Ops” on page 97
- “Data Movement: Shadow Migration Requests” on page 97
- “Data Movement: NDMP Bytes Statistics” on page 98
- “Data Movement: NDMP Operations Statistics” on page 99
- “Data Movement: Replication Bytes” on page 99
- “Data Movement: Replication Operations” on page 100
- “Disk: Disks” on page 101
- “Disk: I/O Bytes” on page 102
- “Disk: I/O Operations” on page 104
- “Network: Device Bytes” on page 106
- “Network: Interface Bytes” on page 107
- “Protocol: SMB/SMB2 Bytes” on page 123
- “Protocol: SMB Operations” on page 108
- “Protocol: Fibre Channel Bytes” on page 109
- “Protocol: Fibre Channel Operations” on page 110
- “Protocol: FTP Bytes” on page 112
- “Protocol: HTTP/WebDAV Requests” on page 113
- “Protocol: iSCSI Bytes” on page 114
- “Protocol: iSCSI Operations” on page 115
- “Protocol: NFSv[2-4] Bytes” on page 116
- “Protocol: NFSv[2-4] Operations” on page 117
- “Protocol: OISP Bytes” on page 119
- “Protocol: OISP Operations” on page 120
- “Protocol: SFTP Bytes” on page 122
- “Protocol: SRP Bytes” on page 124
- “Protocol: SRP Operations” on page 124

Storage Performance Impact

Analytics statistic collection comes at some cost to overall performance. This should not be an issue if you understand what that cost will be, and how to minimize or avoid it.

Analytics statistics can be archived, meaning they will be a dataset that is continually read and saved to the system disks in one second summaries. This allows statistics to be viewed month by month, day by day, right down to second by second. Data is not discarded. If an appliance

has been running for two years, you can zoom down to by-second views for any time in the previous two years for your archived datasets. Depending on the type of statistic, this could present an issue with system disk usage.

You can monitor the growing sizes of the datasets and destroy datasets that are growing too large. The system disks have compression enabled, so the sizes visible in the datasets view will be larger than the space consumed on disk after compression. For system disk usage and available space, see [“Viewing System Disks Status” in Oracle ZFS Storage Appliance Customer Service Manual](#).

The following are example sizes taken from an appliance that has been running for over four months:

TABLE 7 Sizes Taken from an Appliance Running for over Four Months

Category	Statistic	Span	Dataset Size*	Disk Consumed*
CPU	Percent utilization	130 days	127 MB	36 MB
Protocol	NFSv3 operations per second	130 days	127 MB	36 MB
Protocol	NFSv3 operations per second broken down by type of operation	130 days	209 MB	63 MB
CPU	Percent utilization broken down by CPU mode	130 days	431 MB	91 MB
Network	Device bytes per second broken down by device	130 days	402 MB	119 MB
Disk	I/O bytes per second broken down by disk	130 days	2.18 GB	833 MB
Disk	I/O operations per second broken down by latency	31 days	1.46 GB	515 MB

* *These sizes will vary depending on your workload; they have been provided as a rough guide.*

It is worth noting that the appliance has been intended to have 500 Gbyte mirrored system disks, most of which will be available to store datasets.

The factors that affect consumed disk space are:

- Type of statistic: raw vs breakdowns
- For breakdowns: number of breakdowns, and breakdown name length

- Activity rate

Keep an eye on the size in the Datasets. If a dataset is growing too large, and you want to stop it from growing but keep the historic data - use the suspend action.

Raw Statistics

Statistics that are a single value (sometimes written "as a raw statistic") will not consume much disk space for these reasons:

- Integer values consume a fixed and small amount of space.
- The archives are compressed when saved - which will significantly reduce the size for statistics that are mostly zero.

Examples:

- CPU: percent utilization
- Protocol: NFSv3 operations per second

Breakdowns

Statistics that have breakdowns can consume much more data, as shown in the previous table, since:

- Each breakdown is saved per second. For by-file and by-hostname breakdowns, the number of breakdowns per second may reach into the hundreds (how many different files or hosts had activity in a one second summary) - all of which must be saved to disk.
- Breakdowns have dynamic names, which themselves can be long. You may only have ten active files in your breakdown by-file statistics, but each pathname could be dozens of characters in size. This doesn't sound like much, but the dataset will grow steadily when this data is saved every second.

Examples:

- CPU: percent utilization broken down by CPU mode
- Protocol: NFSv3 operations per second broken down by type of operation
- Disk: I/O bytes per second broken down by disk
- Disk: I/O bytes per second broken down by latency

Exporting Statistics

There may come a time where you'd like to archive statistics on a different server, either to free up disk space on the appliance or for other purposes. For more information about

exporting statistics or downloading statistic data in CSV format, see [“Working with Analytics” on page 11](#).

Execution Performance Impact

Enabling statistics will incur some CPU cost for data collection and aggregation. In many situations, this overhead will not make a noticeable difference on system performance. However for systems under maximum load, including benchmark loads, the small overhead of statistic collection can begin to be noticeable.

Here are some tips for handling execution overheads:

- For dynamic statistics, only archive those that are important to record 24x7.
- Statistics can be suspended, eliminating data collection and the collection overhead. This may be useful if gathering a short interval of a statistic is sufficient for your needs (such as troubleshooting performance). Enable the statistic, wait some minutes, then click the power icon in the Datasets view to suspend it. Suspended datasets keep their data for later viewing.
- Keep an eye on overall performance via the static statistics when enabling and disabling dynamic statistics.
- Be aware that drilldowns will incur overhead for all events. For example, you may trace "NFSv3 operations per second for client deimos", when there is currently no NFSv3 activity from deimos. This doesn't mean that there is no execution overhead for this statistic. The appliance must still trace every NFSv3 event, then compare the host with "deimos" to see if the data should be recorded in this dataset - however we have already paid most of the execution cost at this point.

Static Statistics

Some statistics are sourced from operating system counters are always maintained, which may be called *static statistics*. Gathering these statistics has negligible effect on the performance of the system, since to an extent the system is already maintaining them (they are usually gathered by an operating system feature called *Kstat*). Examples of these statistics are:

TABLE 8 Static Statistics

Category	Statistic
CPU	percent utilization
CPU	percent utilization broken down by CPU mode
Cache	ARC accesses per second broken down by hit/miss
Cache	ARC size

Category	Statistic
Disk	I/O bytes per second
Disk	I/O bytes per second broken down by type of operation
Disk	I/O operations per second
Disk	I/O operations per second broken down by disk
Disk	I/O operations per second broken down by type of operation
Network	device bytes per second
Network	device bytes per second broken down by device
Network	device bytes per second broken down by direction
Protocol	NFSv3/NFSv4/NFSv4.1 operations per second
Protocol	NFSv3/NFSv4/NFSv4.1 operations per second broken down by type of operation

When seen in the BUI, those from the above list without "broken down by" text may have "as a raw statistic".

Since these statistics have negligible execution cost and provide a broad view of system behaviour, many are archived by default. See [“Default Statistics” on page 78](#).

Dynamic Statistics

These statistics are created dynamically, and are not usually maintained by the system (they are gathered by an operating system feature called *DTrace*). Each event is *traced*, and each second this trace data is aggregated into the statistic. And so the cost of this statistic is proportional to the number of events.

Tracing disk details when the activity is 1000 ops/sec is unlikely to have a noticeable affect on performance, however measuring network details when pushing 100,000 packets/sec is likely to have a negative effect. The type of information gathered is also a factor: tracing file names and client names will increase the performance impact.

Examples of dynamic statistics include:

TABLE 9 Dynamic Statistics

Category	Statistic
Protocol	SMB operations per second
Protocol	SMB operations per second broken down by type of operation
Protocol	HTTP/WebDAV requests per second
Protocol	... operations per second broken down by client

Category	Statistic
Protocol	... operations per second broken down by file name
Protocol	... operations per second broken down by share
Protocol	... operations per second broken down by project
Protocol	... operations per second broken down by latency
Protocol	... operations per second broken down by size
Protocol	... operations per second broken down by offset

"..." denotes any of the protocols.

The best way to determine the impact of these statistics is to enable and disable them while running under steady load. Benchmark software may be used to apply that steady load. See [“Working with Analytics” on page 11](#) for the steps to calculate performance impact in this way.

Statistics Actions

The following actions can be performed on statistics/datasets:

TABLE 10 Actions Performed on Statistics/Datasets

Action	Description
Open	Begin reading from the statistic (every second) and cache values in memory as a dataset. In Open Worksheets, statistics are opened when they are added to the view, allowing them to be graphed in real-time. The data is kept in memory while the statistic is being viewed.
Close	Closes the statistic view, discarding the in memory cached dataset.
Archive	Sets the statistic to be permanently opened and archived to disk. If the statistic had already been opened, then all cached data in memory is also archived to disk. Archiving statistics creates permanent datasets, visible in the Datasets view (those with a non-zero "on disk" value). This is how statistics may be recorded 24x7, so that activity from days, weeks and months in the past can be viewed after the fact.
Discard data	Manage the amount of data stored for a particular statistic. You can select to discard the entire dataset or select to remove one of the following granularities of archived data: Second, Minute, or Hour. Note that if you want to delete a higher level of granularity, you must also delete any lower level of granularity. For example,

Action	Description
	to delete the Minute granularity, you must also delete the Second granularity. If you select to not discard the entire dataset, you can discard older data and retain only newer data. You enter an integer value in the "Older than" text box and then select the unit of time: hours, days, weeks, or months. For example, if you want to only retain three weeks of stored data for the selected statistic, you enter a "3" into the "Older than" text box and then select "weeks" from the drop-down menu.
Suspend	Pause an archived statistic. New data will not be read, but the existing disk archive will be left intact.
Resume	Resumes a previously suspended statistic, so that it will continue reading data and writing to the archive.

Default Statistics

The following statistics are enabled and archived by default on a factory-installed appliance. These are the statistics you see in the Datasets view when you first configure and log in to the appliance:

TABLE 11 Default Statistics

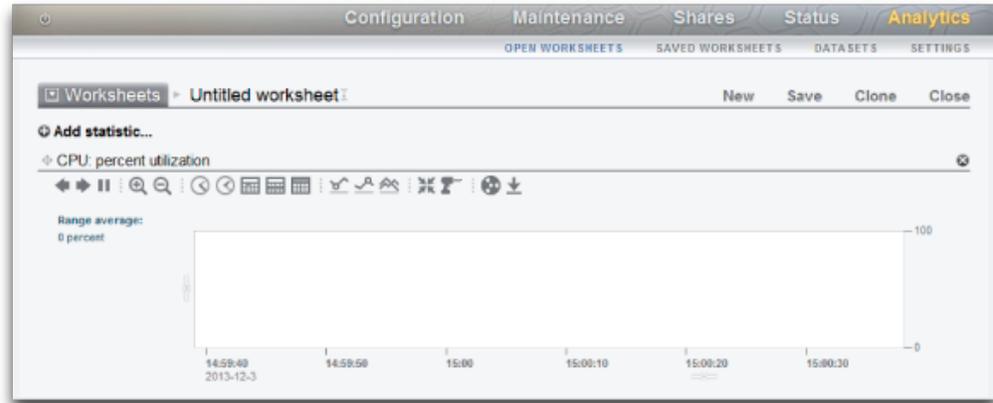
Category	Statistic
CPU	Percent utilization
CPU	Percent utilization broken down by CPU mode
Cache	ARC accesses per second broken down by hit/miss
Cache	ARC size
Cache	ARC size broken down by component
Cache	DNLC accesses per second broken down by hit/miss
Cache	L2ARC accesses per second broken down by hit/miss
Cache	L2ARC size
Data Movement	NDMP bytes transferred to/from disk per second
Disk	Disks with utilization of at least 95 percent broken down by disk
Disk	I/O bytes per second
Disk	I/O bytes per second broken down by type of operation
Disk	I/O operations per second
Disk	I/O operations per second broken down by disk
Disk	I/O operations per second broken down by type of operation

Category	Statistic
Network	Device bytes per second
Network	Device bytes per second broken down by device
Network	Device bytes per second broken down by direction
Protocol	SMB operations per second
Protocol	SMB operations per second broken down by type of operation
Protocol	SMB2 operations per second
Protocol	SMB2 operations per second broken down by type of operation
Protocol	FTP bytes per second
Protocol	Fibre Channel bytes per second
Protocol	Fibre Channel operations per second
Protocol	HTTP/WebDAV requests per second
Protocol	NFSv2 operations per second
Protocol	NFSv2 operations per second broken down by type of operation
Protocol	NFSv3 operations per second
Protocol	NFSv3 operations per second broken down by type of operation
Protocol	NFSv4 operations per second
Protocol	NFSv4 operations per second broken down by type of operation
Protocol	SFTP bytes per second
Protocol	iSCSI operations per second
Protocol	iSCSI bytes per second

These default statistics give broad observability across protocols with minimal statistic collection overhead, and are usually left enabled even when benchmarking. For more discussion on statistic overhead, see [“Storage Performance Impact” on page 72](#) and [“Execution Performance Impact” on page 75](#).

CPU: Percent Utilization

This statistic shows the average utilization of the appliance CPUs. A CPU may be a core on a socket or a hardware thread; the number and type can be seen under analytics interface. For example, a system may have four sockets of quad-core CPUs, meaning there are 16 CPUs available to the appliance. The utilization shown by this statistic is the average across all CPUs.

FIGURE 2 CPU Percent Utilization

The appliance CPUs can reach 100% utilization, which may or may not be a problem. For some performance tests, the appliance is deliberately driven to 100% CPU utilization to measure it at peak performance.

Example

Figure 3 shows CPU: Percent utilization broken down by CPU mode, while the appliance served over 2 Gbytes/sec of cached data over NFSv3.

An average of 82% utilization suggests that there could be more headroom available, and that appliance may be able to serve more than 2 Gbytes/sec (it can). (The breakdowns only add to 81%; the extra 1% is due to rounding.)

The high level of CPU utilization does mean that overall latency of NFS operations may increase, which can be measured by Protocol NFS operations broken down by latency, as operations may be waiting for CPU resources more often.

When to Check CPU Percent Utilization

You can check CPU percent utilization when investigating system bottlenecks. You can also check this statistic when enabling features that consume CPU, such as compression, to gauge the CPU cost of that feature.

CPU: Percent Utilization Breakdowns

The available breakdowns of CPU percent utilization are:

TABLE 12 Breakdowns of Percent Utilization

Breakdown	Description
CPU mode	Either user or kernel. See the CPU modes table below.
CPU identifier	Numeric operating system identifier of the CPU.
application name	Name of the application which is on-CPU.
process identifier	Operating system process ID (PID).
user name	Name of the user who owns the process or thread which is consuming CPU.

The CPU modes are:

TABLE 13 CPU Modes

CPU mode	Description
user	This is a user-land process. The most common user-land process consuming CPU is akd, the appliance kit daemon, which provides administrative control of the appliance.
kernel	This is a kernel-based thread which is consuming CPU. Many of the appliance services are kernel-based, such as NFS and SMB.

Further Analysis

A problem with this CPU utilization average is that it can hide issues when a single CPU is at 100% utilization, which may happen if a single software thread is saturated with work. Use the Advanced Analytic CPUs broken down by percent utilization, which represents utilization as a heat map of CPUs, allowing a single CPU at 100% to be easily identified.

Details

CPU utilization represents the time spent processing CPU instructions in user and kernel code, that are not part of the idle thread. Instruction time includes stall cycles on the memory bus, so high utilization can be caused by the I/O movement of data.

Cache: ARC Accesses

The ARC is the Adaptive Replacement Cache, and is an in-DRAM cache for filesystem and volume data. This statistic shows accesses to the ARC, and allows its usage and performance to be observed.

When to Check ARC Accesses

You can check ARC accesses when investigating performance issues, to understand how well the current workload is caching in the ARC.

ARC Accesses Breakdowns

The available breakdowns of Cache ARC accesses are:

TABLE 14 Breakdowns of ARC Accesses

Breakdown	Description
hit/miss	The result of the ARC lookup. hit/miss states are described in the table below.
filename	The filename that was requested from the ARC. Using this breakdown allows hierarchy mode to be used, so that filesystem directories can be navigated.
L2ARC eligibility	This is the eligibility of L2ARC caching, as measured at the time of ARC access. A high level of ARC misses which are L2ARC eligible would suggest that the workload would benefit from 2nd level cache devices.
project	This shows the project which is accessing the ARC.
share	This shows the share which is accessing the ARC.
LUN	This shows the LUN which is accessing the ARC.

As described in [“Execution Performance Impact” on page 75](#), breakdown such as by filename would be the most expensive to leave enabled.

The hit/miss states are:

TABLE 15 Hit/Miss Breakdowns

Hit/Miss Breakdown	Description
data hits	A data block was in the ARC DRAM cache and returned.
data misses	A data block was not in the ARC DRAM cache. It will be read from the L2ARC cache devices (if available and the data is cached on them) or the pool disks.
metadata hits	A metadata block was in the ARC DRAM cache and returned. Metadata includes the on-disk filesystem framework which refers to the data blocks. Other examples are listed below.
metadata misses	A metadata block was not in the ARC DRAM cache. It will be read from the L2ARC cache devices (if available and the data is cached on them) or the pool disks.

Hit/Miss Breakdown	Description
prefetched data/metadata hits/misses	ARC accesses triggered by the prefetch mechanism, not directly from an application request. More details on prefetch follow.

Metadata

Examples of metadata:

- Filesystem block pointers
- Directory information
- Data deduplication tables
- ZFS uberblock

Prefetch

Prefetch is a mechanism to improve the performance of streaming read workloads. It examines I/O activity to identify sequential reads, and can issue extra reads ahead of time so that the data can be in cache before the application requests it. Prefetch occurs *before the ARC* by performing accesses to the ARC - bear this in mind when trying to understand prefetch ARC activity. For example, if you see:

TABLE 16 Prefetch Types

Type	Description
prefetched data misses	Prefetch identified a sequential workload, and requested that the data be cached in the ARC ahead of time by performing ARC accesses for that data. The data was not in the cache already, and so this is a "miss" and the data is read from disk. This is normal, and is how prefetch populates the ARC from disk.
prefetched data hits	Prefetch identified a sequential workload, and requested that the data be cached in the ARC ahead of time by performing ARC accesses for that data. As it turned out, the data was already in the ARC - so these accesses returned as "hits" (and so the prefetch ARC access wasn't actually needed). This happens if cached data is repeatedly read in a sequential manner.

After data has been prefetched, the application may then request it with its own ARC accesses. Note that the sizes may be different: prefetch may occur with a 128 Kbyte I/O size, while the application may be reading with an 8 Kbyte I/O size. For example, the following doesn't appear directly related:

- Data hits: 368

- Prefetch data misses: 23

However it may be: if prefetch was requesting with a 128 KByte I/O size, $23 \times 128 = 2944$ Kbytes. And if the application was requesting with an 8 Kbyte I/O size, $368 \times 8 = 2944$ Kbytes.

Further Analysis

To investigate ARC misses, check that the ARC has grown to use available DRAM using Cache ARC size.

Cache: L2ARC I/O Bytes

The L2ARC is the 2nd Level Adaptive Replacement Cache, and is an SSD based cache that is accessed before reading from the much slower pool disks. The L2ARC is currently intended for random read workloads. This statistic shows the read and write byte rates to the L2ARC cache devices, if cache devices are present.

When to Check L2ARC I/O Bytes

The Cache: L2ARC I/O Bytes statistic is useful to check during warmup. The write bytes will show the rate of L2ARC warmup time.

L2ARC I/O Bytes Breakdowns

TABLE 17 A Breakdown of L2ARC I/O Bytes

Breakdown	Description
type of operation	Read or write. Read bytes are hits on the cache devices. Write bytes show the cache devices populating with data.

Further Analysis

Also see [“Cache: L2ARC Accesses”](#) on page 84.

Cache: L2ARC Accesses

The L2ARC is the 2nd Level Adaptive Replacement Cache, and is an SSD based cache that is accessed before reading from the much slower pool disks. The L2ARC is currently intended for random read workloads. This statistic shows L2ARC accesses if L2ARC cache devices are present, allowing its usage and performance to be observed.

When to Check L2ARC Accesses

When investigating performance issues, to check how well the current workload is caching in the L2ARC.

L2ARC Accesses Breakdowns

TABLE 18 Breakdowns of L2ARC Accesses

Breakdown	Description
hit/miss	The result of the L2ARC lookup. hit/miss states are described in the table below.
filename	The filename that was requested from the L2ARC. Using this breakdown allows hierarchy mode to be used, so that filesystem directories can be navigated.
L2ARC eligibility	This is the eligibility of L2ARC caching, as measured at the time of L2ARC access.
project	This shows the project which is accessing the L2ARC.
share	This shows the share which is accessing the L2ARC.
LUN	This shows the LUN which is accessing the L2ARC.

As described in [“Execution Performance Impact” on page 75](#), breakdown such as by filename would be the most expensive to leave enabled.

Further Analysis

To investigate L2ARC misses, check that the L2ARC has grown enough in size using the Advanced Analytic Cache L2ARC size. The L2ARC typically takes hours, if not days, to warm up hundreds of Gbytes when feeding from small random reads. The rate can also be checked by examining writes from Cache L2ARC I/O bytes. Also check the Advanced Analytic Cache L2ARC errors to see if there are any errors preventing the L2ARC from warming up.

Cache ARC accesses by L2ARC eligibility can also be checked to see if the data is eligible for L2ARC caching in the first place. Since the L2ARC is intended for random read workloads, it will ignore sequential or streaming read workloads, allowing them to be returned from the pool disks instead.

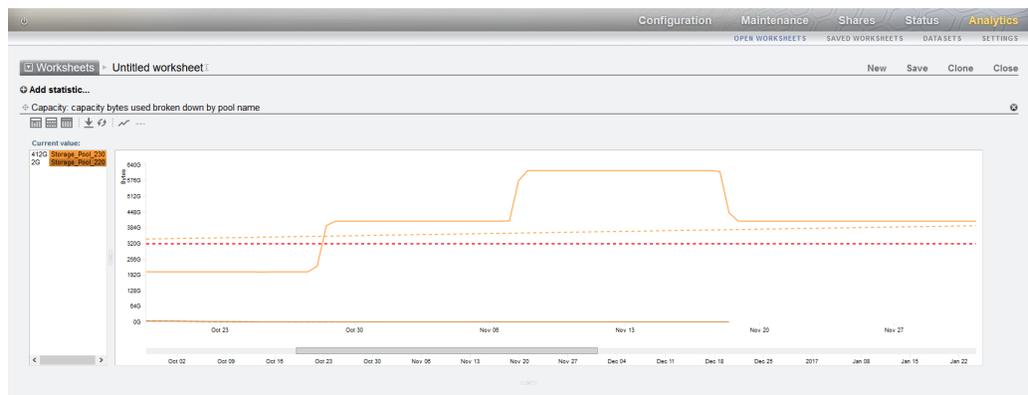
Capacity: Capacity Bytes Used (BUI)

This is a "trendable" statistic that graphically renders datasets, and displays the trends over the full time span of the datasets in the worksheet. These trend graphs can be both viewed and

manipulated. This particular statistic shows the trends for storage capacity usage in Gigabyte units, and is broken down by storage pools. To set a threshold alert for capacity usage by bytes, see “Capacity: Capacity Bytes Used (CLI)” on page 87.

The following figure shows the bytes-used capacity trend broken down by storage pools over the full time span of the datasets.

FIGURE 3 Capacity: Capacity Bytes Used Broken Down by Pool Name

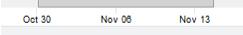


The red dotted line is the reference line, while other dotted lines show the individual trends. The solid lines reflect the actual capacity usage. The dotted lines, except the red-dotted reference line, are color coded with the solid lines to match the color associated with the storage pool, as seen in the Current value pane. When mousing over a pool in the Current value pane, a box displays the pool name, target (or reference) bytes number, and intercept date. The intercept date is when the pool reaches the target (or reference) byte size. When mousing over a pool in the graph, the capacity is displayed for the point in time.

You can control which storage pools are displayed, the time span, the value for the reference line, whether the trend lines are displayed, and when to reload the data, as described in the following table. The reference line is not displayed until a value is set for it.

TABLE 19 Modifying the Capacity Bytes Used Graph

Graph Icon	Element	Description
	Storage Pool	Select a pool in the Current value pane to highlight it and display it in the graph. Select an already highlighted pool to hide it.

Graph Icon	Element	Description
	Reference Line	Select the reference line icon to set a value for the reference line. To hide the reference line, select the icon again. To set a new reference line value, select to show the icon again. The value measurements are: * G = Gigabytes * T = Terabytes
	Trend Line	To show the trend lines, click the trend line icon. To hide the trend lines, select the icon again.
	Reload Data	To reload data from the back-end, and to redisplay data over the full interval on which the data are available, click the reload data icon. Both the reference line and the trend lines are removed.
	Time Span	Select a time span by dragging the mouse on the bar between the dates. After a time span is selected, you can drag the entire, highlighted span to new settings, or you can select an end of the time span and drag it to a new setting.

When to Check Capacity Bytes Used

This statistic can be used to view the capacity usage trend over a set period of time. The information can be used for storage capacity planning on a storage pool basis. To set a threshold alert for capacity usage based on bytes, see [“Capacity: Capacity Bytes Used \(CLI\)” on page 87](#).

Capacity Bytes Used Breakdowns

Pool - Name of the pool for which the capacity trend is shown.

Capacity: Capacity Bytes Used (CLI)

This statistic shows the used bytes, in Gigabyte units, for storage capacity, including data, metadata, and snapshots, except for reservations. It is used as a threshold alert and cannot be displayed in a graph. Unlike other statistics, it is updated every five minutes instead of every second. Various breakdowns are available to show the pool, project, and share capacity used.

To create this capacity alert for datasets in the CLI, navigate to the `analytics datasets` context. Next, use the `create` command to set the alert.

```
hostname:> analytics datasets
hostname:analytics datasets> create cap.bytesused[name]
```

If using worksheets, navigate to the `analytics worksheets` context, select the desired worksheet, and then navigate to the dataset context. Next, use the `set name` command to set the alert. Finally, commit your changes. In the sample below, the `"\n"` character denotes a line break.

```
hostname:> analytics worksheets
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheets worksheet-000> dataset
hostname:analytics worksheets worksheet-000 dataset \
(uncommitted)> set name="cap.bytesused[name]"
hostname:analytics worksheets worksheet-000 dataset \
(uncommitted)> commit
```

For `cap.bytesused`, substitute the appropriate parameters for `[name]` according to the following table.

TABLE 20 Parameters for Capacity Bytes Used Alert

Parameter	Description
[pool]	Selects all pools
[pool= <i>poolname</i>]	Selects the pool named <i>poolname</i>
[project]	Selects all projects
[project= <i>projectname</i>]	Selects the project named <i>projectname</i>
[pool= <i>poolname</i>] [project]	Selects all projects in <i>poolname</i>
[pool= <i>poolname</i>] [project= <i>projectname</i>]	Selects the project named <i>projectname</i> in <i>poolname</i>
[share]	Selects all shares
[share= <i>sharename</i>]	Selects the share named <i>sharename</i>
[pool= <i>poolname</i>] [share]	Selects all shares in <i>poolname</i>
[pool= <i>poolname</i>] [share= <i>sharename</i>]	Selects the share named <i>sharename</i> in <i>poolname</i>
[project= <i>projectname</i>] [share]	Selects all shares in <i>projectname</i>
[project= <i>projectname</i>] [share= <i>sharename</i>]	Selects the share named <i>sharename</i> in <i>projectname</i>
[pool= <i>poolname</i>] [project= <i>projectname</i>] [share]	Selects all shares in <i>projectname</i> in <i>poolname</i>
[pool= <i>poolname</i>] [project= <i>projectname</i>] [share= <i>sharename</i>]	Selects the share named <i>sharename</i> in <i>projectname</i> in <i>poolname</i>

When to Check Capacity Bytes Used

This statistic can be used as a threshold alert for the used storage capacity in bytes. If the threshold is exceeded and the alert is triggered, you can mitigate the situation before storage becomes too full and performance is impacted.

Capacity Bytes Used Breakdowns

- Pool - Name of the pool to set the alert on.
- Project - Name of the project to set the alert on.
- Share - Name of the share to set the alert on.

Further Analysis

See [“Capacity: Capacity Percent Used \(CLI\)” on page 91](#) for a threshold alert for the percentage of storage used capacity.

Capacity: Capacity Percent Used (BUI)

This is a "trendable" statistic that graphically renders datasets, and displays the trends over the full time span of the datasets in the worksheet. These trend graphs can be both viewed and manipulated. This particular statistic shows the trends for storage capacity usage as a percentage, and is broken down by storage pools. To set a threshold alert for capacity usage by percentage, see [“Capacity: Capacity Percent Used \(CLI\)” on page 91](#).

The following figure shows the percent-used capacity trend broken down by storage pools over the full time span of the datasets.

FIGURE 4 Capacity: Capacity Percent Used Broken Down by Pool Name



The red dotted line is the reference line, while other dotted lines show the individual trends. The solid lines reflect the actual capacity usage. The dotted lines, except the red-dotted reference line, are color coded with the solid lines to match the color associated with the storage pool, as seen in the Current value pane. When mousing over a pool in the Current value pane, a box displays the pool name, target (or reference) percentage, and intercept date. The intercept date is when the pool reaches the target (or reference) percentage. When mousing over a pool in the graph, the capacity is displayed for the point in time.

You can control which storage pools are displayed, the time span, the value for the reference line, whether the trend lines are displayed, and when to reload the data, as described in the following table. The reference line is not displayed until a value is set for it.

TABLE 21 Modifying the Capacity Percentage Used Graph

Graph Icon	Element	Description
	Storage Pool	Select a pool in the Current value pane to highlight it and display it in the graph. Select an already highlighted pool to hide it.

Graph Icon	Element	Description
	Reference Line	Select the reference line icon to set a value for the reference line. To hide the reference line, select the icon again. To set a new reference line value, select to show the icon again.
	Trend Line	To show the trend lines, click the trend line icon. To hide the trend lines, select the icon again.
	Reload Data	To reload data from the back-end, and to redisplay data over the full interval on which the data are available, click the reload data icon. Both the reference line and the trend lines are removed.
	Time Span	Select a time span by dragging the mouse on the bar between the dates. After a time span is selected, you can drag the entire, highlighted span to new settings, or you can select an end of the time span and drag it to a new setting.

When to Check Capacity Percent Used

This statistic can be used to view the capacity usage trend over a set period of time. The information can be used for storage capacity planning on a storage pool basis. To set a threshold alert for capacity usage based on percentage, see [“Capacity: Capacity Percent Used \(CLI\)” on page 91](#).

Capacity Percent Used Breakdowns

Pool - Name of the pool for which the capacity trend is shown.

Capacity: Capacity Percent Used (CLI)

This statistic shows the used percentage for storage capacity, including data, metadata, and snapshots, except for reservations. It is used as a threshold alert and cannot be displayed in a graph. Unlike other statistics, it is updated every five minutes instead of every second. Various breakdowns are available to show the pool, project, and share capacity used.

For shares, the storage capacity is the quota, if one exists, or the maximum size on a dynamic LUN. If neither of those exist, the capacity is that of the parent project. For projects, the capacity is the quota, if one exists, or the raw size of the parent pool. For data pools, the capacity is the raw pool size.

To create a capacity alert for datasets in the CLI, navigate to the `analytics datasets` context. Next, use the `create` command to set the alert.

```
hostname:> analytics datasets
hostname:analytics datasets> create cap.percentused[name]
```

If using worksheets, navigate to the `analytics worksheets` context, select the desired worksheet, and then navigate to the dataset context. Next, use the `set name` command to set the alert. Finally, commit your changes. In the example below, the `"\` character denotes a line break.

```
hostname:> analytics worksheets
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheets worksheet-000> dataset
hostname:analytics worksheets worksheet-000 dataset \
(uncommitted)> set name="cap.percentused[name]"
hostname:analytics worksheets worksheet-000 dataset \
(uncommitted)> commit
```

For `cap.percentused`, substitute the appropriate parameters for `[name]` according to the following table.

TABLE 22 Parameters for Capacity Percent Used Alert

Parameter	Description
[pool]	Selects all pools
[pool= <i>poolname</i>]	Selects the pool named <i>poolname</i>
[project]	Selects all projects
[project= <i>projectname</i>]	Selects the project named <i>projectname</i>
[pool= <i>poolname</i>] [project]	Selects all projects in <i>poolname</i>
[pool= <i>poolname</i>] [project= <i>projectname</i>]	Selects the project named <i>projectname</i> in <i>poolname</i>
[share]	Selects all shares
[share= <i>sharename</i>]	Selects the share named <i>sharename</i>
[pool= <i>poolname</i>] [share]	Selects all shares in <i>poolname</i>
[pool= <i>poolname</i>] [share= <i>sharename</i>]	Selects the share named <i>sharename</i> in <i>poolname</i>
[project= <i>projectname</i>] [share]	Selects all shares in <i>projectname</i>
[project= <i>projectname</i>] [share= <i>sharename</i>]	Selects the share named <i>sharename</i> in <i>projectname</i>
[pool= <i>poolname</i>] [project= <i>projectname</i>] [share]	Selects all shares in <i>projectname</i> in <i>poolname</i>
[pool= <i>poolname</i>] [project= <i>projectname</i>] [share= <i>sharename</i>]	Selects the share named <i>sharename</i> in <i>projectname</i> in <i>poolname</i>

When to Check Capacity Percent Used

This statistic can be used as a threshold alert for the percentage of used storage capacity. If the threshold is exceeded and the alert is triggered, you can mitigate the situation before storage becomes too full and performance is impacted.

Capacity Percent Used Breakdowns

- Pool - Name of the pool to set the alert on.
- Project - Name of the project to set the alert on.
- Share - Name of the share to set the alert on.

Further Analysis

See [“Capacity: Capacity Bytes Used \(CLI\)” on page 87](#) for a threshold alert, in bytes, for the storage used capacity.

Capacity: Meta Device Capacity Bytes Used (BUI)

This statistic shows the amount of bytes used on a meta device, broken down by pool name. This statistic provides trending data, similar to the statistic [“Capacity: Capacity Bytes Used \(BUI\)” on page 85](#).

Use this statistic to monitor meta device storage used as bytes, examine the trend of meta device usage, and to alert when capacity reaches a specific threshold.

This statistic is useful only when meta devices are configured for data deduplication.

When to Check Capacity Bytes Used

This statistic can be used as a threshold alert for the used meta device capacity, in bytes. A high capacity of bytes used indicates a need to add more meta devices to the configuration. If the threshold is exceeded and an alert is triggered, you can mitigate the situation before a meta device becomes too full and performance is impacted.

Capacity Bytes Used Breakdown

Pool - Name of the pool for which the capacity trend is shown.

Capacity: Meta Device Capacity Percent Used (BUI)

This statistic shows the percentage of space used on a meta device, broken down by pool name. This statistic provides trending data, similar to the statistic “[Capacity: Capacity Percent Used \(BUI\)](#)” on page 89.

Use this statistic to monitor meta device storage usage as a percentage, examine the trend of meta device usage, and to alert when capacity reaches a specific threshold.

This statistic is useful only when meta devices are configured for data deduplication.

When to Check Capacity Percent Used

This statistic can be used as a threshold alert for the percentage of used meta device capacity. A high percentage use indicates a need to add more meta devices to the configuration. The appliance generates an alert when a meta device, used in any pool, reaches a threshold of 85 percent. If the threshold is exceeded and an alert is triggered, you can mitigate the situation before a meta device becomes too full and performance is impacted.

Capacity Percent Used Breakdown

Pool - Name of the pool for which the capacity trend is shown.

Capacity: System Pool Bytes Used

This statistic shows the used bytes, in Gigabyte units, for the system pool capacity, including data, metadata, and snapshots, except for reservations. It is used as a threshold alert and cannot be displayed in a graph. Unlike other statistics, it is updated every five minutes instead of every second.

To create this capacity alert in the CLI, navigate to the analytics and datasets context. If using worksheets, navigate to analytics, the desired worksheet, and then the dataset context. For datasets, use the "create" command. For worksheets, use the "set name" command. Below, the "\n" character denotes a line break.

```
hostname:> analytics
hostname:analytics> datasets
hostname:analytics datasets> create syscap.bytesused
```

OR

```
hostname:> analytics
```

```
hostname:analytics> worksheets
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheets worksheet-000> dataset
hostname:analytics worksheets worksheet-000 dataset \
(uncommitted)> set name="syscap.bytesused"
hostname:analytics worksheets worksheet-000 dataset \
(uncommitted)> commit
```

When to Check System Pool Bytes Used

This statistic can be used as a threshold alert for the used system pool capacity in bytes. If the threshold is exceeded and the alert is triggered, you can mitigate the situation before the system pool becomes too full and performance is impacted.

System Pool Bytes Used Breakdowns

None.

Further Analysis

See [“Capacity: System Pool Percent Used” on page 95](#) for a threshold alert for the percentage of the system pool used capacity.

Capacity: System Pool Percent Used

This statistic shows the used percentage of the system pool capacity based on the raw pool size. It is used as a threshold alert and cannot be displayed in a graph. Unlike other statistics, it is updated every five minutes instead of every second.

To create this capacity alert in the CLI, navigate to the analytics and datasets context. If using worksheets, navigate to analytics, the desired worksheet, and then the dataset context. For datasets, use the "create" command. For worksheets, use the "set name" command. Below, the "\" character denotes a line break.

```
hostname:> analytics datasets
hostname:analytics datasets> create syscap.percentused
```

OR

```
hostname:> analytics
hostname:analytics> worksheets
hostname:analytics worksheets> select worksheet-000
hostname:analytics worksheets worksheet-000> dataset
hostname:analytics worksheets worksheet-000 dataset \
```

```
(uncommitted)> set name="syscap.percentused"  
hostname:analytics worksheets worksheet-000 dataset \  
(uncommitted)> commit
```

When to Check System Pool Percent Used

This statistic can be used as a threshold alert for the used system pool capacity in a percentage. If the threshold is exceeded and the alert is triggered, you can mitigate the situation before the system pool becomes too full and performance is impacted. It is recommended that you set the threshold alert to 80 percent. For information on setting threshold alerts, see [“Configuring a Threshold Alert \(BUI\)” on page 52](#) or [“Configuring a Threshold Alert \(CLI\)” on page 52](#).

System Pool Percent Used Breakdowns

None.

Further Analysis

See [“Capacity: System Pool Bytes Used” on page 94](#) for a threshold alert, in bytes, for the system pool used capacity.

Data Movement: Shadow Migration Bytes

This statistic tracks total shadow migration bytes per second transferred as part of migrating file or directory contents. This does not apply to metadata (extended attributes, ACLs, etc). It gives a rough approximation of the data transferred, but source datasets with a large amount of metadata will show a disproportionately small bandwidth. The complete bandwidth can be observed by looking at network analytics.

When to Check Shadow Migration Bytes

When investigating Shadow Migration activity.

Shadow Migration Bytes Breakdowns

TABLE 23 Breakdowns of Shadow Migration Bytes

Breakdown	Description
filename	The filename that was migrated. Using this breakdown allows hierarchy mode to be used, so that filesystem directories can be navigated.
project	This shows the project which contains a shadow migration.
share	This shows the share which is being migrated.

Further Analysis

Also see [“Data Movement: Shadow Migration Ops” on page 97](#) and [“Data Movement: Shadow Migration Requests” on page 97](#).

Data Movement: Shadow Migration Ops

This statistic tracks shadow migration operations that require going to the source filesystem.

When to Check Shadow Migration Ops

When investigating Shadow Migration activity.

Shadow Migration Ops Breakdowns

TABLE 24 Breakdowns of Shadow Migration Ops

Breakdown	Description
filename	The filename that was migrated. Using this breakdown allows hierarchy mode to be used, so that filesystem directories can be navigated.
project	This shows the project which contains a shadow migration.
share	This shows the share which is being migrated.
latency	Measure the latency of requests from the shadow migration source.

Further Analysis

Also see [“Data Movement: Shadow Migration Bytes” on page 96](#) and [“Data Movement: Shadow Migration Requests” on page 97](#).

Data Movement: Shadow Migration Requests

This statistic tracks shadow migration requests for files or directories that are not cached and known to be local to the filesystem. It does account for both migrated and unmigrated files and directories, and can be used to track the latency incurred as part of shadow migration, as well as track the progress of background migration. It currently encompasses both synchronous and asynchronous (background) migration, so it's not possible to view only latency visible to clients.

When to Check Shadow Migration Requests

When investigating Shadow Migration activity.

Shadow Migration Requests Breakdowns

TABLE 25 Breakdowns of Shadow Migration Requests

Breakdown	Description
filename	The filename that was migrated. Using this breakdown allows hierarchy mode to be used, so that filesystem directories can be navigated.
project	This shows the project which contains a shadow migration.
share	This shows the share which is being migrated.
latency	Measure the latency incurred as part of shadow migration.

Further Analysis

Also see [“Data Movement: Shadow Migration Ops”](#) on page 97 and [“Data Movement: Shadow Migration Bytes”](#) on page 96.

Data Movement: NDMP Bytes Statistics

This statistic shows total NDMP bytes transferred per second during backup or restore operations. It indicates how much data is being read or written for NDMP backups or restores. This statistic is zero unless NDMP is configured and active.

When to Check NDMP Bytes Statistics

When investigating NDMP backup and restore performance.

NDMP Bytes Statistics Breakdowns

TABLE 26 Breakdowns of NDMP Bytes

Breakdown	Description
type of operation	Read or write
client	Remote hostname or IP address of the NDMP client
session	Set of data streams managed by NDMP
type of I/O	Network, disk, tape, etc.

Breakdown	Description
file	Used with tar and dump

Further Analysis

Also see [“Data Movement: NDMP Operations Statistics”](#) on page 99.

Data Movement: NDMP Operations Statistics

This statistic shows total NDMP backup or restore operations performed per second. This statistic is zero unless NDMP is configured and active.

When to Check NDMP Operations Statistics

When investigating NDMP backup and restore performance.

NDMP Operations Statistics Breakdowns

TABLE 27 Breakdowns of NDMP Operations

Breakdown	Description
type of operation	Read or write
client	Remote hostname or IP address of the NDMP client
session	Set of data streams managed by NDMP
type of I/O	Network, disk, tape, etc.
latency	Time elapsed between operations
size	Number of bytes read/written per operation
offset	Location within a backup stream, buffer, file, etc.

Further Analysis

Also see [“Data Movement: NDMP Bytes Statistics”](#) on page 98.

Data Movement: Replication Bytes

This statistic tracks network data throughput of project/share replication in bytes per second.

When to Check Replication Bytes

When investigating replication activity and replication network usage.

Replication Bytes Breakdowns

TABLE 28 Breakdowns of Replication Bytes

Breakdown	Description
direction	Shows bytes broken down by direction, to or from the appliance.
type of operation	Shows bytes broken down by the type of operation with the remote appliance, read or write.
peer	Shows bytes broken down by the name of the remote appliances.
pool name	Shows bytes broken down by the name of the pools.
project	Shows bytes broken down by the name of the projects.
dataset	Shows bytes broken down by the name of the shares.
as a raw statistic	Shows bytes as raw statistics.

Further Analysis

Also see [“Data Movement: Replication Operations” on page 100](#)

Data Movement: Replication Operations

This statistic tracks Replication read and write operations performed by the Replication service.

When to Check Replication Operations

When investigating replication activity.

Replication Operations Breakdowns

TABLE 29 Breakdowns of Replication Operations

Breakdown	Description
direction	Shows IO operations broken down by direction, to or from the appliance.
type of operation	Shows IO operations broken down by type of operation with the remote appliance, read or write.
peer	Shows IO operations broken down by the name of the remote appliances.

Breakdown	Description
pool name	Shows IO operations broken down by the name of the pools.
project	Shows IO operations broken down by the name of the projects.
dataset	Shows IO operations broken down by the name of the shares.
latency	Measures the current network latency experienced during replication data transfer.
offset	Measures the offset in all replication transfers relative to the start of each individual replication update.
size	Measures the read/write operations size performed by replication service.
as a raw statistic	Shows IO operations as raw statistics.

Further Analysis

Also see [“Data Movement: Replication Bytes”](#) on page 99.

To examine data throughput of project/share replication at the ZFS send/receive internal interface, see [“Data Movement: Replication Send/Receive Bytes”](#) on page 137.

Disk: Disks

The Disks statistic is used to display the heat map for disks broken down by percent utilization. This is the best way to identify when pool disks are under heavy load. It may also identify problem disks that are beginning to perform poorly, before their behavior triggers a fault and automatic removal from the pool.

When to Check Disks

Any investigation into disk performance.

Disks Breakdowns

TABLE 30 A Breakdown of Disks

Breakdown	Description
percent utilization	A heat map with utilization on the Y-axis and each level on the Y-axis colored by the number of disks at that utilization: from light (none) to dark (many).

Interpretation

Utilization is a better measure of disk load than IOPS or throughput. Utilization is measured as the time during which that disk was busy performing requests (see Details below). At 100% utilization the disk may not be able to accept more requests, and additional I/O may wait on a queue. This I/O wait time will cause latency to increase and overall performance to decrease.

In practise, disks with a consistent Utilization of 75% or higher are an indication of heavy disk load.

The heat map allows a particular pathology to be easily identified: a single disk misperforming and reaching 100% utilization (a bad disk). Disks can exhibit this symptom before they fail. Once disks fail, they are automatically removed from the pool with a corresponding alert. This particular problem is during the time *before* they fail, when their I/O latency is increasing and slowing down overall appliance performance, but their status is considered healthy - they have yet to identify any error state. This situation will be seen as a faint line at the top of the heat map, showing that a single disk has stayed at 100% utilization for some time.

Suggested interpretation summary:

TABLE 31 Interpretation Summary

Observed	Suggested Interpretation
Most disks consistently over 75%	Available disk resources are being exhausted.
Single disk at 100% for several seconds	This can indicate a bad disk that is about to fail.

Further Analysis

To understand the nature of the I/O, such as IOPS, throughput, I/O sizes and offsets, see [“Disk: I/O Operations” on page 104](#) and [“Disk: I/O Bytes” on page 102](#).

Details

This statistic is actually a measure of percent busy, which serves as a reasonable approximation of percent utilization since the appliance manages the disks directly. Technically this isn't a direct measure of disk utilization: at 100% busy, a disk may be able to accept more requests which it serves concurrently by inserting into and reordering its command queue, or serves from its on-disk cache.

Disk: I/O Bytes

This statistic shows the back-end throughput to the disks, in I/O bytes per second, after the appliance has processed logical I/O into physical I/O based on share settings and software

RAID settings. To configure the RAID settings, see [“Configuring Storage” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0.](#)

For example, an 8 Kbyte write over NFSv3 may become a 128 Kbyte write after the record size is applied from the share settings, which may then become a 256 Kbyte write to the disks after mirroring is applied, plus additional bytes for filesystem metadata. On the same mirrored environment, an 8 Kbyte NFSv3 read may become a 128 Kbyte disk read after the record size is applied; however, this does not get doubled by mirroring (the data only needs to be read from one half of the mirrored environment). It can help to monitor throughput at all layers at the same time to examine this behavior, for example by viewing:

- [“Network: Device Bytes” on page 106](#) - data rate on the network (logical)
- [“Disk: ZFS Logical I/O Bytes” on page 140](#) - data rate to the share (logical)
- Disk: I/O Bytes - data rate to the disks (physical)

When viewing the I/O bytes per second broken down by operation type, the operation pane shows the read and write statistics. Select an operation in the pane to highlight it and display it separately, by color, in the graph. Select an already highlighted operation to not display it separately in the graph.

When viewing the I/O bytes per second broken down by disk, the disk breakdown pane shows the statistics per the name of the storage pool disk or system disk. Select a disk in the disk breakdown pane to highlight it and display it separately, by color, in the graph. Select an already highlighted disk to not display it separately in the graph. When mousing over a disk in this pane, a box displays the following information:

- Disk name - controller or disk shelf name/label: I/O bytes per second
- Disk Type: typically HDD or SSD
- Type: typically System, Data, Cache, or Log
- Size
- RPM (not displayed for SSDs)

To display the hierarchy view for all disks, click View Hierarchy, below the disk breakdown pane. The I/O bytes per second are shown for the controller and each disk shelf. Click Refresh hierarchy to refresh the hierarchical breakdown visible in the graph. To close this view, click the close icon .

When to Check I/O Bytes

Use this statistic to understand the nature of back-end disk I/O, based on bytes, after an issue based on disk utilization or latency has been observed. It is difficult to identify an issue from disk I/O throughput alone: A single disk may be performing well at 50 Mbytes/sec (sequential I/O), yet poorly at 5 Mbytes/sec (random I/O).

Use the disk breakdown pane and the hierarchy view to determine if the disk shelves are balanced in regard to disk I/O throughput. When examining disk throughput, it is common that cache and log devices have higher throughput than other storage pool disks.

I/O Bytes Breakdowns

TABLE 32 Breakdowns of I/O Bytes

Breakdown	Description
type of operation	Read or write.
disk	Storage pool disk or system disk. This breakdown can identify system disk I/O versus pool disk I/O, and I/O to cache and log devices.

Further Analysis

For the best measure of disk utilization, see [“Disk: Disks” on page 101](#). To examine operations/sec instead of bytes/sec, see [“Disk: I/O Operations” on page 104](#).

Disk: I/O Operations

This statistic shows the back-end I/O operations per second to the disks (disk IOPS) after the appliance has processed logical I/O into physical I/O based on share settings and software RAID settings. To configure the RAID settings, see [“Configuring Storage” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#).

For example, 16 sequential 8 Kbyte NFSv3 writes may become a single 128 Kbyte write sometime later after the data has been buffered in the ARC DRAM cache, which may then become multiple disk writes due to RAID, such as two writes to each half of a mirror. It can help to monitor I/O at all layers at the same time to examine this behavior, for example by viewing:

- [“Protocol: NFSv\[2-4\] Operations” on page 117](#) - NFS writes (logical)
- [“Disk: ZFS Logical I/O Operations” on page 140](#) - share I/O (logical)
- Disk: I/O Operations - I/O to the disks (physical)

This statistic includes a breakdown of disk I/O latency, which is a direct measure of performance for synchronous I/O, and also useful as a measure of the magnitude of back-end disk load. It is difficult to identify issues from disk IOPS alone without considering latency: A single disk may be performing well at 400 IOPS (sequential and small I/O hitting mostly from

the disk's on-board DRAM cache), yet poorly at 110 IOPS (random I/O causing head seek and waiting on disk rotation).

The latency breakdown is presented as a heat map, showing the pattern of I/O latency, along with outliers. Mouse over the outliers icon  to see the current value, and click the icon to toggle between different percentages of outlier elimination. Disk I/O latency is often related to the performance of the delivered logical I/O, such as with synchronous reads (non-prefetch) and synchronous writes. Sometimes the latency is not directly related to logical I/O performance, such as for asynchronous writes that are later flushed to disk, and for prefetch reads.

Because it is difficult to determine a per-disk IOPS limit, also examine the disk IOPS by the offset, which helps identify the IOPS type as either random or sequential, and the I/O size. Use the following breakdowns to observe these attributes:

- Disk: I/O Operations - broken down by offset
- Disk: I/O Operations - broken down by size

When viewing a breakdown, select an individual result in the pane to highlight it and display it separately, by color, in the graph. Select an already highlighted result to not display it separately in the graph.

When viewing the disk IOPS by the disk breakdown, mousing over a disk in the disk breakdown pane displays a box with the following information:

- Disk name - controller or disk shelf name/label: I/O operations per second
- Disk Type: typically HDD or SSD
- Type: typically System, Data, Cache, or Log
- Size
- RPM (not displayed for SSDs)

To display the hierarchy view for all disks, click View Hierarchy, below the disk breakdown pane. The I/O operations per second are shown for the controller and each disk shelf. Click Refresh hierarchy to refresh the hierarchical breakdown visible in the graph. To close this view, click the close icon .

When to Check I/O Operations

Use this statistic to understand the nature of back-end disk I/O, based on disk I/O operations per second (IOPS), after an issue based on disk utilization or latency has been observed.

Use the disk breakdown pane and the hierarchy view to determine if the disk shelves are balanced in regard to disk IOPS. When examining disk IOPS, it is common that cache and log devices have higher throughput than other storage pool disks.

I/O Operations Breakdowns

TABLE 33 Breakdowns of I/O Operations

Breakdown	Description
type of operation	Read or write.
disk	Storage pool disk or system disk. This breakdown can identify system disk I/O versus pool disk I/O, and I/O to cache and log devices.
size	Heat map showing the distribution of I/O sizes.
latency	Heat map showing the latency of disk I/O, as measured from when the I/O was requested to the disk to when the disk returned the completion.
offset	Heat map showing the disk location offset of disk I/O. This can be used to identify random or sequential disk IOPS. To better see details, use the zoom-in icon  .

Further Analysis

For the best measure of disk utilization, see [“Disk: Disks” on page 101](#). To examine bytes/sec instead of operations/sec, see [“Disk: I/O Bytes” on page 102](#).

Network: Device Bytes

This statistic measures network device activity in bytes/sec. Network devices are the physical network ports (see [“Network Configuration” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#)). The measured bytes by this statistic includes all network payload headers (Ethernet, IP, TCP, NFS/SMB/etc.)

When to Check Device Bytes

Network bytes can be used a rough measure of appliance load. It should also be checked whenever performance issues are investigated, especially for 1 Gbit/sec interfaces, in case the bottleneck is the network device. The maximum practical throughput for network devices in each direction (in or out) based on speed:

- 1 Gbit/sec Ethernet: ~120 Mbytes/sec device bytes
- 10 Gbit/sec Ethernet: ~1.16 Gbytes/sec device bytes

If a network device shows a higher rate than these, use the direction breakdown to see the inbound and outbound components.

Device Bytes Breakdowns

TABLE 34 Breakdowns of Device Bytes

Breakdown	Description
direction	In or out, relative to the appliance. For example, NFS reads to the appliance would be show as out(bound) network bytes.
device	Network device (see Devices in Network).

Further Analysis

Also see [“Network: Interface Bytes” on page 107](#) for network throughput at the interface level, instead of the device level.

Network: Interface Bytes

This statistic measures network interface activity in bytes/sec. Network interfaces are the logical network interfaces (see [“Network Configuration” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#)). The measured bytes by this statistic includes all network payload headers (Ethernet, IP, TCP, NFS/SMB/etc.)

Example

See [“Network: Device Bytes” on page 106](#) for an example of a similar statistic with similar breakdowns.

When to Check Interface Bytes

Network bytes can be used a rough measure of appliance load. This statistic can be used to see the rate of network bytes through different interfaces. To examine network devices that make up an interface, especially to identify if there are balancing problems with LACP aggregations, use the Network Device bytes statistic.

Interface Bytes Breakdowns

TABLE 35 Breakdowns of Interface Bytes

Breakdown	Description
direction	In or out, relative to the appliance. For example, NFS reads to the appliance would be show as out(bound) network bytes.

Breakdown	Description
interface	Network interface (see Interfaces in Network).

Further Analysis

Also see [“Network: Device Bytes” on page 106](#) for network throughput at the device level, instead of the interface level.

Protocol: SMB Operations

This statistic shows SMB operations/sec (SMB IOPS) requested by clients to the appliance. Various useful breakdowns are available: to show the client, filename and latency of the SMB I/O.

Example

See [“Protocol: NFSv\[2-4\] Operations” on page 117](#) for an example of a similar statistic with similar breakdowns.

When to Check SMB Operations

SMB operations/sec can be used as an indication of SMB load, and can be viewed on the dashboard.

Use the latency breakdown when investigating SMB performance issues, especially to quantify the magnitude of the issue. This measures the I/O latency component for which the appliance is responsible for, and displays it as a heat map so that the overall latency pattern can be seen, along with outliers. If the SMB latency is high, drill down further on latency to identify the type of operation and filename for the high latency, and, check other statistics for both CPU and Disk load to investigate why the appliance is slow to respond; if latency is low, the appliance is performing quickly, and any performance issues experienced on the client are more likely to be caused by other factors in the environment: such as the network infrastructure, and CPU load on the client itself.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client and filename breakdowns, and the filename hierarchy view. Client and especially filename breakdowns can be very expensive in terms of storage and execution overhead. Therefore, it is not recommended to permanently enable these breakdowns on a busy production appliance.

SMB Operations Breakdowns

TABLE 36 Breakdowns of SMB Operations

Breakdown	Description
type of operation	SMB operation type (read/write/readX/writeX/...)
client	Remote hostname or IP address of the SMB client.
filename	Filename for the SMB I/O, if known and cached by the appliance. If the filename is not known it is reported as "<unknown>".
share	The share for this SMB I/O.
project	The project for this SMB I/O.
latency	A heat map showing the latency of SMB I/O, as measured from when the SMB request arrived on the appliance from the network, to when the response is sent; this latency includes the time to process the SMB request, and to perform any disk I/O.
size	A heat map showing the distribution of SMB I/O sizes.
offset	A heat map showing the file offset of SMB I/O. This can be used to identify random or sequential SMB IOPS. Use the Disk I/O operations statistic to check whether random SMB IOPS maps to random Disk IOPS after the filesystem and RAID configuration has been applied.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: SMB operations per second of type read broken down by latency" (to examine latency for reads only)
- "Protocol: SMB operations per second for file '/export/fs4/10ga' broken down by offset" (to examine file access pattern for a particular file)
- "Protocol: SMB operations per second for client 'phobos.sf.fishpong.com' broken down by filename" (to view which files a particular client is accessing)

Further Analysis

See [“Network: Device Bytes” on page 106](#) for a measure of network throughput caused by the SMB activity; [“Cache: ARC Accesses” on page 81](#) to learn how well an SMB read workload is returning from cache; and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: Fibre Channel Bytes

This statistic shows Fibre Channel bytes/sec requested by initiators to the appliance.

Example

See [“Protocol: iSCSI Bytes” on page 114](#) for an example of a similar statistic with similar breakdowns.

When to Check Fibre Channel Bytes

Fibre Channel bytes/sec can be used as an indication of FC load, in terms of throughput. For a deeper analysis of FC activity, see [“Protocol: Fibre Channel Operations” on page 110](#).

Fibre Channel Bytes Breakdowns

TABLE 37 Breakdowns of Fibre Channel Bytes

Breakdown	Description
initiator	Fibre Channel client initiator
target	Local SCSI target
project	The project for this FC request.
lun	The LUN for this FC request.

See [“Configuring Storage Area Network \(SAN\)” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#) for terminology definitions.

Further Analysis

See [“Protocol: Fibre Channel Operations” on page 110](#) for numerous other breakdowns on FC operations; also see [“Cache: ARC Accesses” on page 81](#) to learn how well an FC read workload is returning from cache, and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: Fibre Channel Operations

This statistic shows Fibre Channel operations/sec (FC IOPS) requested by initiators to the appliance. Various useful breakdowns are available: to show the initiator, target, type and latency of the FC I/O.

Example

See [“Protocol: iSCSI Operations” on page 115](#) for an example of a similar statistic with similar breakdowns.

When to Check Fibre Channel Operations

Fibre Channel operations/sec can be used as an indication of FC load, and can also be viewed on the dashboard.

Use the latency breakdown when investigating FC performance issues, especially to quantify the magnitude of the issue. This measures the I/O latency component for which the appliance is responsible for, and displays it as a heat map so that the overall latency pattern can be seen, along with outliers. If the FC latency is high, drill down further on latency to identify the client initiator, the type of operation and LUN for the high latency, and, check other statistics for both CPU and Disk load to investigate why the appliance is slow to respond; if latency is low, the appliance is performing quickly, and any performance issues experienced on the client initiator are more likely to be caused by other factors in the environment: such as the network infrastructure, and CPU load on the client itself.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client initiator, lun and command breakdowns.

Fibre Channel Operations Breakdowns

TABLE 38 Breakdowns of Fibre Channel Operations

Breakdown	Description
initiator	Fibre Channel client initiator
target	Local SCSI target
project	The project for this FC request.
lun	The LUN for this FC request.
type of operation	FC operation type. This shows how the SCSI command is transported by the FC protocol, which can give an idea to the nature of the I/O.
command	SCSI command sent by the FC protocol. This can show the real nature of the requested I/O (read/write/sync-cache/...).
latency	A heat map showing the latency of FC I/O, as measured from when the FC request arrived on the appliance from the network, to when the response is sent; this latency includes the time to process the FC request, and to perform any disk I/O.
offset	A heat map showing the file offset of FC I/O. This can be used to identify random or sequential FC IOPS. Use the Disk I/O operations statistic to check whether random FC IOPS maps to random Disk IOPS after the LUN and RAID configuration has been applied.
size	A heat map showing the distribution of FC I/O sizes.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: Fibre Channel operations per second of command read broken down by latency" (to examine latency for SCSI reads only)

Further Analysis

See [“Protocol: Fibre Channel Bytes” on page 109](#) for the throughput of this FC I/O; also see [“Cache: ARC Accesses” on page 81](#) to learn how well an FC read workload is returning from cache, and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: FTP Bytes

This statistic shows FTP bytes/sec requested by clients to the appliance. Various useful breakdowns are available: to show the client, user and filename of the FTP requests.

Example

FTP

When to Check FTP Bytes

FTP bytes/sec can be used as an indication of FTP load, and can be viewed on the dashboard.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client, user and filename breakdowns, and the filename hierarchy view. It may be best to enable these breakdowns for short periods only: the by-filename breakdown can be one of the most expensive in terms of storage and execution overhead, and may not be suitable to leave enabled permanently on appliances with high rates of FTP activity.

FTP Bytes Breakdowns

TABLE 39 Breakdowns of FTP Bytes

Breakdown	Description
type of operation	FTP operation type (get/put/...)
user	Username of the client
filename	Filename for the FTP operation, if known and cached by the appliance. If the filename is not known it is reported as "<unknown>".
share	The share for this FTP request.
project	The project for this FTP request.
client	Remote hostname or IP address of the FTP client.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: FTP bytes per second for client 'phobos.sf.fishpong.com' broken down by filename" (to view which files a particular client is accessing)

Further Analysis

See [“Cache: ARC Accesses” on page 81](#) to learn how well an FTP read workload is returning from cache; and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: HTTP/WebDAV Requests

This statistic shows HTTP/WebDAV requests/sec requested by HTTP clients. Various useful breakdowns are available: to show the client, filename and latency of the HTTP request.

When to Check HTTP/WebDAV Requests

HTTP/WebDAV requests/sec can be used as an indication of HTTP load, and can also be viewed on the dashboard.

Use the latency breakdown when investigating HTTP performance issues, especially to quantify the magnitude of the issue. This measures the latency component for which the appliance is responsible for, and displays it as a heat map so that the overall latency pattern can be seen, along with outliers. If the HTTP latency is high, drill down further on latency to identify the file, size and response code for the high latency HTTP requests, and, check other statistics for both CPU and Disk load to investigate why the appliance is slow to respond; if latency is low, the appliance is performing quickly, and any performance issues experienced on the client initiator are more likely to be caused by other factors in the environment: such as the network infrastructure, and CPU load on the client itself.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client, response code and requested filename breakdowns.

HTTP/WebDAV Requests Breakdowns

TABLE 40 Breakdowns of HTTPWebDAV Requests

Breakdown	Description
type of operation	HTTP request type (get/post)
response code	HTTP response (200/404/...)
client	Client hostname or IP address

Breakdown	Description
filename	Filename requested by HTTP
latency	A heat map showing the latency of HTTP requests, as measured from when the HTTP request arrived on the appliance from the network, to when the response is sent; this latency includes the time to process the HTTP request, and to perform any disk I/O.
size	A heat map showing the distribution of HTTP request sizes.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: HTTP/WebDAV operations per second of type get broken down by latency" (to examine latency for HTTP GETs only)
- "Protocol: HTTP/WebDAV requests per second for response code '404' broken down by filename (to see which non-existent files were requested)
- "Protocol: HTTP/WebDAV requests per second for client 'deimos.sf.fishpong.com' broken down by filename" (to examine files requested by a particular client)

Further Analysis

See [“Network: Device Bytes” on page 106](#) for a measure of network throughput caused by HTTP activity; also see [“Cache: ARC Accesses” on page 81](#) to learn how well an HTTP read workload is returning from cache, and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: iSCSI Bytes

This statistic shows iSCSI bytes/sec requested by initiators to the appliance.

When to Check iSCSI Bytes

iSCSI bytes/sec can be used as an indication of iSCSI load, in terms of throughput. For a deeper analysis of iSCSI activity, see [“Protocol: iSCSI Operations” on page 115](#).

iSCSI Bytes Breakdowns

TABLE 41 Breakdowns of iSCSI Bytes

Breakdown	Description
initiator	iSCSI client initiator

Breakdown	Description
target	Local SCSI target
project	The project for this iSCSI request.
lun	The LUN for this iSCSI request.
client	The remote iSCSI client hostname or IP address

See “[Configuring Storage Area Network \(SAN\)](#)” in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0* for terminology definitions.

Further Analysis

See “[Protocol: iSCSI Operations](#)” on page 115 for numerous other breakdowns on iSCSI operations; also see “[Cache: ARC Accesses](#)” on page 81 to learn how well an iSCSI read workload is returning from cache, and “[Disk: I/O Operations](#)” on page 104 for the back-end disk I/O caused.

Protocol: iSCSI Operations

This statistic shows iSCSI operations/sec (iSCSI IOPS) requested by initiators to the appliance. Various useful breakdowns are available: to show the initiator, target, type and latency of the iSCSI I/O.

When to Check iSCSI Operations

iSCSI operations/sec can be used as an indication of iSCSI load, and can also be viewed on the dashboard.

Use the latency breakdown when investigating iSCSI performance issues, especially to quantify the magnitude of the issue. This measures the I/O latency component for which the appliance is responsible for, and displays it as a heat map so that the overall latency pattern can be seen, along with outliers. If the iSCSI latency is high, drill down further on latency to identify the client initiator, the type of operation and LUN for the high latency, and, check other statistics for both CPU and Disk load to investigate why the appliance is slow to respond; if latency is low, the appliance is performing quickly, and any performance issues experienced on the client initiator are more likely to be caused by other factors in the environment: such as the network infrastructure, and CPU load on the client itself.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client initiator, lun and command breakdowns.

iSCSI Operations Breakdowns

TABLE 42 Breakdowns of iSCSI Operations

Breakdown	Description
initiator	iSCSI client initiator
target	Local SCSI target
project	The project for this iSCSI request.
lun	The LUN for this iSCSI request.
type of operation	iSCSI operation type. This shows how the SCSI command is transported by the iSCSI protocol, which can give an idea to the nature of the I/O.
command	SCSI command sent by the iSCSI protocol. This can show the real nature of the requested I/O (read/write/sync-cache/...).
latency	A heat map showing the latency of iSCSI I/O, as measured from when the iSCSI request arrived on the appliance from the network, to when the response is sent; this latency includes the time to process the iSCSI request, and to perform any disk I/O.
offset	A heat map showing the file offset of iSCSI I/O. This can be used to identify random or sequential iSCSI IOPS. Use the Disk I/O operations statistic to check whether random iSCSI IOPS maps to random Disk IOPS after the LUN and RAID configuration has been applied.
size	A heat map showing the distribution of iSCSI I/O sizes.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: iSCSI operations per second of command read broken down by latency" (to examine latency for SCSI reads only)

Further Analysis

See [“Protocol: iSCSI Bytes” on page 114](#) for the throughput of this iSCSI I/O; also see [“Cache: ARC Accesses” on page 81](#) to learn how well an iSCSI read workload is returning from cache, and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: NFSv[2-4] Bytes

This statistic shows NFSv[2-4] bytes/sec transferred between NFS clients and the appliance. Supported NFS versions are: NFSv2, NFSv3, NFSv4.0, and NFSv4.1. Bytes statistics can be broken down by: operation, client, filename, share, and project.

When to Check NFSv[2-4] Bytes

NFSv[2-4] bytes/sec can be used as an indication of NFS load. The best way to improve performance is to eliminate unnecessary work, which may be identified through the client and filename breakdowns, and the filename hierarchy view. Client and especially filename breakdowns can be very expensive in terms of storage and execution overhead. Therefore, it is not recommended to permanently enable these breakdowns on a busy production appliance.

NFSv[2-4] Bytes Breakdowns

TABLE 43 Breakdowns of NFS Bytes

Breakdown	Description
type of operation	NFS operation type (read/write/getattr/setattr/lookup/...)
client	Remote hostname or IP address of the NFS client
filename	Filename for the NFS I/O, if known and cached by the appliance. There are some circumstances where the filename is not known, such as after a cluster failover and when clients continue to operate on NFS filehandles without issuing an open to identify the filename; in these situations the filename reported is "<unknown>".
Application ID	Identity of the client application issuing the I/O. This breakdown is available only for OISP-enabled NFSv4.0 and NFSv4.1 clients.
share	The share for this NFS I/O
project	The project for this NFS I/O

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: NFSv3 bytes per second for client 'phobos.sf.fishpong.com' broken down by filename" (to view which files a particular client is accessing)

Further Analysis

See [“Network: Device Bytes” on page 106](#) for a measure of network throughput caused by the NFS activity; [“Cache: ARC Accesses” on page 81](#) to learn how well an NFS read workload is returning from cache; and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: NFSv[2-4] Operations

This statistic shows NFSv[2-4] operations/sec (NFS IOPS) requested by clients to the appliance. Supported NFS versions are: NFSv2, NFSv3, NFSv4.0 and NFSv4.1. Various breakdowns are available to show the client, filename, and latency of the NFS I/O.

When to Check NFSv[2-4] Operations

NFSv[2-4] operations/sec can be used as an indication of NFS load, and can be viewed on the dashboard.

Use the latency breakdown when investigating NFS performance issues, especially to quantify the magnitude of the issue. This measures the I/O latency component for which the appliance is responsible for, and displays it as a heat map so that the overall latency pattern can be seen, along with outliers. If the NFS latency is high, drill down further on latency to identify the type of operation and filename for the high latency, and, check other statistics for both CPU and Disk load to investigate why the appliance is slow to respond; if latency is low, the appliance is performing quickly, and any performance issues experienced on the client are more likely to be caused by other factors in the environment: such as the network infrastructure, and CPU load on the client itself.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client and filename breakdowns, and the filename hierarchy view. Client and especially filename breakdowns can be very expensive in terms of storage and execution overhead. Therefore, it is not recommended to permanently enable these breakdowns on a busy production appliance.

NFSv[2-4] Operations Breakdowns

TABLE 44 Breakdowns of NFS Operations

Breakdown	Description
type of operation	NFS operation type (read/write/getattr/setattr/lookup/...)
client	Remote hostname or IP address of the NFS client.
filename	Filename for the NFS I/O, if known and cached by the appliance. There are some circumstances where the filename is not known, such as after a cluster failover and when clients continue to operate on NFS filehandles without issuing an open request to identify the filename; in these situations the filename reported is "<unknown>".
Application ID	Identity of the client application issuing the I/O. This breakdown is available only for OISP-enabled NFSv4.0 and NFSv4.1 clients.
share	The share for this NFS I/O.
project	The project for this NFS I/O.
latency	A heat map showing the latency of NFS I/O, as measured from when the NFS request arrived on the appliance from the network, to when the response is sent; this latency includes the time to process the NFS request, and to perform any disk I/O.
size	A heat map showing the distribution of NFS I/O sizes.

Breakdown	Description
offset	A heat map showing the file offset of NFS I/O. This can be used to identify random or sequential NFS IOPS. Use the Disk I/O operations statistic to check whether random NFS IOPS maps to random Disk IOPS after the filesystem and RAID configuration has been applied.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: NFSv3 operations per second of type read broken down by latency" (to examine latency for reads only)
- "Protocol: NFSv3 operations per second for file '/export/fs4/10ga' broken down by offset" (to examine file access pattern for a particular file)
- "Protocol: NFSv3 operations per second for client 'phobos.sf.fishpong.com' broken down by filename" (to view which files a particular client is accessing)

Further Analysis

See [“Network: Device Bytes” on page 106](#) for a measure of network throughput caused by the NFS activity; [“Cache: ARC Accesses” on page 81](#) to learn how well an NFS read workload is returning from cache; and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: OISP Bytes

This statistic shows OISP bytes/sec transferred between OISP clients and the appliance. Bytes statistics can be broken down by client, filename, database name, database filetype, database function, share, and project.

When to Check OISP Bytes

OISP bytes/sec can be used as an indication of OISP load, and can be viewed on the dashboard.

Breakdowns by database filetype and function allow database and storage administrators to correlate database statistics with storage statistics. This provides much better diagnosability to narrow down an abrupt rise not only to a particular database, but also to the database function creating the increase and the filetype associated with it.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client and filename breakdowns, and the filename hierarchy view. Client and especially filename breakdowns can very expensive in terms of storage and execution overhead. Therefore, it is not recommended to permanently enable these breakdowns on a busy production appliance.

OISP Bytes Breakdowns

TABLE 45 Breakdowns of OISP Bytes

Breakdown	Description
client	Remote hostname or IP address of the OISP client.
filename	Filename for the OISP I/O, if known and cached by the appliance.
database name	Name of the database issuing the I/O.
database filetype	Type of file to which the database is writing.
database function	Reason for the database I/O. Acronyms used in this breakdown include RMAN (Oracle Recovery Manager), DBWR (Database Writer for Oracle Database), ARCH (Archiver for Oracle Database), and LGWR (Log Writer for Oracle Database).
share	The share for this OISP I/O.
project	The project for this OISP I/O.

These breakdowns can be combined to produce powerful statistics. For example:

- “Protocol: OISP bytes per second for client ‘phobos.sf.fishpong.com’ broken down by filename” (to view which files a particular client is accessing)

Further Analysis

Also see: [“Protocol: OISP Operations” on page 120](#).

Protocol: OISP Operations

This statistic shows OISP operations/sec requested by clients to the appliance. Operations statistics can be broken down by client, filename, database name, database filetype, database function, share, project, latency, size, and offset.

When to Check OISP Operations

OISP operations/sec can be used as an indication of OISP load, and can be viewed on the dashboard.

Breakdowns by database filetype and function allow database and storage administrators to correlate database statistics with storage statistics. This provides much better diagnosability to narrow down an abrupt rise not only to a particular database, but also to the database function creating the increase and the filetype associated with it.

Use the latency breakdown when investigating OISP performance issues, especially to quantify the magnitude of the issue. This measures the I/O latency component for which the appliance is

responsible, and displays it as a heat map so that the overall latency pattern can be seen, along with outliers. If the OISP latency is high, drill down further on latency to identify the type of operation and filename for the high latency, and check other statistics for both CPU and disk load to investigate why the appliance is slow to respond. If latency is low, the appliance is performing quickly, and any performance issues experienced on the client are more likely to be caused by other factors in the environment, such as the network infrastructure, and CPU load on the client itself.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client and filename breakdowns, and the filename hierarchy view. Client and especially filename breakdowns can be very expensive in terms of storage and execution overhead. Therefore, it is not recommended to permanently enable these breakdowns on a busy production appliance.

OISP Operations Breakdowns

TABLE 46 Breakdowns of OISP Operations

Breakdown	Description
client	Remote hostname or IP address of the OISP client.
filename	Filename for the OISP I/O, if known and cached by the appliance.
database name	Name of the database issuing the I/O.
database filetype	Type of file to which the database is writing.
database function	Reason for the database I/O. Acronyms used in this breakdown include RMAN (Oracle Recovery Manager), DBWR (Database Writer for Oracle Database), ARCH (Archiver for Oracle Database), and LGWR (Log Writer for Oracle Database).
share	The share for this OISP I/O.
project	The project for this OISP I/O.
latency	A heat map showing the latency of OISP I/O, as measured from when the OISP request arrived on the appliance from the network, to when the response is sent. The latency includes the time to process the OISP request, and to perform any disk I/O.
size	A heat map showing the distribution of OISP I/O sizes.
offset	A heat map showing the file offset of OISP I/O. This can be used to identify random or sequential OISP.

These breakdowns can be combined to produce powerful statistics. For example:

- “Protocol: OISP operations per second for file ‘/export/fs4/10ga’ broken down by offset” (to examine file access pattern for a particular file)
- “Protocol: OISP operations per second for client ‘phobos.sf.fishpong.com’ broken down by filename” (to view which files a particular client is accessing)

Further Analysis

Also see Protocol: [“Protocol: OISP Bytes” on page 119](#).

Protocol: SFTP Bytes

This statistic shows SFTP bytes/sec requested by clients to the appliance. Various useful breakdowns are available: to show the client, user and filename of the SFTP requests.

Example

See [“Protocol: FTP Bytes” on page 112](#) for an example of a similar statistic with similar breakdowns.

When to Check SFTP Bytes

SFTP bytes/sec can be used as an indication of SFTP load, and can be viewed on the dashboard.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client, user and filename breakdowns, and the filename hierarchy view. It may be best to enable these breakdowns for short periods only: the by-filename breakdown can be one of the most expensive in terms of storage and execution overhead, and may not be suitable to leave enabled permanently on appliances with high rates of SFTP activity.

SFTP Bytes Breakdowns

TABLE 47 Breakdowns of SFTP Bytes

Breakdown	Description
type of operation	SFTP operation type (get/put/...)
user	Username of the client
filename	Filename for the SFTP operation, if known and cached by the appliance. If the filename is not known it is reported as "<unknown>".
share	The share for this SFTP request.
project	The project for this SFTP request.
client	Remote hostname or IP address of the SFTP client.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: SFTP bytes per second for client 'phobos.sf.fishpong.com' broken down by filename" (to view which files a particular client is accessing)

Further Analysis

See [“Cache: ARC Accesses” on page 81](#) to learn how well an SFTP read workload is returning from cache; and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Since SFTP uses SSH to encrypt FTP, there will be additional CPU overhead for this protocol. To check overall CPU utilization of the appliance, see [“CPU: Percent Utilization” on page 79](#).

Protocol: SMB/SMB2 Bytes

These statistics show SMB bytes/sec transferred between SMB clients and the appliance. Supported SMB versions are SMB and SMB2. Bytes statistics can be broken down by: operation, client, filename, share, and project.

When to Check SMB/SMB2 Bytes

SMB bytes/sec can be used as an indication of SMB load. The best way to improve performance is to eliminate unnecessary work, which may be identified through the client and filename breakdowns, and the filename hierarchy view. Client and especially filename breakdowns can be very expensive in terms of storage and execution overhead. Therefore, it is not recommended to permanently enable these breakdowns on a busy production appliance.

SMB/SMB2 Bytes Breakdowns

TABLE 48 Breakdowns of SMB Bytes

Breakdown	Description
type of operation	SMB/SMB2 operation type (read/write/getattr/setattr/lookup/...)
client	Remote hostname or IP address of the SMB client
filename	Filename for the SMB I/O, if known and cached by the appliance. There are some circumstances where the filename is not known, such as after a cluster failover and when clients continue to operate on SMB filehandles without issuing an open to identify the filename; in these situations the filename reported is "<unknown>".
share	The share for this SMB I/O
project	The project for this SMB I/O

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: SMB2 bytes per second for client 'phobos.sf.fishpong.com' broken down by filename" (to view which files a particular client is accessing)

Protocol: SRP Bytes

This statistic shows SRP bytes/sec requested by initiators to the appliance.

Example

See [“Protocol: iSCSI Bytes” on page 114](#) for an example of a similar statistic with similar breakdowns.

When to Check SRP Bytes

SRP bytes/sec can be used as an indication of SRP load, in terms of throughput. For a deeper analysis of SRP activity, see [“Protocol: SRP Operations” on page 124](#).

SRP Bytes Breakdowns

TABLE 49 Breakdowns of SRP Bytes

Breakdown	Description
initiator	SRP client initiator
target	Local SCSI target
project	The project for this SRP request.
lun	The LUN for this SRP request.

See [“Configuring Storage Area Network \(SAN\)” in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0*](#) for terminology definitions.

Further Analysis

See [“Protocol: SRP Operations” on page 124](#) for numerous other breakdowns on SRP operations; also see [“Cache: ARC Accesses” on page 81](#) to learn how well an SRP read workload is returning from cache, and [“Disk: I/O Operations” on page 104](#) for the back-end disk I/O caused.

Protocol: SRP Operations

This statistic shows SRP operations/sec (SRP IOPS) requested by initiators to the appliance. Various useful breakdowns are available: to show the initiator, target, type and latency of the SRP I/O.

Example

See [“Protocol: iSCSI Bytes” on page 114](#) for an example of a similar statistic with similar breakdowns.

When to Check SRP Operations

SRP operations/sec can be used as an indication of SRP load.

Use the latency breakdown when investigating SRP performance issues, especially to quantify the magnitude of the issue. This measures the I/O latency component for which the appliance is responsible for, and displays it as a heat map so that the overall latency pattern can be seen, along with outliers. If the SRP latency is high, drill down further on latency to identify the client initiator, the type of operation and LUN for the high latency, and, check other statistics for both CPU and Disk load to investigate why the appliance is slow to respond; if latency is low, the appliance is performing quickly, and any performance issues experienced on the client initiator are more likely to be caused by other factors in the environment: such as the network infrastructure, and CPU load on the client itself.

The best way to improve performance is to eliminate unnecessary work, which may be identified through the client initiator, lun and command breakdowns.

SRP Operations Breakdowns

TABLE 50 Breakdowns of SRP Operations

Breakdown	Description
initiator	SRP client initiator
target	Local SCSI target
project	The project for this SRP request.
lun	The LUN for this SRP request.
type of operation	SRP operation type. This shows how the SCSI command is transported by the SRP protocol, which can give an idea to the nature of the I/O.
command	SCSI command sent by the SRP protocol. This can show the real nature of the requested I/O (read/write/sync-cache/...).
latency	A heat map showing the latency of SRP I/O, as measured from when the SRP request arrived on the appliance from the network, to when the response is sent; this latency includes the time to process the SRP request, and to perform any disk I/O.
offset	A heat map showing the file offset of SRP I/O. This can be used to identify random or sequential SRP IOPS. Use the Disk I/O operations statistic to check whether random SRP IOPS maps to random Disk IOPS after the LUN and RAID configuration has been applied.
size	A heat map showing the distribution of SRP I/O sizes.

These breakdowns can be combined to produce powerful statistics. For example:

- "Protocol: SRP operations per second of command read broken down by latency" (to examine latency for SCSI reads only)

Further Analysis

See ["Protocol: SRP Bytes" on page 124](#) for the throughput of this SRP I/O; also see ["Cache: ARC Accesses" on page 81](#) to learn how well an SRP read workload is returning from cache, and ["Disk: I/O Operations" on page 104](#) for the back-end disk I/O caused.

Using Advanced Analytics Statistics

These statistics are only visible if Advanced Analytics is enabled in Preferences (see [“Setting Appliance Preferences”](#) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0*). These are statistics of lesser interest and are not typically needed for system observability. They are often dynamic which can induce higher overhead, and expose more complex areas of the system which require additional expertise to understand properly.

To learn more about advanced analytics statistics, use the following tasks:

- [“CPU: CPUs”](#) on page 128
- [“CPU: Kernel Spins”](#) on page 129
- [“Cache: ARC Adaptive Parameter”](#) on page 129
- [“Cache: ARC Evicted Bytes”](#) on page 130
- [“Cache: ARC Size”](#) on page 130
- [“Cache: ARC Target Size”](#) on page 131
- [“Cache: DNLC Accesses”](#) on page 132
- [“Cache: DNLC Entries”](#) on page 132
- [“Cache: L2ARC Errors”](#) on page 133
- [“Cache: L2ARC Size”](#) on page 134
- [“Data Movement: NDMP Bytes Transferred to/from Disk”](#) on page 134
- [“Data Movement: NDMP Bytes Transferred to/from Tape”](#) on page 135
- [“Data Movement: NDMP File System Operations”](#) on page 135
- [“Data Movement: NDMP Jobs”](#) on page 136
- [“Data Movement: Replication Latencies”](#) on page 136
- [“Data Movement: Replication Send/Receive Bytes”](#) on page 137
- [“Disk: Average Number of I/O Operations”](#) on page 138
- [“Disk: Percent Utilization”](#) on page 138
- [“Disk: ZFS DMU Operations”](#) on page 139
- [“Disk: ZFS Logical I/O Bytes”](#) on page 140
- [“Disk: ZFS Logical I/O Operations”](#) on page 140
- [“Memory: Dynamic Memory Usage”](#) on page 141
- [“Memory: Kernel Memory”](#) on page 141

- “Memory: Kernel Memory in Use” on page 142
- “Memory: Kernel Memory Allocated, But Not in Use” on page 142
- “Network: Datalink Bytes” on page 143
- “Network: IP Bytes” on page 144
- “Network: IP Packets” on page 144
- “Network: TCP Bytes” on page 145
- “Network: TCP Packets” on page 145
- “Network: TCP Retransmissions” on page 146
- “System: NSCD Backend Requests” on page 146
- “System: NSCD Operations” on page 147

CPU: CPUs

The CPUs statistic is used to display the heat map for CPUs broken down by percent utilization. This is the most accurate way to examine how CPUs are utilized.

When to Check CPUs

When investigating CPU load, after checking the utilization average from CPU Percent Utilization.

This statistic is particularly useful for identifying if a single CPU is fully utilized, which can happen if a single thread is saturated with load. If the work performed by this thread cannot be offloaded to other threads so that it can be run concurrently across multiple CPUs, then that single CPU can become the bottleneck. This will be seen as a single CPU stuck at 100% utilization for several seconds or more, while the other CPUs are idle.

CPUs Breakdowns

TABLE 51 A Breakdown of CPUs

Breakdown	Description
percent utilization	A heat map with utilization on the Y-axis and each level on the Y-axis colored by the number of CPU at that utilization: from light (none) to dark (many).

Details

CPU utilization includes the time to process instructions (that are not part of the idle thread); which includes memory stall cycles. CPU utilization can be caused by:

- Executing code (including spinning on locks)

- Memory load

Since the appliance primarily exists to move data, memory load often dominates. So a system with high CPU utilization may actually be high as it is moving data.

CPU: Kernel Spins

This statistic counts the number of spin cycles on kernel locks, which consume CPU.

An understanding of operating system internals is required to properly interpret this statistic.

When to Check Kernel Spins

When investigating CPU load, after checking CPU Percent Utilization and CPUs broken down by percent utilization.

Some degree of kernel spins is normal for processing any workload, due to the nature of multi-threaded programming. Compare the behavior of kernel spins over time, and for different workloads, to develop an expectation for what is normal.

Kernel Spins Breakdowns

TABLE 52 Breakdowns of CPU Kernel Spins

Breakdown	Description
type of synchronization primitive	Type of lock (mutex/...)
CPU identifier	CPU identifier number (0/1/2/3/...)

Cache: ARC Adaptive Parameter

This is `arc_p` from the ZFS ARC. This shows how the ARC is adapting its MRU and MFU list size depending on the workload.

An understanding of ZFS ARC internals may be required to properly interpret this statistic.

When to Check ARC Adaptive Parameter

Rarely; this may be useful for identifying internal behavior of the ARC, however there are other statistics to check before this one.

If there are caching issues on the appliance, check the Cache ARC accesses statistic to see how well the ARC is performing, and the Protocol statistics to understand the requested workload.

Then, check the Advanced Analytics Cache ARC size and Cache ARC evicted bytes for further details on the ARC behavior.

ARC Adaptive Parameter Breakdowns

None.

Cache: ARC Evicted Bytes

This statistic shows bytes that were evicted from the ZFS ARC, as part of its usual housekeeping. The breakdown allows L2ARC eligibility to be examined.

An understanding of ZFS ARC internals may be required to properly interpret this statistic.

When to Check ARC Evicted Bytes

This could be checked if you were considering to install cache devices (L2ARC), as this statistic can be broken down by L2ARC state. If L2ARC eligible data was frequently being evicted from the ARC, then the presence of cache devices could improve performance.

This may also be useful to check if you have issues with cache device warmup. The reason may be that your workload is not L2ARC eligible.

If there are caching issues on the appliance, check the Cache ARC accesses statistic to see how well the ARC is performing, and the Protocol statistics to understand the requested workload. Then, check the Advanced Analytics Cache ARC size and Cache ARC evicted bytes for further details on the ARC behavior.

ARC Evicted Bytes Breakdowns

TABLE 53 A Breakdown of ARC Evicted Bytes

Breakdown	Description
L2ARC state	Shows L2ARC cached or uncached, L2ARC eligible or ineligible.

Cache: ARC Size

This statistic shows the size of the primary filesystem cache, the DRAM based ZFS ARC.

An understanding of ZFS ARC internals may be required to properly interpret this statistic.

When to Check ARC Size

When examining the effectiveness of the ARC on the current workload. The ARC should automatically increase in size to fill most of available DRAM, when enough data be accessed by the current workload to be placed in the cache. The breakdown allows the contents of the ARC to be identified by type.

This may also be checked when using cache devices (L2ARC) on systems with limited DRAM, as the ARC can become consumed with L2ARC headers.

If there are ARC caching issues on the appliance, also check the Cache ARC accesses statistic to see how well the ARC is performing, and the Protocol statistics to understand the requested workload.

ARC Size Breakdowns

Available breakdowns:

TABLE 54 A Breakdown of ARC Size

Breakdown	Description
component	Type of data in the ARC. See table below

ARC component types:

TABLE 55 ARC Component Types

Component	Description
ARC data	Cached contents, including filesystem data and filesystem metadata.
ARC headers	Space consumed by metadata of the ARC itself. The ratio of headers to data is relative to the ZFS record size used; a small record size may mean more ARC headers to refer to the same volume.
ARC other	Other kernel consumers of the ARC
L2ARC headers	Space consumed by tracking buffers stored on L2ARC devices. If the buffer is on the L2ARC and yet still in ARC DRAM, it is considered "ARC headers" instead.

Cache: ARC Target Size

This is `arc_c` from the ZFS ARC. This shows the target size that the ARC is attempting to maintain. For the actual size, see the Advanced Analytic [“Cache: ARC Size” on page 130](#).

An understanding of ZFS ARC internals may be required to properly interpret this statistic.

When to Check ARC Target Size

Rarely; this may be useful for identifying internal behavior of the ARC, however there are other statistics to check before this one.

If there are caching issues on the appliance, check the Cache ARC accesses statistic to see how well the ARC is performing, and the Protocol statistics to understand the requested workload. Then, check the Advanced Analytics Cache ARC size and Cache ARC evicted bytes for further details on the ARC behavior.

ARC Target Size Breakdowns

None.

Cache: DNLC Accesses

This statistic shows accesses to the DNLC (Directory Name Lookup Cache). The DNLC caches pathname to inode lookups.

An understanding of operating system internals may be required to properly interpret this statistic.

When to Check DNLC Accesses

This may be useful to check if a workload accesses millions of small files, for which the DNLC can help.

If there are generic caching issues on the appliance, first check the Cache ARC accesses statistic to see how well the ARC is performing, and the Protocol statistics to understand the requested workload. Then, check the Advanced Analytic Cache ARC size for the size of the ARC.

DNLC Accesses Breakdowns

TABLE 56 A Breakdown of DNLC Accesses

Breakdown	Description
hit/miss	Shows counts for hits/misses, allowing the effectiveness of the DNLC to be checked.

Cache: DNLC Entries

This shows the number of entries in the DNLC (Directory Name Lookup Cache). The DNLC caches pathname to inode lookups.

An understanding of operating system internals may be required to properly interpret this statistic.

When to Check DNLC Entries

This may be useful to check if a workload accesses millions of small files, for which the DNLC can help.

If there are generic caching issues on the appliance, first check the Cache ARC accesses statistic to see how well the ARC is performing, and the Protocol statistics to understand the requested workload. Then, check the Advanced Analytic Cache ARC size for the size of the ARC.

DNLC Entries Breakdowns

None.

Cache: L2ARC Errors

This statistic shows L2ARC error statistics.

When to Check L2ARC Errors

This may be useful to leave enabled when using cache devices, for when troubleshooting L2ARC issues beyond the standard statistics.

L2ARC Errors Breakdowns

Available breakdowns:

TABLE 57 A Breakdown of L2ARC Errors

Breakdown	Description
error	L2ARC error type. See table below.

L2ARC error types:

TABLE 58 L2ARC Error Types

Error	Description
memory abort	The L2ARC choose not to populate for a one second interval due to a shortage of system memory (DRAM)

Error	Description
	which holds the L2ARC metadata. Continual memory aborts will prevent the L2ARC from warming up.
bad checksum	A read from a cache device failed the ZFS ARC checksum. This may be an indicator that a cache device is beginning to fail.
io error	A cache device returned an error. This may be an indicator that a cache device is beginning to fail.

Cache: L2ARC Size

This shows the size of data stored on the L2ARC cache devices. This is expected to increase in size over a period of hours or days, until the amount of amount of constant L2ARC eligible data is cached, or the cache devices are full.

When to Check L2ARC Size

When troubleshooting L2ARC warmup. If the size is small, check that the workload applied should be populating the L2ARC using the statistic Cache ARC evicted bytes broken down by L2ARC state, and use the Protocol breakdowns such as by size and by offset to confirm that the workload is of random I/O. Sequential I/O does not populate the L2ARC. Another statistic to check is Cache L2ARC errors.

The L2ARC size does shrink, if data that was cached is deleted from the filesystem.

L2ARC Size Breakdowns

None.

Data Movement: NDMP Bytes Transferred to/from Disk

This statistic indicates disk throughput during a NDMP backup or restore operation. This statistic is zero unless NDMP is configured and active.

When to Check NDMP Bytes Transferred to/from Disk

When investigating NDMP backup and restore performance. This can also be checked when trying to identify an unknown disk load, some of which may be caused by NDMP.

NDMP Bytes Transferred to/from Disk Breakdowns

TABLE 59 Breakdowns of NDMP Bytes Transferred to/from Disk

Breakdown	Description
type of operation	Read or write
raw statistic	Raw statistic

Further Analysis

Also see [“Data Movement: NDMP Bytes Transferred to/from Tape”](#) on page 135.

Data Movement: NDMP Bytes Transferred to/from Tape

This statistic indicates tape throughput during a NDMP backup or restore operation. This statistic is zero unless NDMP is configured and active.

When to Check NDMP Bytes Transferred to/from Tape

When investigating NDMP backup and restore performance.

NDMP Bytes Transferred to/from Tape Breakdowns**TABLE 60** Breakdowns of NDMP Bytes Transferred to/from Tape

Breakdown	Description
type of operation	Read or write
raw statistic	Raw statistic

Further Analysis

Also see [“Data Movement: NDMP Bytes Transferred to/from Disk”](#) on page 134.

Data Movement: NDMP File System Operations

This statistic shows access per second to the file system during NDMP backup or restore. This statistic is only meaningful for tar-based backups because they occur at the file, and not block, level.

When to Check NDMP File System Operations

This could be useful to check when investigating the source of ZFS load. First, check all other sources of file system activity via the Protocol statistics. Also see the advanced Analytics statistics [“Data Movement: NDMP Bytes Transferred to/from Disk” on page 134](#) and [“Data Movement: NDMP Bytes Transferred to/from Tape” on page 135](#).

NDMP File System Operations Breakdowns

TABLE 61 Breakdowns of NDMP File System Operations

Breakdown	Description
type of operation	Read or write
raw statistic	Raw statistic

Data Movement: NDMP Jobs

This statistic shows active NDMP job counts.

When to Check NDMP Jobs

When monitoring NDMP progress, and troubleshooting NDMP. Also see the standard Analytics statistic [“Data Movement: NDMP Bytes Transferred to/from Disk” on page 134](#) and [“Data Movement: NDMP Bytes Transferred to/from Tape” on page 135](#).

NDMP Jobs Breakdowns

TABLE 62 A Breakdown of NDMP Jobs

Breakdown	Description
type of operation	Type of job: backup/restore.

Data Movement: Replication Latencies

This statistic shows the average latency per second as a single value per unit of time instead of a heat map which shows multiple values per unit of time. In many cases this statistic provides sufficient detail without going to the heat map. Typically heat-map-based statistics are more expensive.

When to Check Replication Latencies

When monitoring replication progress, and troubleshooting replication. Also see the standard Analytics statistic [“Data Movement: Replication Bytes” on page 99](#) and [“Data Movement: Replication Operations” on page 100](#).

Replication Latencies Breakdowns

TABLE 63 Breakdowns of Replication Latencies

Breakdown	Description
direction	Shows latency broken down by direction, to or from the appliance.
type of operation	Shows latency broken down by type of operation with the remote appliance, read or write.
peer	Shows latency broken down by the name of the remote appliances.
pool name	Shows latency broken down by the name of the pools.
project	Shows latency broken down by the name of the projects.
dataset	Shows latency broken down by the name of the shares.
as a raw statistic	Shows latency as raw statistics.

Data Movement: Replication Send/Receive Bytes

This statistic tracks data throughput of project/share replication in bytes per second at the ZFS send/receive internal interface. This statistic omits the effect of the compression/decompression stage in the replication data pipeline.

When to Check Replication Send/Receive Bytes

When investigating replication activity or evaluating the benefit of enabling compression for a replication action.

Replication Send/Receive Bytes Breakdowns

TABLE 64 Breakdowns of Replication Send/Receive Bytes

Breakdown	Description
direction	Shows send/receive bytes broken down by direction, to or from the appliance.
type of operation	Shows send/receive bytes broken down by the type of operation with the remote appliance, read or write.
peer	Shows send/receive bytes broken down by the name of the remote appliances.

Breakdown	Description
pool name	Shows send/receive bytes broken down by the name of the pools.
project	Shows send/receive bytes broken down by the name of the projects.
dataset	Shows send/receive bytes broken down by the name of the shares.
as a raw statistic	Shows send/receive bytes as raw statistics.

Further Analysis

Also see [“Data Movement: Replication Bytes” on page 99](#) for replication data throughput after any compression in the replication data pipeline.

Disk: Average Number of I/O Operations

This statistic shows the average number of disk I/O operations per second.

When to Check Average Number of I/O Operations

This statistic shows how busy the I/O subsystem is on the appliance. High wait times could indicate that you need more disks or a different RAID profile.

Average Number of I/O Operations Breakdowns

TABLE 65 Breakdowns of Average Number of Disk Operations

Breakdown	Description
state of operation	Average number of I/O operations for active and waiting disks. Drill down on a selected state to examine further.
by disk	Average number of I/O operations for system and data disks. Drill down on a selected disk to examine further.
as a raw statistic	Average number of I/O operations as raw statistics.

Disk: Percent Utilization

This statistic shows average utilization across all disks. The per-disk breakdown shows the utilization that that disk contributed to the total average, not the utilization of that disk.

When to Check Percent Utilization

This statistic may be useful to trigger an alert based on the average for all disks.

Investigating disk utilization is usually much more effective using the standard Analytics statistic [“Disk: Disks” on page 101](#) broken down by percent utilization - which instead of averaging utilization, presents it as a heat map. This allows individual disk utilization to be examined.

Percent Utilization Breakdowns

TABLE 66 A Breakdown of Percent Utilization

Breakdown	Description
disk	Disks, including system and pool disks

The disk breakdown shows the contribution to the average percent which each disk made.

Notes

A system with 100 disks would never show more than 1 for any disk breakdown, unless that disk was selected and displayed separately as a raw statistic. Such a system would also show 0 percent utilization for disks less than 50% busy, due to rounding. Since this may be a source of confusion, and that there is a better statistic available for most situations ([Disk: Disks](#)), this statistic has been placed in the Advanced category.

See [“Disk: Disks” on page 101](#) for a different and usually more effective way to display this data.

Disk: ZFS DMU Operations

This statistic shows ZFS DMU (Data Management Unit) operations/sec.

An understanding of ZFS internals is required to properly interpret this statistic.

When to Check ZFS DMU Operations

Troubleshooting performance issues, after all relevant standard Analytics have been examined.

The DMU object type breakdown can identify if there is excessive DDT (Data Deduplication Table) activity. For more information about data deduplication, see [“About the Oracle ZFS Storage Appliance” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#).

ZFS DMU Operations Breakdowns

TABLE 67 Breakdowns of ZFS DMU Operations

Breakdown	Description
type of operation	Read/write/...
DMU object level	Integer
DMU object type	ZFS plain file/ZFS directory/DMU dnode/SPA space map/...

Disk: ZFS Logical I/O Bytes

This statistic shows logical access to the ZFS file system as bytes/sec. Logical I/O refers to the type of operations as those that are requested to the file system, such as by NFS; as opposed to physical I/O, which are the requests by the file system to the back-end pool disks.

When to Check ZFS Logical I/O Bytes

This could be useful while investigating how I/O is processed between the Protocol layer and pool disks.

ZFS Logical I/O Bytes Breakdowns

TABLE 68 Breakdowns of ZFS Logical I/O Bytes

Breakdown	Description
type of operation	Read/write/...
pool name	Name of the disk pool

Disk: ZFS Logical I/O Operations

This statistic shows logical access to the ZFS file system as operations/sec. Logical I/O refers to the type of operations as those that are requested to the file system, such as by NFS; as opposed to physical I/O, which are the requests by the file system to the back-end pool disks.

When to Check ZFS Logical I/O Operations

This could be useful while investigating how I/O is processed between the Protocol layer and pool disks.

ZFS Logical I/O Operations Breakdowns

TABLE 69 Breakdowns of ZFS Logical I/O Operations

Breakdown	Description
type of operation	Read/write/...
pool name	Name of the disk pool

Memory: Dynamic Memory Usage

This statistic gives a high level view of memory (DRAM) consumers, updated every second.

When to Check Dynamic Memory Usage

This can be used to check that the filesystem cache has grown to consume available memory.

Dynamic Memory Usage Breakdowns

Available breakdowns:

TABLE 70 A Breakdown of Dynamic Memory Usage

Breakdown	Description
application name	See table below.

Application names:

TABLE 71 Application Names

Application Name	Description
cache	The ZFS filesystem cache (ARC). This will grow to consume as much of available memory as possible, as it caches frequently accessed data.
kernel	The operating system kernel.
mgmt	The appliance management software.
unused	Unused space.

Memory: Kernel Memory

This statistic shows kernel memory which is allocated, and can be broken down by kernel cache (kmem cache).

An understanding of operating system internals is required to understand this statistic.

When to Check Kernel Memory

Rarely. If the dashboard were to show kernel memory as a large consumer of available DRAM (in the Usage: Memory section), then this may be used when troubleshooting the cause. Also see [“Memory: Kernel Memory in Use” on page 142](#) and [“Memory: Kernel Memory Allocated, But Not in Use” on page 142](#).

Kernel Memory Breakdowns

TABLE 72 A Breakdown of Kernel Memory

Breakdown	Description
kmem cache	Kernel memory cache name.

Memory: Kernel Memory in Use

This statistic shows kernel memory which is in use (populated), and can be broken down by kernel cache (kmem cache).

An understanding of operating system internals is required to understand this statistic.

When to Check Kernel Memory in Use

Rarely. If the dashboard were to show kernel memory as a large consumer of available DRAM (in the Usage: Memory section), then this may be used when troubleshooting the cause. Also see [“Memory: Kernel Memory Allocated, But Not in Use” on page 142](#).

Kernel Memory in Use Breakdowns

TABLE 73 A Breakdown of Kernel Memory in Use

Breakdown	Description
kmem cache	Kernel memory cache name.

Memory: Kernel Memory Allocated, But Not in Use

This statistic shows kernel memory which is allocated, but not in use, and can be broken down by kernel cache (kmem cache). Such a state can occur when memory is freed (for example, when cached file system data is deleted), and the kernel has yet to recover the memory buffers.

An understanding of operating system internals is required to understand this statistic.

When to Check Kernel Memory Allocated, But Not in Use

Rarely. If the dashboard were to show kernel memory as a large consumer of available DRAM (in the Usage: Memory section), then this may be used when troubleshooting the cause. Also see [“Memory: Kernel Memory in Use” on page 142](#).

Kernel Memory Allocated, But Not in Use Breakdowns

TABLE 74 A Breakdown of Kernel Memory Allocated, But Not in Use

Breakdown	Description
kmem cache	Kernel memory cache name.

Network: Datalink Bytes

This statistic measures network datalink activity in bytes/sec. Network datalinks are logical entities constructed from network devices (see [“Network Configuration” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#)). The measured bytes by this statistic includes all network payload headers (Ethernet, IP, TCP, NFS/SMB/etc.)

Example

See [“Network: Device Bytes” on page 106](#) for an example of a similar statistic with similar breakdowns.

When to Check Datalink Bytes

Network bytes can be used a rough measure of appliance load. This statistic can be used to see the rate of network bytes through different datalinks.

Datalink Bytes Breakdowns

TABLE 75 Breakdowns of Datalink Bytes

Breakdown	Description
direction	In or out, relative to the appliance. For example, NFS reads to the appliance would be show as out(bound) network bytes.
datalink	Network datalink (see Datalink in Network).

Further Analysis

Also see [“Network: Device Bytes” on page 106](#) and [“Network: Interface Bytes” on page 107](#) for network throughput at the device and interface level, respectively.

Network: IP Bytes

This statistic shows IP payload bytes/second, excluding the Ethernet/IB and IP headers.

When to Check IP Bytes

Rarely. Network throughput monitoring can be achieved using the standard Analytics statistic Network Device bytes, which is enabled and achieved by default. Examining by-client throughput can usually be achieved through the Protocol statistic (for example, Protocol iSCSI bytes, which allows other useful breakdowns based on the protocol). This statistic is most useful if the previous two were not appropriate for some reason.

IP Bytes Breakdowns

TABLE 76 Breakdowns of IP Bytes

Breakdown	Description
hostname	Remote client, either as a hostname or IP address
protocol	IP protocol: tcp/udp
direction	Relative to the appliance. in/out

Network: IP Packets

This statistic shows IP packets/second.

When to Check IP Packets

Rarely. Since packets usually map to protocol operations, it is often more useful to examine these using the Protocol statistics (for example, Protocol iSCSI operations, which allows other useful breakdowns based on the protocol).

IP Packets Breakdowns

TABLE 77 Breakdowns of IP Packets

Breakdown	Description
hostname	Remote client, either as a hostname or IP address

Breakdown	Description
protocol	IP protocol: tcp/udp
direction	Relative to the appliance. in/out

Network: TCP Bytes

This statistic shows TCP payload bytes/second, excluding the Ethernet/IB, IP and TCP headers.

When to Check TCP Bytes

Rarely. Network throughput monitoring can be achieved using the standard Analytics statistic Network Device bytes, which is enabled and achieved by default. Examining by-client throughput can usually be achieved through the Protocol statistic (for example, Protocol iSCSI bytes, which allows other useful breakdowns based on the protocol). This statistic is most useful if the previous two were not appropriate for some reason.

TCP Bytes Breakdowns

TABLE 78 Breakdowns of TCP Bytes

Breakdown	Description
client	Remote client, either as a hostname or IP address
local service	TCP port: http/ssh/215(administration)/...
direction	Relative to the appliance, in/out
local IP address	The appliance IP address that packets are sent from or received on

Network: TCP Packets

This statistic shows TCP packets/second.

When to Check TCP Packets

Rarely. Since packets usually map to protocol operations, it is often more useful to examine these using the Protocol statistics (for example, Protocol iSCSI operations, which allows other useful breakdowns based on the protocol).

TCP Packets Breakdowns

TABLE 79 Breakdowns of TCP Packets

Breakdown	Description
client	Remote client, either as a hostname or IP address
local service	TCP port: http/ssh/215(administration)/...
direction	Relative to the appliance, in/out
local IP address	The appliance IP address that packets are sent from or received on
packet size	Size of the transmitted packet
latency	Time between when a packet is sent and when the ack is received

Network: TCP Retransmissions

This statistic shows TCP retransmissions.

When to Check TCP Retransmissions

Rarely. Since packets usually map to protocol operations, it is often more useful to examine these using the Protocol statistics. However, certain network types, as well as client receive buffers sizes, are more prone to TCP retransmissions, which can manifest as high latency.

TCP Retransmissions Breakdowns

TABLE 80 Breakdowns of TCP Retransmissions

Breakdown	Description
client	Remote client, either as a hostname or IP address
local service	TCP port: http/ssh/215(administration)/...
local IP address	The appliance IP address that packets are sent from or received on
packet size	Size of the transmitted packet

System: NSCD Backend Requests

This statistic shows requests made by NSCD (Name Service Cache Daemon) to back-end sources, such as DNS, NIS, etc.

An understanding of operating system internals may be required to properly interpret this statistic.

When to Check NSCD Backend Requests

It may be useful to check the latency breakdown if long latencies were experienced on the appliance, especially during administrative logins. The breakdowns for the database name and source will show what the latency is for, and which remote server is responsible.

NSCD Backend Requests Breakdowns

TABLE 81 Breakdowns of NSCD Backend Requests

Breakdown	Description
type of operation	Request type
result	Success/fail
database name	NSCD database (DNS/NIS/...)
source	Hostname or IP address of this request
latency	Time for this request to complete

System: NSCD Operations

This statistic shows requests made to NSCD (Name Service Cache Daemon).

An understanding of operating system internals may be required to properly interpret this statistic.

When to Check NSCD Operations

This can be used to check the effectiveness of the NSCD cache, by using the hit/miss breakdown. Misses become backend requests to remote sources, which can be examined using System NSCD backend requests.

NSCD Operations Breakdowns

TABLE 82 Breakdowns of NSCD Operations

Breakdown	Description
type of operation	Request type
result	Success/fail
database name	NSCD database (DNS/NIS/...)
latency	Time for this request to complete
hit/miss	Result from the cache lookup: hit/miss

