



Cisco Mobility Unified Reporting System Installation and Administration Guide

Version 14.0

Last Updated April 30, 2015

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-27216-09

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Mobility Unified Reporting System Installation and Administration Guide

© 2015 Cisco Systems, Inc. All rights reserved.

CONTENTS

About this Guide	vii
Conventions Used	viii
Contacting Customer Support	ix
Additional Information	x
Mobility Unified Reporting System Overview	11
Introduction	12
Report Types	13
Exporting Reports to Other File Formats	20
License Requirements	20
MUR Architecture	21
Distributed Architecture of MUR	25
How RDP works with MUR	26
MUR Features	28
Clustering Support for High Availability	28
Operation	28
HTTPS Access	28
Creation of Security Certificates	29
Implementation on RHEL	30
Implementation on Solaris	31
LDAP Authentication in MUR	32
Region-based Reporting	33
Load Distribution Based on Number of Files	35
Tethering Detection Feature	36
MUR Support for Tethering Detection	36
Tethering Detection Databases	37
Loading and Upgrading Tethering Detection Databases	38
MUR Deployment	39
MUR System Requirements	41
Server Recommendations for Use in Solaris Environment	41
Server Recommendations for Use in RHEL Environment	42
Storage RAID recommendation for MUR Application	43
Storage Recommendation for MUR Application	43
Hardware Requirements for Scalable Model of MUR	43
Software Requirements for Scalable Model of MUR	43
MUR Ports	45
Firewall Settings	45
Using Apache Port	46
Using Apache in Solaris	46
Using Apache in RHEL	46
Configuring Chassis for Mobility Unified Reporting System	47
Initial Configuration	48
Installing the ECS License	48
Creating the ECS Administrative User Account	48
Enabling Active Charging	49

Creating the Active Charging Service	49
Configuration	50
Activating P2P Analyzer	50
Configuring the EDR Flow Format	50
Verifying your Configuration.....	52
Configuring Deep Packet Inspection	53
Configuring Routing Rule Definition	53
Configuring Rulebase.....	55
Configuring Charging Action	58
Configuring Tethering Detection Feature	58
Upgrading Tethering Detection Databases.....	59
Sample Configurations.....	59
EDR Module Configuration.....	63
Verifying your Configuration.....	65
Pushing EDR/UDR Files Manually	65
Configuring EDR Download Permission.....	66
Configuring Bulkstats Schemas Using GUI.....	66
Supported Bulkstat Schemas	68
Supported SNMP Traps.....	71
Configuring MegaRAID for MUR Applications.....	75
Option Recommendations	76
Sample Configuration	77
Creating Virtual Drives.....	77
Setting the Boot Drive	81
Cisco UCS Server Hardware Configuration for MUR Applications	83
Overview	84
Prerequisites.....	85
Storage Recommendations	86
Disk Partitioning	87
Creating Volume Groups and Partitions	88
Configuring the XFS File System	91
Mobility Unified Reporting System Clustering Support for High Availability	93
System and Hardware Recommendations for HA Deployment	94
Configuring the External Storage Disk on UCS for HA Deployment	95
Tuning the VxFS File System	98
Configuring Resources for High Availability	100
Recovering MUR in HA.....	104
Managing Mobility Unified Reporting System Installation	105
Installing MUR	106
Setting the Database Environment Strings	107
Settings for Solaris	107
Settings for RHEL	108
Pre-installation Checks.....	108
MUR Installation	110
Installing MUR Using GUI/Console based Installer	111
Confirming Successful Installation	117
Upgrading MUR	118
Uninstalling MUR	121
Uninstallation Using GUI/Console-based Uninstaller.....	121
Mobility Unified Reporting System Administration and Management.....	123
Launching the MUR GUI.....	124

Administration	126
Managing User Accounts	126
Managing Gateways.....	126
Modifying Gateway in GUI in Hierarchical Setup	127
Managing Archive Directory	128
Configuring Logging	128
Configuring Purging Feature	128
Configuring Backup Functionality.....	129
Configuring Recovery Functionality	130
Configuring Offline Mode	130
Operations and Management.....	131
Generating Reports in Excel Format.....	131
Generating Unknown URL Files.....	131
Loading Blacklist and Whitelist File for Tethering Detection	132
Loading HTTP Group Data Using csv File	132
Purging Tethering Database	133
Resetting GUI Administrator User Password	134
Using the generate_dns_mapp_sql.sh Script	134
Using the getSupportDetails Script	135
Requirements.....	135
Supported Levels	136
Using the Maintenance Utility.....	137
Using the PSMON Script.....	139
Using the Purging Script	139
Using the unanonymize_msisdn.sh Script	139
Server Script Parameters.....	140
Scalable Solution for MUR.....	141
Scalable MUR Overview.....	142
Basic Scalability Model.....	142
System Requirements for Scalable MUR.....	144
Hardware Requirements	144
Software Requirements.....	144
Deployment of Scalable MUR	145
Hardware Configurations.....	146
Heartbeat Link between RDP Hosts	146
Installing RHEL on RDP Hosts.....	146
Configurations Post RHEL Installation.....	146
Disk Partitioning for RDP	147
Disk Partitioning for MUR Master	147
Cabling between StorageTek and RDP	159
Configuring StorageTek	160
Installing the Management Software (CAM)	161
Incoming Data Partition.....	176
Archival Data Partition	176
Configuring Multipaths on RDP Hosts.....	177
Configuring Multipath Service	177
Configuring Veritas.....	180
Hardware Setup for Veritas Cluster File System	180
Installing Veritas Cluster File System	180
Configuring Shared Volume.....	375
I/O Fencing.....	379
Understanding Split Brain and the Need for I/O Fencing	379
Configuring Disk-based I/O Fencing Using installsfcfs.....	380
Verifying I/O Fencing Configuration.....	392

Testing I/O Fencing Configuration	392
Configurations in MUR for Clustering Support	394
Configuring MUR for Use in Scalable model	394
Post Installation Actions	400
Installation / Upgrade to Scalable Model	401
Upgrade from Standalone Deployment to Scalable Model	401
Upgrade from Hierarchical Deployment to Scalable Model	402
Troubleshooting MUR System	405
MUR Preventive and Control Measures	406
MUR Install and Upgrade Prerequisites	407
Issues Pertaining to MUR Installation / Upgrade	408
Issues Pertaining to MUR Processes	409
Issues Related to MUR GUI	412
Issues Related to Master-RDP Communication	415
Issues Related to BS-KPI Alarms	418
Issues Related to Backup and Restore	420
Miscellaneous Issues	421

About this Guide

This document pertains to the features and functionality that are related to the Cisco® ESS Installation and Administration Guide.

Conventions Used

The following tables describe the conventions used throughout this documentation.

Icon	Notice Type	Description
	Information Note	Provides information about important features or instructions.
	Caution	Alerts you of potential damage to a program, device, or system.
	Warning	Alerts you of potential personal injury or fatality. May also alert you of potential electrical hazards.

Typeface Conventions	Description
Text represented as a <code>screen display</code>	This typeface represents displays that appear on your terminal screen, for example: <code>Login:</code>
Text represented as <code>commands</code>	This typeface represents commands that you enter, for example: <code>show ip access-list</code> This document always gives the full form of a command in lowercase letters. Commands are not case sensitive.
Text represented as a <code>command variable</code>	This typeface represents a variable that is part of a command, for example: <code>show card slot_number</code> <code>slot_number</code> is a variable representing the desired chassis slot number.
Text represented as menu or sub-menu names	This typeface represents menus and sub-menus that you access within a software application, for example: Click the File menu, then click New

Contacting Customer Support

Use the information in this section to contact customer support.

Refer to the support area of <http://www.cisco.com> for up-to-date product documentation or to submit a service request. A valid username and password are required to access this site. Please contact your Cisco sales or service representative for additional information.

Additional Information

Refer to the following guides for supplemental information about the system:

- *Cisco ASR 5000 Installation Guide*
- *Cisco ASR 5000 System Administration Guide*
- *Cisco ASR 5x00 Command Line Interface Reference*
- *Cisco ASR 5x00 Thresholding Configuration Guide*
- *Cisco ASR 5x00 SNMP MIB Reference*
- *Web Element Manager Installation and Administration Guide*
- *Cisco ASR 5x00 AAA Interface Administration and Reference*
- *Cisco ASR 5x00 GTPP Interface Administration and Reference*
- *Cisco ASR 5x00 Release Change Reference*
- *Cisco ASR 5x00 Statistics and Counters Reference*
- *Cisco ASR 5x00 Gateway GPRS Support Node Administration Guide*
- *Cisco ASR 5x00 HRPD Serving Gateway Administration Guide*
- *Cisco ASR 5000 IP Services Gateway Administration Guide*
- *Cisco ASR 5x00 Mobility Management Entity Administration Guide*
- *Cisco ASR 5x00 Packet Data Network Gateway Administration Guide*
- *Cisco ASR 5x00 Packet Data Serving Node Administration Guide*
- *Cisco ASR 5x00 System Architecture Evolution Gateway Administration Guide*
- *Cisco ASR 5x00 Serving GPRS Support Node Administration Guide*
- *Cisco ASR 5x00 Serving Gateway Administration Guide*
- *Cisco ASR 5000 Session Control Manager Administration Guide*
- *Cisco ASR 5000 Packet Data Gateway/Tunnel Termination Gateway Administration Guide*
- Release notes that accompany updates and upgrades to the StarOS for your service and platform

Chapter 1

Mobility Unified Reporting System Overview

This chapter provides an overview of the Mobility Unified Reporting (MUR) application.

This chapter describes the following topics:

- [Introduction](#)
- [MUR Architecture](#)
- [Distributed Architecture of MUR](#)
- [MUR Features](#)
- [MUR Deployment](#)
- [MUR System Requirements](#)
- [MUR Ports](#)

Introduction

The Cisco Mobility Unified Reporting (MUR) system is a Web-based application providing a unified reporting interface for diverse data from Cisco Systems In-line service and storage applications.

The MUR application enables:

- Generating customized reports and comparison charts.

This release of MUR supports generating HTML-based historical canned reports displaying data in graphical—graphs/charts—and tabular formats. Reports for ad-hoc periods are not supported. For information on the various reports supported, see the [Report Types](#) section.

- Analyzing the reporting data and enabling the operator to get a full understanding of the performance of the network, enabling operators to optimally configure and plan their network.
- Supporting distributed installation which allows to view reports from multiple sites.
- Rich visualization (Graphs/tabular form).
- Exporting reports in Microsoft Excel, Adobe PDF, and CSV formats.
- Report scheduling, notification, and distribution. The report notification can be in the form of alarms/traps.

In this release, MUR supports sending e-mails to registered users' IDs for all the alarms including the KPI alarms.

- Capacity monitoring and planning of system supporting a suite of products such as PDSN, GGSN, SGSN, and inline service applications like Content Filtering, Stateful Firewall, Application Detection and Control (ADC).

The MUR application is available for report generation only when you install the software application on to your local server. For information on the server recommendations, refer to [MUR System Requirements](#) section in this guide. For information on how to install the MUR application, refer to the *Managing Mobility Unified Reporting System Installation* chapter in this guide.

The MUR application provides comprehensive and consistent set of statistics and customized reports, report scheduling and distribution from ASR chassis / in-line service product. For example, a subscriber's Quality of Experience, top 10 sites visited, top 10 users, and so on.

The MUR application provides reporting capability for Content Filtering (CF) data, bulk statistics, Key Performance Indicators (KPIs), EDR data from in-line service and storage applications. The MUR application facilitates and enhances the operators' ability to simply and easily determine the health and usage of the network.

The chassis directly pushes the bulkstat files and EDR data to the reporting server through SFTP. MUR receives the input data from the chassis only when the Enhanced Charging Services (ECS) module is enabled and configured to generate reporting EDRs. To enable this, you must purchase and install ECSv2 license on the chassis.

 **Important:** In RHEL-based deployment of MUR, L-ESS is NOT required as the ASR chassis Enhanced Charging Services (ECS) module can be configured to push the xDRs directly to the MUR reporting server. Push from ASR chassis is the Cisco recommended deployment model. Currently L-ESS is supported only on Solaris platforms. For information on the L-ESS installation instructions, refer to the *ESS Installation and Administration Guide*. Existing deployments where L-ESS is installed, to pull EDRs from the chassis, may continue with their deployment model in the 12.0 version of MUR Software Release and later.

For information on obtaining and installing the license, see *System Administration Guide* and *Enhanced Charging Services Administration Guide*. For information on configuring the ECS module, see *Configuring Chassis for Mobility Unified Reporting System* chapter in this guide.

MUR receives the following types of EDRs for report processing:

- CF-EDRs
- Flow EDRs
- HTTP EDRs

To reduce disk space and improve performance, MUR limits the bucket distribution for EDR data to ONLY last 2 days in case a EDR is spanning across more than 2 days or so.

For example, if the following EDR is received:

```
#sn-start-time,sn-end-time,radius-calling-station-id,ip-subscriber-ip-address,sn-
subscriber-port,ip-server-ip-address,sn-server-port,sn-app-protocol,p2p-protocol,traffic-
type,voip-duration,sn-volume-amt-ip-bytes-uplink,sn-volume-amt-ip-bytes-downlink,sn-
volume-amt-ip-pkts-uplink,sn-volume-amt-ip-pkts-downlink,bearer-3gpp rat-type,radius-
called-station-id,bearer-3gpp imei,ip-protocol,bearer-3gpp sgsn-address,sn-flow-start-
time,sn-flow-end-time
1275330600,1275334200,9689944191,19.19.1.1,35111,1.1.1.1,21,8,,0,52428800,1048576,100,20
0,1,apn.org1,35302703-090362-52,6,1.1.1.3,1275330600,1275334200
```

MUR determines the difference between the starttime and endtime attributes and limits the bucket distribution as shown here.

```
#starttime,endtime,protocol,rxbytes,txbytes 2011/02/26 10:00:00,2011/02/28
10:00:00,HTTP,100MB,100MB
```

 **Important:** The bucket distribution calculation will remain intact i.e. the volume will be distributed equally among all the half-hour's buckets that fall in the starttime and endtime.

 **Important:** The MUR receives the data in terms of EDRs which are generated based on the flow. As the EDRs are flow-based and the bulkstats is a real-time data, the volumes reported in the EDR are different from the volumes reported by bulkstats.

For more information on using the MUR application to generate reports, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Report Types

The MUR application supports generation of canned statistical reports that can be used to analyze network performance, and decide the policies for users, and identify the customer trends, network usage patterns, network categorization, etc. The reports can be per gateway, or multiple gateways (region), or for the overall network. The reports can be generated for the usage of different entities such as gateway, content type, etc on an hourly, daily, weekly, or monthly basis.

The typical canned reports that are supported for the MUR application include:

- Historical summary reports (Daily/Weekly/Monthly)
 - Half-hourly Reports: Usage reporting for the specified time period

- Daily Reports: Usage reporting for the past 24-hour period (midnight through midnight)
- Weekly Reports: Usage reporting for the past seven day period (Monday through Sunday)
- Monthly Reports: Usage reporting for the past 30-day period (1 day of the month through the last day of the month)
- Top “N” Reports
- Statistical and analytical reports
- Bulkstats and KPI reports

The static report layout comprises the following sections:

- The report name
- The report ownership: the user account that requested the report
- The date and time of generation
- The list of report parameters
- The chart legend (displayed under the chart)

On the interactive layout the user can set a series of preferences in a specific manner. The user has the flexibility to change the type of chart from Bar to Pie (supported output types depend on the selected report). Changing the preferences like the chart type or report parameters will cause the report to refresh in the same window.

The interactive chart layout provides the following list of features:

- Tool tip: When the mouse pointer stops over a chart series, after a short time a tool tip is displayed showing the information of the targeted sample.
- Dynamic legend: The legend is located beneath the chart and is used to recognize the series plotted on the screen. In case of series representing either network services or subscriber packages, the colors are bound to the service/package names. This means that, for example, the HTTP Service will be rendered with a specific color for the reports. The legend is usually displayed with check-boxes associated to each color.

The MUR application provides the following reports:

- Traffic Analysis Report: The Traffic Analysis report provides the total usage traffic (including uplink and downlink traffic) details for the following application categories:
 - Video
 - Filesharing
 - Web
 - IM
 - VOIP
 - Standard
 - Streaming
 - Tunnel
 - Gaming
 - Unclassified

MUR supports traffic type detection for P2P protocols such as Skype, Gtalk, MSN, Yahoo, and Oscar with the use of “traffic-type” attribute present in the EDR fields. Based on the value of this EDR attribute, the data will be classified to respective protocols.

The usage traffic is expressed in terms of megabytes (MB) or Megabits per second (Mbps) and percentage (%). The traffic can also be in gigabytes (GB) / kilobytes (KB) / bytes depending on the magnitude.

- **Traffic Distribution Report:** The Traffic Distribution report provides the summary of total traffic distribution for all the protocols application categories over a specified time period. The usage traffic is represented in MB/Mbps and percentage.
- **Active Flow Count Report:** The Active Flow Count report provides the details of traffic distribution flow count against the different application categories. This report also provides the summary of maximum number of flows in the EDR records.



Important: Active Flow Count report for current date will not be available because daily tables used to fetch this report are generated only at the end of the day. Also when the user selects a date range, for example, 10/1/2011 to 10/5/2011 where 10/5/2011 is the current date, then the report will be shown for the period 10/1/2011 to 10/4/2011 i.e. up till 10/4/2011.

Release 12.2 onwards, the Active Flow Count report will show flow counts for a sample/bucket (as per the configured granularity) that has maximum number of flows for selected filters in flow count summary. This new behavior is applicable to data ONLY after upgrading MUR to 12.2 version. Previous data will be shown as per the old reporting behavior.

- **Unique Subscriber Hits Report:** The Unique Subscriber Hits report provides an overview of the usage patterns of the entire subscriber population per protocol, for example, how many people are actually using VoIP.



Important: Unique Subscriber Hits report can be generated ONLY for a single date/week/month and not for any date-range. Also, note that the time selection is also disabled for this report.

Typically, this report provides the total number of times a subscriber is using a specific protocol. These reports are displayed for all configured gateways.



Important: Unique Subscriber Hits report for current date will always be available on the subsequent date because unique subscribers hits calculation will be performed at the end of the day.

- **TopN versus Total Traffic Report:** This report provides the summary of total usage traffic and Top N subscriber traffic for all the protocols over a specified time period. The usage traffic is represented in MB/Mbps and packets.
- **TopN Subscribers Report:** The TopN Subscribers report simply counts the number of bytes per subscriber for different time intervals. It displays the top 10/100/1000 subscribers for each day/week/month. This report is displayed across all configured gateways, per region or per NOC.



Important: This report is not available for a multiple date range selection.

After identifying the total amount of transferred data per subscriber, and identifying the top users, to understand the protocol and services breakdown for each subscriber, this report allows listing the different applications used by the top 10/100/1000 subscribers based on the selection of top subscriber per day/week/month.

- **TopN VCD Subscribers Report:** The TopN Voice Call Duration (VCD) Subscribers report displays the top N subscribers based on their voice usage (voice duration) for Yahoo, MSN and Skype voice protocols. The summary report displays the voice summary (voice duration) for VoIP category.



Important: This report is also available per week or month.

- **Weekly Report:** The weekly report provides details of the following:
 - Total traffic
 - Total traffic by category
 - VOIP Call Duration
 - Total unclassified traffic (TCP and UDP)
 - Top N subscribers
- **Monthly Report:** The monthly report provides the details of total traffic across the top N protocols / application categories in a month.
- **Custom Reporting:** MUR supports on-demand offline reporting of subscriber specific information to operators. This ad-hoc request could be a subscriber search request or top N search request.

Offline Subscriber Report: The MUR aids in searching individual subscribers' data based on certain parameters like IMSI, MSISDN, NAI, IMEI and Public and Private (NAT) subscriber IP address with ports, individually or in combination, and generates a subscriber-specific report showing the list of URLs visited by the subscriber, and other details like QoS, usage traffic, aggregate application/protocol breakdown, etc for the specified time period. MUR mainly supports this search functionality to track a subscriber or a set of subscribers for lawful intercept.

To use this Offline Reporting feature seamlessly, you must configure the EDR Filename Format appropriately through the Gateway configuration from **ADMIN** tab, and organize the archive directory date-wise. For information on how to manage the archive directory, see the *Managing Archive Directory* section in the *MUR Administration and Management* chapter of this guide.

Offline Top N Subscribers Report: MUR also facilitates to generate a report that covers the % of volume/duration used by N% subscribers. This report provides information on the absolute number of subscribers and the list of MSISDNs to facilitate correlation with the provisioning data. In this release, this ad-hoc report is available per APN group, Device Group, Location Group, and Service Profile.

Through this custom TopN reporting feature, it is possible to monitor and report the video traffic usage as and when needed. This report is mainly required to identify TopN hosts for video traffic and also to determine the biggest sources of video traffic, which drives the network load at a greater extent.

HTTP content type will be used to identify the video traffic. Ideally video traffic should be derived from flow-EDRs. Since the video usage monitoring report is generated based on HTTP content type, only HTTP traffic will be counted.

For more information on these features, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Reports based on Tethering Configuration: Tethering refers to the use of a mobile smartphone as a USB dongle/modem to provide Internet connectivity to PC devices (laptops, PDAs, tablets, and so on) running on the smartphone's data plan. Typically, for smartphone users, most operators have in place an unlimited data plan, the usage of which is intended to be from the smartphone as a mobile device. However, some subscribers use the low cost / unlimited usage data plan to provide Internet connectivity to their laptops in places where normal Internet connection via broadband/WiFi may be costly, unavailable, or insecure.

The ASR chassis works in conjunction with the MUR application to facilitate tethering detection on the chassis. The EDRs generated by the chassis will be enhanced to include OS signatures.

MUR processes flow-EDR files containing OS signature and IMEI field, HTTP files containing OS signature, User Agent, and IMEI field, and populates the tethering data in database files.

For more information on this feature, see the *Cisco Mobility Unified Reporting System Online Help* documentation and *Enhanced Charging Services Administration Guide*.

- **DPI Report:** The Deep Packet Inspection (DPI) reports are the canned statistical reports at the gateway level and region level. You can configure the MUR application to generate the reports for any of the available gateways.

In this release, MUR supports generating daily, weekly and monthly summary details and busy hour traffic usage details for the following report categories:

- Traffic Analysis Report
- Traffic Distribution Report
- Active Flow Count Report
- Unique Subscribers Hits Report
- TopN Reports — Report on Top N vs Total Traffic, TopN subscribers, TopN VCD subscribers



Important: Release 12.2 onwards, users with only administrative privileges can decrypt the subscriber's MSISDN to make it appear in the clear text format in the weekly reports.

MUR has the capability to report the following details per protocol:

- Total volume for the day/week/month
- Volume distribution in the busy hour
- Peak performance for the day/week/month
- Maximum number of unique subscribers

MUR supports additional information breakdown by network characteristics. These include Application Category, Protocol Groups, IP Protocol, Device Group, RAT (Radio Access Type i.e 2G vs 3G), APN (Access Point Name), SGSN group, Service Profile, Roaming Partner, and Location Group. During its development, a device may have several TAC codes and there may be a need to report devices by broader device type such as "Blackberry" or "Smartphone". Device groups allow the operator to combine a range of TACs into a single named group for reporting purposes.

Busy Hour Reporting: Busy Hour (BH) reporting is mainly useful for the users to monitor different traffic flows in their network during the busy hour. BH indicates the sliding 60-minute period during which occurs the maximum total traffic load in a given 24-hour period.

Please note the following key points:

- BH reporting is available ONLY on the GUI and not in xls format.
- BH reporting is available only under the **DPI** tab.
- BH radio button is available on the date panel.
- BH reporting is available for a date, date range, week and month.
- Busy hour reports are currently available ONLY at the NOC level.

DSL Reports: The current release of MUR provides the following details for Digital Subscriber Line (DSL) traffic reports:

- Traffic analysis — uplink DSL, downlink DSL and total DSL traffic including daily weekly, and monthly aggregation/distribution.
- DSL traffic categorization — total P2P traffic over DSL, IP traffic, web traffic, etc.
- Top N% DSL subscribers
- Comparison of total DSL traffic versus total UMTS traffic

For information on additional reports supported through DPI, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

- CF-RE Report: Content Filtering (CF) solution enables operators to filter HTTP and WAP requests from mobile subscribers based on the URLs in the requests, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences.

The CF-RE report provides the summary of traffic over CF categories, CF actions, and CF ratings. The CF actions that can be taken on the URL are as follows:

- allow
- discard
- redirect-url
- content-insert
- terminate-flow
- reply-code-terminate-flow

The CF ratings can be one of the following:

- dynamic
- static
- blacklisted

The CF-RE report also provides the list of top N subscribers and URLs based on their unique subscriber's hit count and total usage.

- HTTP Reports: The HTTP report provides summary and top N information on the traffic distribution, number of hits, number of subscribers using the HTTP service based on the HTTP content type groups and its sub-types. The HTTP summary reports are available at NOC level and region level whereas HTTP Top N reports are available ONLY at NOC level. Date range support is not available for Top N reports.

Typically, MUR supports the following categories of HTTP reports:

- Summary reports — Content type/subtype group distribution available for daily, weekly, monthly, and date range
- Top N reports
 - HTTP Group Aggregation — TopN HTTP Hosts by Volume; TopN HTTP Hosts by Hit count; TopN HTTP Hosts by Unique subscriber hits
 - Top N Referrer Group Aggregation by Hit count
 - User Agent (UA) / UA Group Aggregation — TopN UA reports also available for APN-TAC combination in addition with individual per APN, per TAC reports
 - HTTP Services Aggregation — TopN HTTP Services by Volume; TopN HTTP Services by Hit count

The top N referrers' report provides details of the total hit count for top N referrers and their sub-domain wise traffic distribution.



Important: In the distributed model of MUR, the data received from RDP is populated and TopN referrer report is available only at NOC level.



Important: It is mandatory to configure *http-url* and *http-referer* fields in the EDR records for top N HTTP referrers report generation.

- **Top N Unknown Ports:** This report highlights the top N ports for which traffic is classified as either unidentified or unknown. This report also lists the underlying IP protocol, downlink volume (in Megabytes), uplink volume (in Megabytes) and total volume (in Megabytes).

The report on top N unknown ports can be viewed through the link **Edr unknown port infos** under the **System** menu.

- **Bulkstat Report:** The Bulkstat report provides details of the processed bulk statistics from any application (PDSN, GGSN, SGSN, and so on) on the managed nodes in a timely manner.



Important: Make sure that you configure the bulkstats schemas through the GUI to generate bulkstats reports for any of the available gateways. For more information on schema configuration, refer to the *Configuring Bulkstats Schemas Using GUI* section in this guide and also *Cisco Mobility Unified Reporting System Online Help* documentation.

The bulkstat data is sent from the gateway to the MUR server with GMT (UTC Time stamps). The bulkstat file processing is triggered by the MUR scheduler engine. The scheduler processes the bulkstat files line by line for each gateway, and gets the schema, timestamp, and key index. If the index does not exist, the parser creates index and inserts data into bulkstats data table. Once the processing is complete, this data file is moved to the archive directory. Summarization must happen as the user moves from gateway to higher levels.



Important: For Bulkstat, there is no support for distributed model and all the bulkstat input files will be parsed by master MUR system only.

MUR supports generation of busy hour reports, top N Min/Max reports, performance aggregation reports i.e. daily, weekly and monthly summary reports.

Please keep the following key points in mind for bulkstats reporting:

- The gateway(s) and MUR server need to be NTP synced for accurate BS aggregation reports.
- Hourly aggregation reports are triggered at 50th minute of every hour.
- Daily reports are scheduled at 3:45 PM the next day.
- Weekly reports are scheduled at 5:00 PM every Monday.
- Monthly reports are scheduled at 06:15 PM on 1st of every month.
- **KPI Report:** The KPI report provides details of the KPIs for each selected schema. KPIs are the formula-based calculations of selected bulk statistics counters. You can configure the MUR application to generate the reports for any of the available gateways. For a complete listing of supported KPIs and its associated formulas/descriptions, see the *Cisco Mobility Unified Reporting System Online Help* documentation.
- **KPI Canned Report:** Canned report can be enabled/disabled for any of the available KPIs. This can be configured through the MUR GUI under **System > Kpis** menu. This will display hourly reports at NOC level. These KPI values will be pre-calculated and stored in the DB at the end of each day.

Once a KPI is enabled/disabled, it will start generating the canned report from the immediate next day.

 **Important:** Please note that the Bulkstats and KPI reports are displayed based on the gateway's time zone.

 **Important:** Please note that the subscriber's private data like Mobile Station Integrated Services Digital Network (MSISDN) will appear encrypted in all the subscribers reporting. Users with administrative privilege can only decrypt the MSISDNs using a shell script utility. For information on how to use this script, refer to the *MUR Administration and Management* chapter in this guide. The MSISDN decryption can also be accomplished through **Admin > Users** menu in the GUI. For decryption through the GUI, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

 **Important:** Please note that the availability of any report is typically based on the date/date range configurations and purging interval. If you are trying to view a report beyond the configured purging interval, MUR system will display an error message indicating that the report is unavailable.

For more information on each of these reports, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Exporting Reports to Other File Formats

The MUR application supports exporting reports to the following file formats:

- Microsoft Excel format: To export a report to Microsoft Excel format, use the `get_excel_report` script in the CLI. For more information about this script, refer to the *Generating Reports in Excel Format* section in the *MUR Administration and Management* chapter of this guide.
Exporting of reports to Excel format is also possible through the GUI by clicking the excel icon present in the tabular view of each of the reports under **HOME** and **DPI** tabs.
- Comma Separated Value (CSV) file format: To view reports in CSV format, in the HOME and DPI tabs, click the csv icon present in the tabular view of each of the reports.
- PDF format: To export a report to PDF format, in the **HOME** and **DPI** tabs of the MUR GUI, click the **PDF** button. The PDF file is displayed in a new window and can be saved for future reference.

If there is no data available for a report, the **PDF** button is disabled.

- Text File format: This format is applicable only to HTTP User Agent (UA) reports. To export this report in a text file, click **Export to Text** button available in the HTTP UA reporting page.

For more information, see the *Cisco Mobility Unified Reporting System Online Help* documentation.


License Requirements

The MUR system is a licensed Cisco product. Contact your Cisco account representative for detailed information on specific licensing requirements.

MUR Architecture

The MUR solution consists of two components — a server and a GUI client. The following figure shows a typical organization of the MUR solution.

Figure 1. Internal Architecture of MUR



The server components include:

- DB Server: This is the standard PostGreSQL 8.3 database server. This is started at the time of application startup.

MUR uses pgbouncer utility for postgres connection pooling. This utility gets started/stopped with Postgres Server.

- Quartz Scheduling Engine: This is the core of the MUR reporting solution. It is used to schedule different tasks such as parsing of incoming data files (bulkstat, EDR, etc.), trigger various canned reports on a periodic basis, cleaning up of stored outdated data and files, and so on.
- Generators: These are python based scripts that are used for parsing various CSV files. The files are parsed to an extent where generated files (or data in database) themselves represent meaningful data. This is a very powerful concept introduced for faster processing of information.

The generators archive the files once they are parsed. In archival, the files are zipped and placed in the configured location.

- KPI Parser: The KPI Alarm Generator uses the information stored by bulkstat parser in the database for KPI calculations and then, based on the calculations, generates the alarms that are subsequently sent to Network Node Manager (NNM).
- Notif Server: This stands as a separate entity that collects information from the MUR system and generates alarms which are then sent to the NNM for further analysis.
- Loggers: The MUR application uses various loggers so that application logs with various severities are made available for debugging purpose.
- MUR Parser Server: This will be running as daemons, and it will be spawned at the time of **serv start** command. Parser server will keep running in background and will perform the parsing activity for all gateways.

The following is a sample output of the **serv status** command:

```

-----
----- MUR Process Status -----
PID                Process                Status
-----
4245                Process Monitor          Running
4256                Scheduling server        Running
4267                Postgres Server          Running
4289                Apache Server            Running
3249                Notif Server             Running
3243                Parser Server            Running
2430                Cache Server             Running
-----

```

The following describes the sequential steps associated with the functioning of RPC parser daemons.

1. For each configured gateway, RPC Parser daemon will check if the appropriate reporting (Flow/HTTP/CF) is enabled or not.

If say, Flow-EDR reporting is enabled for GW1, RPC Parser daemon will check the Process Count configured for Flow-EDR under **System** menu.

2. Depending on the number of processes configured, RPC Parser daemon will spawn those many RPC server instances for GW1. Also, it will update each RPC server URL in DB as shown below:

3. RPC Server Instances for Gateways

ID	Gateway ID	Reporting Type	RPC Server URL	Process ID
1	1	Flow-EDR	http://localhost:8000	7643
2	1	Flow-EDR	http://localhost:8001	8756
3	1	Flow-EDR	http://localhost:8002	9054
4	1	Http-EDR	http://localhost:8003	5645
5	1	Http-EDR	http://localhost:8004	6576
6	1	Http-EDR	http://localhost:8005	8678

4. Steps 1 through 3 are repeated for each configured gateway and reporting type.
5. Normalization daemon will pick up the set of files to be parsed. Depending on the number of files to be parsed, it will get the corresponding RPC server information from DB from the above table.
6. Depending on the number of files to be parsed, normalization daemon will spawn those many threads. Each thread will allocate its bunch of files to corresponding RPC server instance. The RPC server instance will parse and store the normalized data in DB and the corresponding thread will exit.
7. If the Process count is increased/reduced, additional RPC server instances will be fired/closed as and when required.
8. Both the normalization daemon and RPC Parser daemon will be continuously running in background.
9. Normalization daemon will be spawned by the scheduler initially. RPC Parser daemon will be spawned through `serv start` command.

Some of the components at the client side include Django and Mod_python.

Distributed Architecture of MUR

MUR supports the distributed model to allow the deployment which enables network wide view or work load balancing. Newly introduced component, Remote Data Processor (RDP), plays the role of pre-processing the input files from gateways. One or more RDPs, installed separately on remote machines can be registered to a master MUR and one RDP can process files from one or more gateways.

RDP periodically sends the intermediate data to registered master MUR. The role of MUR in such deployments is mostly for report generation, report viewing, RDP management and optionally data processing.

 **Important:** RDP installation and registration is required only for network wide deployments. For standalone installation no RDP is required. For information on how to install the RDP, refer to the *Managing MUR Installation* chapter of this guide.

 **Important:** Make sure that you first install the master MUR system and then proceed with the RDP installation. Also, note that the RDP and MUR must be installed, upgraded, and uninstalled separately.


 **Important:** Before registering RDP with the master MUR, ensure that the RDP is installed and running.

 **Important:** The RDP management like configuration and removal is possible from MUR GUI only. For information on managing the RDPs, refer to the *Cisco Mobility Unified Reporting System Online Help* documentation.

 **Important:** For Bulkstat, there is no support for distributed model and all the bulkstat input files will be parsed by master MUR only.

The following figure illustrates the distributed architecture of MUR.

Figure 2. Distributed Architecture of MUR



How RDP works with MUR

This section describes how the RDP works with the MUR application.


The RDP parses the raw data or EDR files from one or more GGSNs and populates the database for required reports. The RDP pre-processes the data and then periodically forwards them to the master MUR through SFTP for report generation.



Important: If the distributed model of MUR is used, then the SFTP user name and password should be the same as the MUR Administrator user's login name and password provided during installation. For information on configuring SFTP details, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Each of the RDP and MUR will be assigned a unique ID during installation and will be used for identification of each RDP along with its gateway and data.

Figure 3. MUR with RDPs in Distributed Model



In 12.0 and earlier releases, each of the registered RDPs will form a new region. RDP region can be a child of the root of the MUR (NOC) or can be the child of another region. The gateways associated with a RDP will always be the children of RDP region.

Release 12.2 onwards, users can create individual regions and add RDPs to the regions. All the gateways must be associated with RDP(s) or NOC and not to a region directly.



Important: Only single MUR can communicate with an RDP simultaneously.

MUR Features

This section provides information on the basic features of MUR application and its implementation.

Clustering Support for High Availability

MUR application consists of important internal entities such as Apache, Postgres, scheduling, cache, parser, notif servers and process monitor which run on a machine and communicate with the external entities such as ASR 5000 chassis and in turn provide Web Reporting capability to operators.

Whenever the machine or MUR process gets crashed/stopped, there are chances of loss of communication between internal and external entities. To avoid downtime and ensure continuous availability of MUR application, High Availability (HA) support using Veritas Clustering has been provided.

Using Veritas Cluster, there will be two machines configured to run MUR application. In case of machine failure, MUR Server will failover (move on to) from the current machine (Active Node) to another machine (called Passive node in cluster terminology) and that node will handle further processing of data. Operator will not be aware of this failover except that he will be asked to re-login on MUR UI after failover. A floating IP (shared IP), which is common to both the nodes, will be used for accessing MUR GUI. High Availability is supported for both standalone as well as distributed architecture of MUR. So in case of distributed architecture, there will be extra nodes for Master and each RDP to support failover mechanism.



Important: High Availability mode installation is only supported on RHEL platform.

Operation

Veritas cluster continuously monitors the MUR Process Monitor to ensure application availability. The shared IP address is floated across the Cluster nodes, which will be used by chassis for sending the EDR or Bulkstat data files and also to invoke the MUR GUI. MUR application is data centric and will have to be installed on the shared storage. Shared storage will be mounted on active node that will host MUR application. So the cluster service has to make the IP address, shared storage and MUR processes highly available.

The resources monitored by Veritas cluster are:

- NIC - Monitors a NIC (Network Interface Card)
- IP - Monitors the shared IP address
- Disk Group, Volume and Mount - for shared storage
- MUR Application - comprising of all the MUR related processes

For information on minimum system requirements and MUR configuration for HA deployment, please see the *Mobility Unified Reporting System Clustering Support for High Availability* chapter of this guide.

HTTPS Access

Hypertext Transfer Protocol Secure (HTTPS) is a combination of the Hypertext Transfer Protocol (HTTP) with Secure Sockets Layer (SSL) / Transport Layer Security (TLS) protocol to provide encrypted communication and secure identification of a network Web server.

HTTPS is a URI scheme that is, aside from the scheme token, syntactically identical to the HTTP scheme used for normal HTTP connections, but which signals the browser to use an added encryption layer of SSL/TLS to protect the traffic. SSL is especially suited for HTTP since it can provide some protection even if only one side of the communication is authenticated. This is the case with HTTP transactions over the Internet, where typically only the server is authenticated (by the client examining the server's certificate).

In Release 14.0 and later, MUR supports HTTPS communication between MUR UI (client browser) and MUR server which will allow addressing the following important security considerations:

- **Authentication:** During the initial attempt to communicate with a Web server over a secure connection, that server will present the Web browser with a set of credentials in the form of a server certificate. The purpose of the certificate is to verify that the site is who and what it claims to be. In some cases, the server may request a certificate that the client is who and what it claims to be (which is known as client authentication).
- **Confidentiality:** When data is being passed between the MUR client and the MUR server on a network, third parties can view and intercept this data. SSL responses are encrypted so that the data cannot be deciphered by the third party and the data remains confidential.
- **Integrity:** When data is being passed between the MUR client and the MUR server on a network, third parties can view and intercept this data. SSL helps guarantee that the data will not be modified in transit by that third party.

Apache's HTTPS capability will be leveraged for this. The SSL server certificate will be self generated. The certificate usually contains the server name, the trusted certificate authority (CA) and the server's public encryption key.

 **Important:** To effectively use HTTPS support, users are encouraged to create their own custom certification. If they do so, there will not be any address mismatch for MUR GUI and hence no warning. However, if the user wants to use the default Cisco certificate/key, then the user should perform a one-time process of adding an exception like disabling mismatch warning in the Settings in all the Web browsers such as Mozilla, IE, and Chrome.

 **Important:** During MUR upgrades, please make sure to upgrade Apache Server and Certificates if needed. For example, when upgrading from an MUR version with no HTTPS support to a new MUR version with HTTPS implementation support, you should copy the whole Apache server with HTTPS support and also the default self signed certificate/key to <MUR_install_dir>/starbi/certificate directory. If you are using your own custom certificates, please remember not to update the certificate directory.

Creation of Security Certificates

The user should create self signed certificate and key to authenticate the client. Certificate and key can be generated by OpenSSL (www.openssl.org). Use the following command to generate self signed certificate and key.

```
openssl req -new -x509 -nodes -out server.crt -keyout server.key
```

This command will ask some user inputs. For example:

```
Country Name (2 letter code) [US]:IN
```

```
State or Province Name (full name) [Some-State]:Maharashtra
```

```
Locality Name (eg, city) []:Pune
```

```
Organization Name (eg, company) [Unconfigured OpenSSL Installation]:Cisco Systems
```

Organizational Unit Name (eg, section) []:MITG

Common Name (eg, YOUR name) []: cisco.com

Email Address []: (you can leave it blank)



Important: Once both the certificate and key are created, they will be populated and stored in the `<MUR_install_dir>/starbi/certificate` directory.

Enabling Certificates on Browser

This section describes the procedure to enable SSL certificate on different browsers.

To enable certificate on IE:

1. Copy certificate file on Windows machine and open it with IE.
2. Click **Tools -> Internet Options -> Content -> Certificates -> Trusted Root Certification Authorities** and click **Import**.
3. Follow the wizard and import the *server.crt* into Trusted Root CA.
4. Open new IE and access the page.

To enable certificate on Mozilla:

1. Click **Tools -> Options -> Advanced -> Encryption tab -> View Certificates -> Authorities**. Click **Import** and browse to select *server.crt* file and click **OK**. In the window that appears, select "This certificate can identify websites".
2. Check and ensure that the selected certificate appears under the tab **Authorities**. Open new Mozilla and access the page.

To enable certificate on Chrome:

1. Click **Settings -> Under the Bonnet -> HTTPS/SSL -> Manage Certificates**.
2. Import the certificate in the Trusted Root Certification Authorities.

Implementation on RHEL

SSL is already enabled in Apache Server for RHEL. That means the required libraries are already present in the Apache Server. You should only configure HTTPS in the configurations files *httpd.conf* and *httpd-ssl.conf* in the respective directories `<APACHE>/conf` and `<APACHE>/conf/extra`.

1. Edit the *httpd.conf* file in the `<APACHE>/conf` directory.
Uncomment "Include conf/extra/httpd-ssl.conf"
2. Edit the *httpd-ssl.conf* file in the `<APACHE>/conf/extra` directory.
 - Listen to port 9443 (Change your HTTPS port)
 - SSLSessionCache
"shmcb:<MUR_install_dir>/starbi/apache2/logs/ssl_scache(512000)"
 - <VirtualHost *:9443>
 - ServerName cisco.com
 - ServerSignature On

- SSLEngine on
- SSLProtocol all -SSLv2
- SSLCipherSuite ALL:!ADH:!EXPORT:!SSLv2:RC4+RSA:+HIGH:+MEDIUM:+LOW
- SSLCertificateFile “<MUR_install_dir>/starbi/certificate/server.crt” (Path of certificate file created above)
- SSLCertificateKeyFile “<MUR_install_dir>/starbi/certificate/server.key” (Path of key file created above)



Important: All the changes mentioned in steps 1 and 2 will be done by installer scripts.

Implementation on Solaris

The Solaris version of Apache does not have SSL enabled. That means the required libraries are not present in Apache Server.

To make Apache Server enabled, you should install openssl from IPCentral. And then recompile the Apache Server. The following are the steps to be followed to perform this action:

1. Compile and install OpenSSL.

- Get Open SSL from IPCentral. Download openssl-0.9.8i-9-2.tar.
- Run the command `tar xvf openssl-0.9.8i-9-2.tar`
- Run the command `PREFIX=openssl-0.9.8i-9-2`
- Run the command `export LD_OPTIONS="-R/<INSTALL_DIR>/${PREFIX}/lib"`
- Run the command `cd openssl-0.9.8i`
- Run the command `./Configure solaris-x86-gcc --prefix=<INSTALL_DIR>/${PREFIX} shared -R/<INSTALL_DIR>/${PREFIX}/lib`
- Run the command `make`
- Run the command `make install`
- Create a link using the following commands:

```
cd <INSTALL_DIR>
ln -s openssl-0.9.8i ssl
```

2. Compile and install the Apache Server.

- Get Apache2.2.14 from IPCentral. (httpd-2.2.14-5-2.tar)
- Run the command `tar xvf httpd-2.2.14-5-2.tar`
- Run the command `cd httpd-2.2.21`
- Run the command `./configure --prefix=<INSTALL_DIR>/apache2 --enable-ssl`
- Run the command `make`
- Run the command `make install`



Important: The above steps (1 and 2) will be executed manually on Solaris system to get the Apache Server with SSL libraries.

Now, you should configure HTTPS in the configurations files *httpd.conf* and *httpd-ssl.conf* in the respective directories `<APACHE2>/conf` and `<APACHE2>/conf/extra`.

3. Edit the *httpd.conf* file in the `<APACHE2>/conf` directory.

Uncomment “Include conf/extra/httpd-ssl.conf”

4. Edit the *httpd-ssl.conf* file in the `<APACHE2>/conf/extra` directory.

- Listen to port 9443 (Change your HTTPS port)
- SSLSessionCache
“shmcb:<MUR_install_dir>/starbi/apache2/logs/ssl_scache(512000)”
- <VirtualHost *:9443>
- ServerName cisco.com
- ServerSignature On
- SSLEngine on
- SSLProtocol all -SSLv2
- SSLCipherSuite ALL:!ADH:!EXPORT:!SSLv2:RC4+RSA:+HIGH:+MEDIUM:+LOW
- SSLCertificateFile “<MUR_install_dir>/starbi/certificate/server.crt”
- SSLCertificateKeyFile “<MUR_install_dir>/starbi/certificate/server.key” (Path of key file created above)



Important: All the changes mentioned in steps 3 and 4 will be done by installer scripts.

LDAP Authentication in MUR

The Lightweight Directory Access Protocol, known as LDAP, is based on the X.500 standard, but significantly simpler and more readily adapted to meet custom needs. Unlike X.500, LDAP supports TCP/IP, which is necessary for Internet access.

LDAP is used as a central repository for user information and as an authentication service. It can also be used to store the attribute based data and the role information for application users. The LDAP maintains data in a hierarchical structure wherein the entries are in a tree-like structure called Directory Information Tree (DIT).

Prior to Release 14.0, MUR authenticates users against the MUR local DB information. In Release 14.0 and later, users can be authenticated against the LDAP directory.

For this, user should configure the following parameters for communication with LDAP server:

- LDAP server as authentication backend
- LDAP server hostname, LDAP server port
- The Base Distinguished Name (DN) to start the search for users
- User's Relative Distinguished Name (RDN) to be used for search
- Type of mapping to be used to assign a role to the LDAP user
- Miscellaneous configuration based on the selected type



Important: For LDAP user, the **Change password** option will not be available in the MUR GUI.

The following procedure describes how to set up LDAP authentication in MUR.


1. User can select LDAP server as authentication backend. If the user information does not exist in LDAP directory then it will be authenticated against local MUR DB.
2. If the user selects LDAP as authentication backend then the above mentioned parameters need to be specified.
3. If LDAP backend is selected then when the user tries to login the authentication request is sent to the LDAP server by forming appropriate LDAP URL from LDAP hostname and port configured by the user.
4. MUR tries to bind LDAP server with the provided user credentials. A proper DN is formed from User DN and RDN and is used while binding to the LDAP server.
5. If binding is successful (user authentication is successful) a success message is sent to MUR.
6. Perform the authorization in either of the two ways:
 - Attribute Based: The value of the attribute which is used to map the MUR role/privilege to the LDAP user is compared with already configured values for MUR administrator and operator roles.
For example: In LDAP directory there is a user attribute title, if the value of title attribute equals 'admin' then this user will be mapped to administrator role in MUR.
For customer information: Already existing attribute can be configured or a new attribute for MUR specific usage can be added in LDAP directory.
 - Group Based: All users from an existing group from LDAP directory will be classified as MUR administrator or as MUR operator.
For example: There is 'sysadmin' group in LDAP directory. All users from this 'sysadmin' group can login to MUR as administrator.
If the authorization fails then this LDAP user will not be able to login to the MUR.
7. If the authentication fails at LDAP server then MUR will authenticate the user against its local DB.
8. If LDAP user logs into MUR then this user will be restricted for user administration activities in MUR (like user creation, modification and deletion activities).

Region-based Reporting

In MUR versions prior to 12.2, RDP is considered as region, hence all reports were based on RDP. Whenever an RDP is configured, internally MUR used to create corresponding region for the same. However, with the introduction and need of scalable MUR, one gateway's files would be processed by two or multiple RDPs. In that case, RDP does not stand as a region. So, reports would be required across all the RDPs under one particular region. Hence, reports would be available per region.

The following figure shows an example of scalable and region based reporting model. This example considers two regions, New York and Chicago, both of them having one gateway each. Each location has two RDPs.

Figure 4. Scalable Setup of MUR



The following procedure outlines the steps to be followed for configuring the scalable network model.

1. Install MUR in RDP mode on four nodes NYRDP1, NYRDP2, CHRDP1, CHRDP2, and in Master mode on fifth node.
2. Connect shared storage 1 to NYRDP1 and NYRDP2. Similarly, connect shared storage 2 to CHRDP1 and CHRDP2.
3. After configuring mount point and permissions, configure GWNY1 to push files to NYRDP1 on the shared storage 1 and GWCH1 to push files to CHRDP1 on shared storage 2.

Important: Ensure to check if EDR push functionality is configured on the gateway to send files to the RDP and shared path.

4. Configure two regions, New York and Chicago. To perform this, invoke **System** menu from MUR GUI and then click **Regions**.

Important: While configuring regions, select Parent as NOC.

5. Configure the RDPs. For example, to configure NYRDP1, on the GUI select **Admin** tab, click **RDP** from the left navigation pane and then click **Add RDP**.

6. Once all the RDPs are configured, add GWNY1 to NYRDP1 through the **Admin** tab. During configuration, select New York region and NYRDP1 for this gateway and make sure that path for incoming files is specified from the shared disk.
7. Similarly, configure GWCH1 on Chicago region on CHRDP1. During the configuration, the pattern for file name is given such that it picks up only required set of files. This can be achieved using regular expressions in the file name pattern. That is, configure the Flow EDR Filename Pattern as `edr_flow_*[1,3,5,7,9]` or `edr_flow_*[2,4,6,8]` so that the gateway picks up all odd or even numbered files accordingly.
8. Configure GWNY1's pseudo gateway on NYRDP2. During the configuration, enter NYRDP2 as RDP, region as New York, and mark this as pseudo with the check box. Similarly, configure a pseudo gateway for GWCH1.
9. The configuration setup is now complete and will start to function as needed.

When user clicks on region, say, Chicago, reports will be shown for GWCH1 (aggregated from both the RDPs under it). When user clicks on NOC, reports will be seen in consolidate manner for all regions.

On gateway, apart from mandatory configuration for EDR generation, it is must to have the sequence numbers enabled for the file generation. This would help MUR split the files into different RDPs.

 **Important:** In the gateway tree in **DPI**, **HTTP**, **CF** and **Bulkstats** tab, the pseudo gateway is NOT shown. This is because, there are no specific reports to the gateway, it is just a pseudo to original gateway and all the data is coming from the original gateway only.

If there are multiple RDPs parsing the traffic from same gateway, multiple pseudo gateways should be created for each RDP. While entering the file name pattern, user can use regular expressions as needed.

Load Distribution Based on Number of Files

Following are sample file pattern configurations that can be used to distribute load among configured RDPs.

Load distribution for 10-20 Gbps throughput — 50-50(%) files distribution for both flow/HTTP among two RDPs.

Gateways	HTTP Pattern	Flow Pattern
GW1	<code>*http*[0-4].*</code> or <code>*http*[0,2,4,6,8].*</code>	<code>*flow*[0-4].*</code> or <code>*flow*[0,2,4,6,8].*</code>
GW2(pseudo)	<code>*http*[5-9].*</code> or <code>*http*[1,3,5,7,9].*</code>	<code>*flow*[5-9].*</code> or <code>*flow*[1,3,5,7,9].*</code>

Load distribution for 20-30 Gbps throughput — 30-30-35(%) files distribution for both flow/HTTP among three RDPs.

Gateways	HTTP Pattern	Flow Pattern
GW1	<code>*http*[0-6][0-4].*</code>	<code>*flow*[0-6][0-4].*</code>
GW2(pseudo)	<code>*http*[0-6][5-9].*</code>	<code>*flow*[0-6][5-9].*</code>
GW3(pseudo)	<code>*http*[7-9][0-9].*</code>	<code>*flow*[7-9][0-9].*</code>

Load distribution for 30-40 Gbps throughput — 25-25-25-25(%) files distribution for both flow/HTTP among four RDPs.

Gateways	HTTP Pattern	Flow Pattern
GW1	<code>*http*[0-4][0-4].*</code>	<code>*flow*[0-4][0-4].*</code>

Gateways	HTTP Pattern	Flow Pattern
GW2(pseudo)	*http*[0-4][5-9].*	*flow*[0-4][5-9].*
GW3(pseudo)	*http*[5-9][0-4].*	*flow*[5-9][0-4].*
GW4(pseudo)	*http*[5-9][5-9].*	*flow*[5-9][5-9].*

Tethering Detection Feature

Tethering refers to the use of a mobile smartphone as a USB dongle/modem to provide Internet connectivity to PC devices (laptops, PDAs, tablets, and so on) running on the smartphone's data plan. Typically, for smartphone users, most operators have in place an unlimited data plan, the usage of which is intended to be from the smartphone as a mobile device. However, some subscribers use the low cost / unlimited usage data plan to provide Internet connectivity to their laptops in places where normal Internet connection via broadband/WiFi may be costly, unavailable, or insecure.



Important: In this release, the Tethering Detection feature is supported only on the GGSN, HA, and P-GW.

The Tethering Detection feature enables detection of subscriber data traffic originating from PC devices tethered to mobile smartphones, and also provides effective reporting to enable service providers take business decisions on how to manage such usage and to bill subscribers accordingly.



Important: Use of Smartphone tethering detection feature requires that a valid license key be installed. Contact your Cisco account representative for information on how to obtain a license.

For detailed information on this feature, refer to *Enhanced Charging Services Administration Guide*.

MUR Support for Tethering Detection

The ASR chassis works in conjunction with the MUR application to facilitate tethering detection on the chassis.

Upon enabling tethering detection feature through the GUI, MUR collects samples of HTTP and TCP signatures from live traffic to create a database of OS and UA signatures for assorted devices accessing the network through the gateways. For this, offline TAC-device mappings are fed to MUR, and MUR generates the signature databases based on EDRs generated by the chassis for various TAC groups.

MUR processes flow-EDR files containing OS signature and IMEI field, HTTP files containing OS signature, User Agent, and IMEI field, and populates the following set of data in the respective database files.

- Laptop (USB Dongles device group) - User Agent data
- Laptop (USB Dongles device group) - OS Signature data
- Smartphone - TAC data

MUR is configured in such a way that the database files are pushed to the ASR chassis under the `/hd-raid/databases/` directory.

For information on how to configure tethering detection feature, refer to *Configuring Chassis for Mobility Unified Reporting System* chapter in this guide.

Tethering Detection Databases

The Tethering Detection feature uses the OS signature, UA signature, and TAC databases.

These database files must be populated and loaded on to the chassis by the administrator. The procedure to load the databases is the same for all the three types of databases.

Before the database(s) can be loaded for the first time, tethering detection must be enabled using the **tethering-database** CLI command in the Active Charging Service Configuration Mode.

For all three databases, only a full upgrade of a database file is supported. Incremental upgrade is not supported. If, for any particular database, the upgrade procedure fails, the system will revert back to the previous working version of that database.

OS Signature Database

The OS signature database file is named “os-db”. The file contains OS fingerprint signatures that have been identified as non-smartphone signatures.

The OS fingerprint signature string is a null-terminated ASCII string of maximum 32 bytes in the following format:

```
<tl>|<ttl>|<d>|<wlen>|<mss>|<wss>|STEN
```

Where:

- *tl*: Total IP Packet Length
- *ttl*: Initial TTL
- *d*: IP DF bit
- *wlen*: TCP Window Length
- *mss*: TCP Maximum Segment Size
- *wss*: TCP option Window Size Scale
- *S*: TCP option Selective ACK OK
- *T*: TCP option Timestamp
- *E*: TCP option EOL
- *N*: TCP option NOP (count)

The maximum number of entries permitted in the os-db file is 16384.

The maximum size of the os-db file can be 524KB + 50 bytes for header and trailer.

In the 12.2 release, the file is in plain text format and contains one TCP signature in ASCII format, one entry per line.

The following is the content of a sample os-db file:

```
VERSION 1.1

BEGIN OS-DB

48|128|1|5840|1460|1|1112

44|128|0|5840|1460|1|1011

END OS-DB
```

UA Signature Database

The UA signature database file is named “ua-db”. The file contains UA signatures that have been identified as non-smartphone signatures.

The UA signatures are stored in plain text format in the database file so that manual modification of the database is possible.

The maximum number of entries permitted in the ua-db file is 16384.

The maximum size of the ua-db file can be 67MB + 50 bytes for header and trailer.

The following is the content of a sample ua-db file:

```
VERSION 1.1

BEGIN UA-DB

Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; Trident/4.0; SLCC2; .NET CLR
2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729; Media Center PC 6.0; InfoPath.2)

END UA-DB
```

TAC Database

The TAC database file is named “tac-db”. The file contains smartphone TACs that are uploaded in MUR by the operator.

The maximum number of entries permitted in the tac-db file is 16384.

The maximum size of the tac-db file can be 147KB + 50 bytes for header and trailer.

The following is the content of a sample tac-db file:

```
VERSION 1.1

BEGIN TAC-DB

01194800

01194801

END TAC-DB
```

Loading and Upgrading Tethering Detection Databases

This section provides an overview of loading and upgrading the OS, UA, and TAC databases used in tethering detection.

The database files from MUR must be copied onto the chassis to the following directory path designated for storing the database files:


/hd-raid/databases/

Any further upgrades to the database files can be done by placing the file named *new-filename* in the designated directory path. ACS auto-detects the presence of files available for upgrade daily. When a new version of a file is found, the upgrade process is triggered. The upgrade can also be forced by running the upgrade command in the CLI. On a successful upgrade this file is renamed to *filename*.

MUR Deployment

The following figure illustrates how the MUR reporting server interacts with the gateways and generates the reports.

Figure 5. End-to-end Component Mapping



The chassis / gateway supports on board Hard Disk Drive (HDD) for extended storage of the xDR files such as EDR, UDR, CDR, and NBR. If the HDD is configured, then the gateway pushes the files to an external entity like External Storage Server (ESS) for short-term storage. In case of no HDD support on the gateway, the Local, short-term External Storage Server (L-ESS) has the capability of pulling the files from gateways via SFTP, and send it for report processing. For more information on L-ESS, refer to the *ESS Installation and Administration Guide*.

The MUR server collects the EDRs, and bulkstats from gateways or L-ESS server, and processes the incoming data files and presents reports on Web-based GUI. The MUR application can generate reports in Excel, CSV, and PDF formats, and present them to users on a request basis.

Important: L-ESS is NOT required as the ASR5K EDR module can be configured to push the xDRs directly to the MUR reporting server. Push from ASR chassis is the Cisco recommended deployment model. Currently, L-ESS is supported only on Solaris platforms. For information on the L-ESS installation instructions, refer to the *ESS Installation and Administration Guide*. Existing deployments where L-ESS is installed, to pull EDRs from chassis, may continue with their deployment model in the 12.0 version of MUR Software Release and later.

For information on how to configure the chassis to push the xDRs, refer to the *Configuring Chassis for Mobility Unified Reporting System* chapter in this guide.

MUR System Requirements

This section identifies the minimum system requirements that are required for the deployment of MUR at the operator's premises.



Important: The hardware required for MUR may vary depending on incoming EDR generation, subscriber count, and number of gateways.

Server Recommendations for Use in Solaris Environment

This section identifies the minimum system requirements recommended when installing the MUR application in Solaris environment.

NEBS Requirements:

The following are the server specifications for MUR when an additional external storage is required:

- Sun Microsystems Netra™ X4270 server
 - Quad-Core two socket Intel Xeon L5518 processor
 - 32GB RAM
 - 2 * 300GB 10K RPM SAS disks
 - SATA DVD drive
 - 8 internal port SAS HBA
 - Choice of AC or DC power supplies
- Sun StorageTek 2540 SAS Array, Rack-Ready Controller Tray
 - 12 * 300GB 15K RPM SAS drives
 - Two redundant AC power supplies
- Operating system:
 - Sun Solaris 10 with latest patches installed
 - Sun Solaris 10 with patch 126547-07 towards the SUN bash vulnerability fix

Non-NEBS Requirements:

The following are the server specifications with only the internal storage used:

- Sun Fire X4270 server
 - Intel Xeon processor 5500 series
 - 32GB RAM
 - 16 * 300GB 10K RPM SAS disks
 - SATA DVD drive
- Operating system:
 - Sun Solaris 10 with latest patches installed
 - Sun Solaris 10 with patch 126547-07 towards the SUN bash vulnerability fix

 **Important:** It is strongly recommended to update the Operating System with the latest security patches.

 **Important:** The number of disks recommended is purely based on the throughput of the network and data retention configuration. Please contact Cisco Advanced Service Team for data sizing.

ZFS Pooling Recommendations:

This section provides information on the recommendations for ZFS pooling.

- OS pool: This mirrored ZFS pool shall be created for Solaris OS installation.
- MUR pool: This standard ZFS pool shall be created for MUR i.e. MUR installation, incoming data files.
- Postgres pool: This standard ZFS pool shall be created for MUR postgres database.
- Archive pool: This standard ZFS pool shall be created for retaining archived and data backed up files.

 **Important:** ZFS pool shall NOT be created with RAID-Z since ZFS does not allow attaching an additional disk to an existing RAID-Z pool. Hence, this freezes the chances of data scaling.

Server Recommendations for Use in RHEL Environment

This section identifies the requirements of server recommended when installing the MUR application in RHEL environment.

- UCS C460 M2 server
 - 4 x Intel® Xeon® E7-4860 @ 2.26 GHz, 130W 10 Core CPU / 24 MB Cache
 - 128GB RAM
 - 12 * 600 GB SAS 6G, 10K RPM
 - RAID Controller
 - 4Gb Dual port FC Host Bus Adapter

 **Important:** The number of disks recommended is purely based on the throughput of the network and data retention configuration. Please contact Cisco Advanced Service Team for data sizing.

- Operating System
 - Cisco UCS running OS version 'Cisco MITG RHEL 5.5'

For information related to OS installation, refer to the *Cisco MITG RHEL OS v5.5 Application Note*.

 **Important:** The Cisco MITG RHEL v5.5 OS is a custom image that contains only those software packages required to support compatible Cisco MITG external software applications. Users must not install any other applications on servers running the Cisco MITG v5.5 OS. For detailed software compatibility information, refer to the *Cisco MITG RHEL v5.5 OS Application Note*.



Important: ZFS Pooling recommendations are applicable ONLY for Solaris hardware.

- XFS/EXT-3 File System Volumes & RAID Recommendations

XFS file system is recommended for MUR application (/apps), Postgres (/db) and archive (/archive) partitions.

Storage RAID recommendation for MUR Application

CISCO UCS machine supports MegaRAID controller. This allows configuring the UCS hard disks into hardware RAID arrays (disk groups). The MegaRAID controller provides the BIOS utility for configuring the RAID.

The RAID recommendations for MUR are as follows:

- Separate disk arrays for OS, MUR and postgres (data directory).
- RAID Level - Combination of 5 and 0 depending upon the fault tolerance.
- Stripe size should be 256KB
- RAID Controller parameters —
 - Read Policy - Select Adaptive read ahead
 - Write Policy - Select Write Back
 - I/O Policy - Select Direct I/O

For information on configuring the RAID arrays using MegaRAID BIOS, refer to the *Configuring Cisco UCS Servers for MUR System Application Note*.

Storage Recommendation for MUR Application

This section provides the storage recommendations needed for the MUR application.

- Separate storage (single disk or RAID array) for OS. (root and swap space partitions)
- **Two RAID arrays:** RAID-0 for MUR application and RAID-5 for database (Postgres data directory).
- **LVM:** Separate physical volume and volume groups for the three RAID array disk groups.
- **XFS file-system:** block size 4 KB, s-unit in terms of RAID stripe size (256KB) and s-width in terms of span of disks in the RAID array.

For information on how to partition storage disk and configure XFS file system, refer to the *Configuring Cisco UCS Servers for MUR System Application Note*.

Hardware Requirements for Scalable Model of MUR

- UCS 460 Server or Oracle x86 4270 Server
- Sun StorageTek 2540 SAS Array
- Other miscellaneous components such as cables, switches, etc same as mentioned in the previous sections of this guide.

Software Requirements for Scalable Model of MUR

- Veritas Cluster 5.1

- Veritas cluster suite
- VxFS as cluster file system
- Veritas Cluster Logical Volume Manager

NOTE: Solaris OS clustering is NOT recommended for the following reasons:

- If there are any deployments that have a Solaris machine running Solaris OS (either as Master or as RDP or both), and if you want to upgrade to scalable model, following cases arise.
 - In a standalone installation —
 - Keep existing machine as master and have Linux machines as RDPs. If the RDPs are Linux, then Veritas cluster would do.
 - Another option is to add all Solaris machines with Solaris OS. This is not recommended as Solaris machine would be much more costlier than UCS.
 - If hierarchical installation exists —
 - Keep master as Solaris OS. For the RDP, there are two possibilities —
 - Add another RDP as Solaris OS and have Solaris cluster. This is not recommended as Solaris machine would be costlier than UCS.
 - If the existing Solaris machine is NOT Linux compatible (e.g. SPARC machine), it would be good to replace current Solaris machine with UCS and purchase another UCS for scalability purpose.
 - If the existing Solaris machine is Linux compatible (e.g. x86 architecture), it would be good to upgrade it to RHEL and have it as cluster.

MUR Ports

This section provides information on various ports and their corresponding port numbers used by the MUR application.

Various ports are used by the MUR for both client-server communication and communication with ASR chassis. If firewalls are used on these interfaces, these ports need to be opened.

The following table lists the ports that are used by MUR.

Table 1. Default Port Utilization

Port Name	Port Number	Usage
TCP Port	22	This port is used by MUR administrator to connect via SSH to UNIX command line on MUR servers for system administration. This port is also used by gateway to upload files via SFTP to MUR servers (stand-alone master and RDPs), and also by RDPs to upload files to the master. In the case of pull model, the L-ESS process on the RDPs or stand-alone master will use SFTP to connect to this port on the gateway. This port is also used between master MUR server and gateway to configure and upload bulkstat files.
TCP Port	25	This port is used to send e-mails to a mail server in case these are configured to deliver reports and alarms.
UDP Port	162	This port is used to send traps to the northbound network management system.
Postgres Port	5432	This port is used by the local processes to access the PostgreSQL server and can be restricted to prevent external access.
Apache Port	8080	For a standalone model: This port is used for communication between client workstation and Apache Webserver on MUR via HTTP. For distributed model: This port is used for both Master to RDP and RDP to Master RPC communication.  Important: When firewall is used, Apache is the only port that should be kept opened.

Typically, MUR starts all its related services with non-root (i.e. muradmin) privileges.

Firewall Settings

When MUR is running on RHEL platform, Firewall is ON by default. In that case, user will NOT be able to get access to MUR GUI. The Firewall MUST be disabled with the following commands:

```
service iptables save

service iptables stop

chkconfig iptables off
```

Using Apache Port

This section provides information on how to configure the Apache port to use in conjunction with the MUR reporting server.

Using Apache in Solaris

In case the user wants to configure Apache port as 80 (i.e. < 1024), it is necessary to run the following command as **root** user so that *muradmin* can start the services on ports < 1024.

```
usermod -K defaultpriv=basic,net_privaddr <mur admin user>
```

Using Apache in RHEL



Important: Make sure that you disable Firewall before using the Apache port in the RHEL environment.

RHEL does not allow port 80 to be used by non-root users. However, Apache Web server requests made on port 80 can be redirected to a port >1024 defined by the operator, with the following two commands:

```
iptables -t nat -A PREROUTING -p tcp --dport 80 -i eth0 -j REDIRECT --to-port <user  
defined port> 1024>
```

```
iptables -t nat -A OUTPUT -p tcp -d 127.0.0.1 --dport 80 -j REDIRECT --to-port <user  
defined port> 1024>
```

For example, to redirect requests made on port 80 to port 8080:

```
iptables -t nat -A PREROUTING -p tcp --dport 80 -i eth0 -j REDIRECT --to-port 8080
```

```
iptables -t nat -A OUTPUT -p tcp -d 127.0.0.1 --dport 80 -j REDIRECT --to-port 8080
```

Once this is done, user will be able to access the MUR GUI directly, without specifying the port in the Web browser URL *http://<serveripaddress>*.

Chapter 2

Configuring Chassis for Mobility Unified Reporting System

This chapter describes the configurations required to source data for the MUR application.



Important: These configurations are on the chassis.

For more information on ECS configurations, see the *Enhanced Charging Services Administration Guide*.

This chapter describes the following topics:

- [Initial Configuration](#)
- [Configuration](#)

Initial Configuration

If the configurations described in this section are not already available on the system, these must be configured.

Initial configuration steps:

- Step 1** Ensure that ECS license is installed on the system.
- Step 2** Create the ECS administrative user account as described in the [Creating the ECS Administrative User Account](#) section.
- Step 3** Enable Active Charging as described in the [Enabling Active Charging](#) section.
- Step 4** Create Active Charging Service as described in the [Creating the Active Charging Service](#) section.
- Step 5** Save your configuration to flash memory, an external memory device, and/or a network location using the Exec mode command **save configuration**. For additional information on how to verify and save configuration files, refer to the *System Administration Guide* and the *Command Line Interface Reference*.



Important: Commands used in the configuration examples in this section provide base functionality to the extent that the most common or likely commands and/or keyword options are presented. In many cases, other optional commands and/or keyword options are available. Refer to the *Command Line Interface Reference* for complete information regarding all commands.

Installing the ECS License

The ECS in-line service is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on licensing requirements.

For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Creating the ECS Administrative User Account

At least one administrative user account that has ECS functionality privileges must be configured on the system. This is the account that is used to log on and execute ECS-related commands. For security purposes, it is recommended that these user accounts be created along with general system functionality administration.

Use the following configuration example to create the ECS Administrative user account:

```
configure
  context local
    administrator <user_name> password <password> ecs
  end
```

Notes:

- Aside from having ECS capabilities, an ECS Administrator account also has the same capabilities and privileges as any other system-level administrator account.
- You can also create system ECS user account for a config-administrator, operator, or inspector. ECS accounts have all the same system-level privileges of normal system accounts except that they have full ECS command execution capability. For example, an ECS has rights to execute every command that a regular administrator can in addition to all of the ECS commands.
- Note that only Administrator and Config-administrator-level users can provision ECS functionality. Refer to the *Configuring System Settings* chapter of the *System Administration Guide* for additional information on administrative user privileges.

Enabling Active Charging

Active Charging must be enabled before configuring charging services.

Use the following configuration example to enable Active Charging:

```
configure

require active-charging

context local

interface <interface_name>

    ip address <ipv4/ipv6_address> <ipv4/ipv6_address/mask>

    exit

server ftpd

end
```

For more information, refer to the *Enhanced Charging Services Administration Guide*.

Creating the Active Charging Service

Use the following configuration example to create an Active Charging Service:

```
configure

active-charging service <service_name> [ -noconfirm ]

end
```

Configuration

The following is the sequence of configurations necessary to source data to the MUR application:

- Step 1** Activate P2P analyzer as described in the [Activating P2P Analyzer](#) section.
- Step 2** Configure EDR flow format as described in the [Configuring the EDR Flow Format](#) section.
- Step 3** Configure routing ruledefs and rulebase for deep-packet inspection as described in the [Configuring Deep Packet Inspection](#) section.
- Step 4** Optional. Configure Smartphone tethering detection feature as described in the [Configuring Tethering Detection Feature](#) section.
- Step 5** Configure EDR module as described in the [EDR Module Configuration](#) section.
- Step 6** Configure user as described in the [Configuring EDR Download Permission](#) section.
- Step 7** Configure the bulkstat schemas and then load it onto the gateway.
- Step 8** Save your configuration to flash memory, an external memory device, and/or a network location using the Exec mode command **save configuration**. For additional information on how to verify and save configuration files, refer to the *System Administration Guide* and the *Command Line Interface Reference*.

Activating P2P Analyzer

Use the following configuration example to activate P2P protocol detection:

```
configure

active-charging service <service_name>

    p2p-detection protocol all

    rulebase <rulebase_name>

        p2p dynamic-flow-detection

    end
```

Notes:

- P2P protocol detection must be activated only within rulebases used by the APNs for which P2P detection is applicable. P2P detection must not be applied to the rulebases used for APNs where such reporting is either not useful or is not possible.

Configuring the EDR Flow Format

Use the following configuration example to configure the EDR format generated for flows:

```
configure
```

```

active-charging service <service_name>

    edr-format <edr_format_name> [ -noconfirm ]

        attribute <attribute> { [ format { MM/DD/YY-HH:MM:SS | MM/DD/YYYY-HH:MM:SS |
YYYY/MM/DD-HH:MM:SS | YYYYMMDDHHMMSS | seconds } ] [ localtime ] | [ { ip | tcp } { bytes
| pkts } { downlink | uplink } ] priority <priority> }

        rule-variable <protocol> <rule> priority <priority>

        rule-variable traffic-type priority <priority>

        rule-variable voip-duration priority <priority>

        event-label <event-label> priority <priority>

    end

```

Notes:

- The **rule-variable traffic-type** and **rule-variable voip-duration** keywords must be configured to enable voice-call-duration (VCD) based reporting.
- *p2p-protocol* is a mandatory field in a flow-edr configurations. However, this field cannot be added to the edr-format configuration unless P2P is licensed. Contact your Cisco account representative for information on how to obtain a license.
- For information on EDR format configuration and rule variables, refer to the *EDR Format Configuration Mode Commands* chapter of the *Command Line Interface Reference*.

The following is a sample flow end EDR configuration.

```

configure

active-charging service ecs_svc1

    edr-format edr_flow_format

        attribute sn-start-time format seconds priority 10

        attribute sn-end-time format seconds priority 20

        attribute radius-calling-station-id priority 30

        rule-variable ip server-ip-address priority 60

        attribute sn-server-port priority 70

        attribute sn-app-protocol priority 80

        attribute sn-parent-protocol priority 81

        rule-variable ip protocol priority 82

        rule-variable p2p protocol priority 90

        attribute sn-volume-amt ip bytes uplink priority 100

```

■ Configuration

```

attribute sn-volume-amt ip bytes downlink priority 110
attribute sn-volume-amt ip pkts uplink priority 120
attribute sn-volume-amt ip pkts downlink priority 130
rule-variable bearer 3gpp charging-id priority 140
rule-variable bearer 3gpp imei priority 141
rule-variable bearer 3gpp rat-type priority 142
rule-variable bearer 3gpp user-location-information priority 143
rule-variable traffic-type priority 160
rule-variable voip-duration priority 170
end

```

The following is a sample HTTP EDR configuration.

```

configure
active-charging service ecs_svc1
edr-format edr_http_format
attribute sn-start-time format seconds priority 10
attribute sn-end-time format seconds priority 20
attribute radius-calling-station-id priority 30
rule-variable ip server-ip-address priority 50
rule-variable http host priority 70
rule-variable http content type priority 80
attribute transaction-downlink-bytes priority 90
attribute transaction-uplink-bytes priority 100
attribute transaction-downlink-packets priority 110
attribute transaction-uplink-packets priority 120
rule-variable bearer 3gpp charging-id priority 130
end

```

Verifying your Configuration

To verify your configuration, in the Exec Mode, enter the following command:

```
show active-charging edr-format name <edr_format_name>
```

Configuring Deep Packet Inspection

This section provides the example configurations that are required for deep packet inspection.

Configuring Routing Rule Definition

Use the following configuration example to create and configure a routing ruledef:

```
configure

active-charging service <service_name>

    ruledef <ruledef_name>

        <protocol> <expression> <operator> <condition>

        rule-application routing

    end
```

Notes:

- The **rule-application routing** command specifies the ruledef type. If not specified, by default, the system configures the ruledef as a charging ruledef.
- For information on all the protocol types, expressions, operators, and conditions supported, refer to the *Ruledef Configuration Mode Commands* chapter of the *Command Line Interface Reference*.
- Up to 10 rule matches can be configured in one ruledef.
- MMS rules must be set appropriately and MMS should be activated at ECS to support MMS reporting in MUR.

The following is a sample ruledef configuration.

```
configure

active-charging service srv1

    ruledef http_anymatch

        http any-match = TRUE

    exit

    ruledef icmp_anymatch

        icmp any-match = TRUE

    exit

    ruledef ip_anymatch

        ip any-match = TRUE
```

■ Configuration

```
exit

ruledef mms_anymatch

mms any-match = TRUE

exit

ruledef rr_http_80

tcp either-port = 80

rule-application routing

exit

ruledef rr_http_8080

tcp either-port = 8080

rule-application routing

exit

ruledef rr_mms_http_ct

http content type = application/vnd.wap.mms-message

rule-application routing

exit

ruledef rr_mms_http_url

http url ends-with .mms

rule-application routing

exit

ruledef rr_mms_wsp_ct

wsp content type = application/vnd.wap.mms-message

rule-application routing

exit

ruledef rr_mms_wsp_ct_uri

rule-application routing

exit

ruledef rr_mms_wsp_url

wsp url ends-with .mms
```

```
rule-application routing
exit

ruledef rr_wsp_cl_dst_port
udp dst-port = 9200
rule-application routing
exit

ruledef rr_wsp_cl_src_port
udp src-port = 9200
rule-application routing
exit

ruledef rr_wsp_co_dst_port
udp dst-port = 9201
rule-application routing
exit

ruledef rr_wsp_co_src_port
udp src-port = 9201
rule-application routing
exit

end
```

Verifying your Configuration

To verify your configuration, in the Exec Mode, enter the following command:

```
show active-charging ruledef routing
```

Configuring Rulebase

Use the following configuration example to route traffic to the appropriate analyzer within each rulebase where the reporting is applicable.

```
configure
```

```
active-charging service <service_name>
```

```
rulebase <rulebase_name> [ -noconfirm ]
```

```

    route priority <priority> ruledef <ruledef_name> analyzer <analyzer> [ description
]

    rtp dynamic-flow-detection

    flow end-condition handoff timeout normal-end-signaling session-end [ charging-edr
| edr | reporting-edr edr_format_name ]

    edr transaction-complete http [ charging-edr | edr-format | reporting-edr
edr_format_name ]

end

```

Notes:

- **charging-edr** will be used as the default option in the **flow end-condition** and **edr transaction-complete** command configurations.
- The **edr** and **edr-format** options are available only in 12.1 and earlier releases. In 12.2 and later releases, these options are deprecated and are replaced by the **charging-edr** option.
- For MUR reporting needs, use the **reporting-edr** keyword in the rulebase configuration.

The following is a sample rulebase configuration for reporting EDRs.

configure

```

    active-charging service ecs_svc1

    rulebase p2p-rb

        flow end-condition handoff timeout normal-end-signaling session-end reporting-edr
edr_flow_format

        action priority 4 ruledef rtsp_setup charging-action standard
        action priority 5 ruledef rtsp_play charging-action standard
        action priority 6 ruledef rtsp_tearardown charging-action standard
        action priority 7 ruledef rtsp_anymatch charging-action standard
        action priority 10 ruledef sip_anymatch charging-action handshake
        action priority 11 ruledef rtp_anymatch charging-action handshake
        action priority 12 ruledef udp_anymatch charging-action handshake
        action priority 13 ruledef tcp_anymatch charging-action handshake
        action priority 16 ruledef mms_anymatch charging-action policy1
        action priority 60 ruledef http_anymatch charging-action standard
        action priority 95 ruledef udp_anymatch charging-action standard
        action priority 99 ruledef icmp_anymatch charging-action standard

```



```

action priority 100 ruledef ip_anymatch charging-action handshake
action priority 990 ruledef tcp_anymatch charging-action standard
action priority 1000 ruledef ip_anymatch charging-action standard
route priority 1 ruledef rr_wsp_co_src_port analyzer wsp-connection-oriented
route priority 2 ruledef rr_wsp_co_dst_port analyzer wsp-connection-oriented
route priority 3 ruledef rr_wsp_cl_src_port analyzer wsp-connection-less
route priority 4 ruledef rr_wsp_cl_dst_port analyzer wsp-connection-less
route priority 5 ruledef rr_http_80 analyzer http
route priority 6 ruledef rr_http_8080 analyzer http
route priority 7 ruledef rr_mms_http_ct analyzer mms
route priority 8 ruledef rr_mms_http_url analyzer mms
route priority 9 ruledef rr_mms_wsp_ct analyzer mms
route priority 10 ruledef rr_mms_wsp_url analyzer mms
route priority 11 ruledef rr_mms_wsp_ct_uri analyzer mms
route priority 60 ruledef sip_src analyzer sip
route priority 65 ruledef sip_dst analyzer sip
route priority 70 ruledef rtsp_src analyzer rtsp
route priority 75 ruledef rtsp_dst analyzer rtsp
route priority 250 ruledef sdp_route analyzer sdp
rtp dynamic-flow-detection
edr transaction-complete http reporting-edr edr_http_format
edr voip-call-end reporting-edr edr_flow_format
udr threshold interval 60
udr threshold volume total 100000
p2p dynamic-flow-detection
end

```

Verifying your Configuration

To verify your configuration, in the Exec Mode, enter the following command:

```
show active-charging rulebase name <rulebase_name>
```

Configuring Charging Action

Use the following configuration example to configure a charging action:

```
configure

  active-charging service <service_name>

    charging-action <charging_action_name> [ -noconfirm ]

    content-id <content_id>

    retransmissions-counted

    billing-action [ edr <edr_format> [ wait-until-flow-ends ] | egcdr | exclude-from-
udrs | radius ] +

    flow idle-timeout <idle_timeout>

  end
```

Verifying your Configuration

To verify your configuration, in the Exec Mode, enter the following command:

```
show active-charging charging-action name <charging_action_name>
```

Configuring Tethering Detection Feature

This section describes how to configure the Tethering Detection feature to detect subscriber flows from PC devices tethered to mobile smartphones. For details on how this feature is implemented, see the *Enhanced Charging Services Administration Guide*.

To enable and configure the Tethering Detection feature, use the following configuration:

```
configure

  active-charging service <ecs_service_name>

    tethering-database [ os-signature <os_signature_db_file_name> | tac
<tac_db_file_name> | ua-signature <ua_signature_db_file_name> ] +

    ruledef <tethering_detection_ruledef_name>

      tethering-detection { flow-not-tethered | flow-tethered }

      exit

    rulebase <rulebase_name>

      tethering-detection [ os-db-only | ua-db-only ]

      action priority <priority> ruledef <tethering_detection_ruledef_name> charging-
action <charging_action_name>
```

```
...
end
```

Upgrading Tethering Detection Databases

To upgrade the Tethering Detection feature databases, in the Exec mode, use the following CLI command:

```
upgrade tethering-detection database { all | os-signature | tac | ua-signature } [ -
noconfirm ]
```

Sample Configurations

The following examples illustrate two different implementations of the Tethering Detection feature's configuration.

- The following type of configuration is suitable where ECS performance is critical and the operator wants to put in a flat charging plan in place for all the tethered traffic. In such a scenario, addition of a single new ruledef to the configuration suffices. Placing this ruledef at the highest priority in the rulebase will ensure all the tethered flows are charged as per the tariff plan for tethered traffic.

```
configure

  active-charging service ecs_service

    tethering-database

    ruledef tethered-traffic

      tethering-detection flow-tethered

      tcp any-match = TRUE

    exit

  ruledef ftp-pkts

    ftp any-match = TRUE

  exit

  ruledef http-pkts

    http any-match = TRUE

  exit

  ruledef tcp-pkts

    tcp any-match = TRUE

  exit

  ruledef ip-pkts

    ip any-match = TRUE
```

```

        exit
    ruledef http-port
        tcp either-port = 80
        rule-application routing
    exit
    ruledef ftp-port
        tcp either-port = 21
        rule-application routing
    exit
    charging-action premium
        content-id 1
        retransmissions-counted
        billing-action egcdr
    exit
    charging-action standard
        content-id 2
        retransmissions-counted
        billing-action egcdr
    exit
    rulebase consumer
        tethering-detection
        action priority 10 ruledef tethered-traffic charging-action
premium
        action priority 20 ruledef ftp-pkts charging-action standard
        action priority 30 ruledef http-pkts charging-action standard
        action priority 40 ruledef tcp-pkts charging-action standard
        action priority 50 ruledef ip-pkts charging-action standard
        route priority 80 ruledef http-port analyzer http
    exit

```

```
rulebase default

end
```

- The following type of configuration is suitable when operators want to apply differentiated charging to various flows that are found to be tethered. In this case, traffic that requires different charging action or content ID when it is tethered will be identified using two ruledefs, one with “flow-is-tethered = TRUE” option and another without this option. This configuration provides finer granularity of control but results in higher performance degradation because the rule matching tree size increases.

```
configure

active-charging service ecs_service

tethering-database

ruledef ftp-pkts

    ftp any-match = TRUE

    exit

ruledef ftp-pkts-tethered

    ftp any-match = TRUE

    tethering-detection flow-tethered

    exit

ruledef http-pkts

    http any-match = TRUE

    exit

ruledef http-pkts-tethered

    http any-match = TRUE

    tethering-detection flow-tethered

    exit

ruledef tcp-pkts

    tcp any-match = TRUE

    exit

ruledef tcp-pkts-tethered

    tcp any-match = TRUE

    tethering-detection flow-tethered
```

```
exit
ruledef ip-pkts
    ip any-match = TRUE
    exit
ruledef ip-pkts-tethered
    ip any-match = TRUE
    tethering-detection flow-tethered
    exit
ruledef http-port
    tcp either-port = 80
    rule-application routing
    exit
ruledef ftp-port
    tcp either-port = 21
    rule-application routing
    exit
charging-action premium-http
    content-id 10
    retransmissions-counted
    billing-action egcdr
    exit
charging-action premium-ftp
    content-id 20
    retransmissions-counted
    billing-action egcdr
    exit
charging-action premium
    content-id 1
    retransmissions-counted
```

```

        billing-action egcdr
    exit

    charging-action standard
        content-id 2
        retransmissions-counted
        billing-action egcdr
    exit

    rulebase consumer
        tethering-detection
        action priority 10 ruledef ftp-pkts-tethered charging-action
premium-ftp
        action priority 20 ruledef ftp-pkts charging-action standard
        action priority 30 ruledef http-pkts-tethered charging-action
premium-http
        action priority 40 ruledef http-pkts charging-action standard
        action priority 50 ruledef tcp-pkts-tethered charging-action
premium
        action priority 60 ruledef tcp-pkts charging-action standard
        action priority 70 ruledef ip-pkts-tethered charging-action
premium
        action priority 80 ruledef ip-pkts charging-action standard
        route priority 80 ruledef http-port analyzer http
    exit

    rulebase default
end

```

EDR Module Configuration

Use the following configuration example to configure the EDR module:

configure

context <context_name>

edr-module active-charging-service [charging | reporting]

```

    file name <file_name> rotation volume <file_size_bytes> rotation time
    <file_complete_seconds> rotation num-records <records_number> storage-limit
    <storage_limit_bytes> headers reset-indicator edr-format-name trap-on-file-delete
    compression gzip file-sequence-number rulebase-seq-num

    cdr [ push-interval <interval> | remove-file-after-transfer | transfer-mode { pull
    | push primary { encrypted-url <enc_url> | url <url> } [ secondary { encrypted-secondary-
    url <enc_sec_url> | url <sec_url> } ] } + | use-harddisk ]

    end

```

Notes:

- The `<context_name>` must be the context specified for accounting.
- EDR type configuration is optional. The EDR types can be either **charging** or **reporting**. The **charging** keyword is the default setting.
For MUR reporting needs, use the **reporting** keyword for the EDR type.
- The **cdr use-harddisk** command is only available on the ASR 5000 platform.
- The **cdr use-harddisk** command specifies storing files on the hard disk. The reporting server will download these files through the SPIO interface on the SMC and will delete the files after successful retrieval.
- The **edr-format-name** keyword must be configured to distinguish between different EDRs. The EDR file name must be configured in an accepted format so that the Offline Subscriber Reporting functionality can be used effectively. For information on this functionality and the EDR file name configuration recommendations, see the *Cisco Mobility Unified Reporting System Online Help* documentation.
- The files will be compressed to save storage and transmission bandwidth.
- For the PULL model, an external device like L-ESS is used to pull the EDR files from the chassis via SFTP. Whereas, for the PUSH model, the chassis is configured to push the files to the required destination.



Important: The chassis automatically creates `/edr` and `/udr` directories on the destined path on MUR server when you configure it to push the files.

- The values recommended for **rotation volume** and **rotation time** keywords are 40 MB and 300 seconds respectively.



Important: In RHEL-based deployments, L-ESS is NOT required as the Enhanced Charging Services (ECS) module can be configured to push the xDRs directly to the MUR reporting server. Push from ASR chassis is the Cisco recommended deployment model. Currently L-ESS is supported only on Solaris platforms. For information on the L-ESS installation instructions, refer to the *ESS Installation and Administration Guide*. Existing deployments where L-ESS is installed, to pull EDRs from chassis, may continue with their deployment model in the 12.0 version of MUR Software Release and later.

The following is a sample EDR PUSH configuration.

```

configure

context test

    edr-module active-charging-service reporting

```



```

file name EDRFILE rotation num-records 10000 storage-limit 268435456 headers reset-
indicator trap-on-file-delete compression gzip file-sequence-number rulebase-seq-num

cdr transfer-mode push primary url sftp://root:nulink@10.4.72.54/inpilot-
local/Ash_Test/starbi/server/data via local-context

cdr push-interval 60

cdr remove-file-after-transfer

cdr use-harddisk

end

```

The following is a sample EDR PULL configuration.

```

configure

context local

edr-module active-charging-service

file name EDRFILE1 rotation time 300 rotation num-records 10000 storage-limit
268435456 headers reset-indicator trap-on-file-delete compression gzip file-sequence-
number rulebase-seq-num

cdr remove-file-after-transfer

cdr use-harddisk

end

```

Verifying your Configuration

To verify your configuration, in the Exec Mode, enter the following command:

```
show configuration context <context_name>
```

Pushing EDR/UDR Files Manually

To manually push EDR/UDR files to the configured L-ESS, in the Exec mode, enter the following command:

```
cdr-push { all | local-filename <file_name> }
```

Notes:

- Before you can use this command, in the EDR/UDR Configuration Mode, the CDR transfer mode and file locations must be set to push.
- <file_name> must be absolute path of the local file to push.

For more information on the **cdr** command, please refer to the *Command Line Interface Reference*.

Configuring EDR Download Permission

Use the following configuration example to configure EDR download permission:

```
configure

context local

    administrator <administrator_id> password <password> ftp nocli

end
```

Notes:

- The user must be configured in the local context with administrative privileges to download and delete EDRs from the hard disk. The **ftp nocli** options restrict access to FTP only.

Configuring Bulkstats Schemas Using GUI

MUR provides a user interface to configure bulk statistics schemas on chassis / gateway via SSH and SFTP. The Client sends HTTP request to MUR to configure schemas on a particular gateway after providing inputs to the parameters needed for schema configuration. MUR server receives the HTTP request, generates a configuration file on the fly, sends the configuration file to the gateway via SFTP and loads it on to the gateway through SSH.

 **Important:** In StarOS 10.0 and earlier releases, WEM is used to configure the bulkstats schemas on the chassis if user has deployed WEM. In case if WEM has not been deployed, then please contact your Cisco account representative for obtaining the embedded bulkstats configuration file.

 **Important:** In StarOS 11.0 and later releases, you can configure Bulkstat schemas only through the MUR GUI by selecting **ADMIN > BULKSTATS** menu.

Prior to configuring the bulkstats schemas, ensure that the following checks are performed:

- The gateway must be running and active.
- Enable SFTP and FTP services

For Solaris setup:

- FTP must be enabled on the MUR server.

To enable the FTP daemon, use the following command:

```
/usr/sbin/svcadm/ enable ftp
```

To disable the FTP daemon, use the following command:

```
/usr/sbin/svcadm/ disable ftp
```

For RHEL setup:

- FTP must be enabled on the MUR server.

To enable the FTP daemon, use the following command:

```
service vsftpd start
```

To disable the FTP daemon, use the following command:

```
service vsftpd stop
```

- SSH version 2.0 key must be generated on the gateway. To generate the SSH version 2.0 key through the CLI, enter the following command:

```
configure

context local

    ssh generate key type v2-rsa

    ssh generate key type v2-dsa

end
```

- Secure Shell (SSH) configuration mode must be enabled on the gateway. To enable the SSH configuration mode, enter the following command:

```
configure

context local

    server sshd

end
```

- FTP/SFTP must be allowed on the gateway for the “SSH Username” that will be entered in the Bulkstat Schema Configuration screen. For example, if the username is *staradmin* and password is *test* then the following commands should be used to enable FTP/SFTP for *staradmin* user.

```
configure

context local

    administrator staradmin password test ftp

end
```



Important: The bulkstats report will be visible to users only when the schemas are configured successfully.

For information on how to configure the bulkstats schemas, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Supported Bulkstat Schemas

This section provides the list of bulk statistics schemas that are supported in MUR for reporting.

- SS7RD
- MME
- SGW
- SAEGW
- MIPV6HA
- PGW
- IMSA
- NAT_REALM
- ASNGW
- PORT
- SGSN
- MISC
- CARD
- MIPFA
- GTPP
- PHSGW
- CSCFINTF
- RADIUS
- RADIUS_GROUP
- APN
- CLOSEDRP
- LAC
- SGTP
- IPPPOOL
- SCCP
- GPRS
- SS7LINK
- CSCF
- MAG
- CONTEXT
- SYSTEM
- ECS
- PHSPC

- EGTPC
- RP
- PPP
- MIPHA
- PDG
- GTPC
- PDIF
- IPST
- LMA
- AAL2
- ALCAP
- ASNPC
- BCMCS
- CS_NW_RANAP
- CS_NW_RTP
- DCCA
- DPCA
- GTPU
- HNBGW_HNBAP
- HNBGW_RANAP
- HNBGW_RTP
- HNBGW_RUA
- HNBGW_SCTP
- LNS
- PCC_POLICY
- PCC_QUOTA
- PCC_SERVICE
- PCC_SP_ENDPT
- PS_NW_RANAP
- MVS
- Diameter
- SBC
- SLS
- map_bk
- hnbgw_cbs_sabp
- hnbgw_cbs_tcp

■ Supported Bulkstat Schemas

- hnbgw_sabp
- hnbgw_sabp_closed
- hnbgw_sabp_hybrid
- hnbgw_sabp_open
- vlan_npu
- ps_nw_gtpu
- mme_tai
- iups_bk
- rlf
- rlf_detailed
- gprs_bk
- sgtp_bk
- mme_bk
- APN QCI Duration
- ePDG APN QCI Duration
- S102

For more information on these bulkstats, refer to the *Statistics and Counters Reference*.

All the bulkstats counters mentioned in the *Statistics and Counters Reference* for particular schema might not be present due to the limitation of ASR 5K bulkstat implementation.

Supported SNMP Traps

The alarm generation feature aids in proactively monitoring the nodes and important resources of MUR. This feature also provides configuration interface for setting up thresholds and other key information related to critical resources. Alarms are generated when these thresholds are exceeded and various actions can be performed such as sending e-mail, syslog messages, Simple Network Management Protocol (SNMP) traps.

It is necessary to configure the SNMP manager or Network Node Manager (NNM) to receive these notifications. The SNMP server and SNMP event configurations can be made through the **System** menu in the Web-based MUR GUI.

Threshold values should be configured for the following event identifiers (event IDs):

- CPU Usage - *CPU* — This alarm is generated when the CPU resource usage exceeds the preset threshold value.
- Postgres Disk Usage - *PostgresDisk* — This alarm is generated when the Postgres disk usage exceeds the threshold.
- Memory (Swap) Usage - *Mem* — This alarm is generated when the memory swap usage exceeds the threshold.
- Unprocessed Files - *UnprocFiles* — This alarm is generated when the count of (HTTP-EDR/EDR/CF-EDR) files pending for getting parsed (in their respective directories), exceeds the threshold value.
- Erroneous Files - *ErrFiles* — This alarm is generated whenever the count of invalid files exceeds the threshold value. The file is considered as invalid either due to missing headers or the file being corrupted.
- Erroneous Records - *ErrRecords* — This alarm is generated when the number of erroneous records breaches the threshold. The EDR records are considered as erroneous when any of the fields are missing in the EDR or when an invalid data is present in a particular field.
- Application Disk - *AppDisk* — This alarm is generated when the MUR application disk usage exceeds the threshold.
- Archive Disk - *ArchiveDisk* — This alarm is generated when the disk space in the Archive Data directory path reaches the threshold.
- Additional Disk Path - *AdditionalDiskPath* — This alarm is generated when the usage of any other additional disks exceeds the threshold.

In this release, MUR monitors the disk utilization on both MASTER and RDP, and generates SNMP alarms accordingly. When the threshold in disk alarm is set to the value of MASTER, which has already been crossed, this threshold value will not be propagated to RDP, as scheduler will be stopped as soon as it detects that disk threshold is breached in MASTER. Due to which the communication from MASTER and RDP will be stopped and the RDP will not have postgres threshold of 99. But if MASTER/RDP with values prior to scheduler stopped scenario is tested, generation of alarms in both MASTER and RDP will be successful.

In addition to these events, MUR also supports *AppStatus* and *TaskLag* event identifiers; however, these are NOT configurable.

- Application Status - *AppStatus* — This alarm is generated when the MUR application is started or stopped.



Important: Please note that the alarms are sent when Scheduling server/Apache server is started or stopped. However, in the case of Postgres server, alarms are sent only when it is started.

- Task Lag - *TaskLag* — This alarm is generated when a particular script like normalization/aggregation takes more time than expected to complete the job.

The following scripts have been added for the Task Lag alarms, which play an important role in parsing EDR/HTTP-EDR/CF-EDR. Each of these scripts handle a specific task which is either part of aggregation or normalization.

Script	Default Value of Tasklag Time (in sec/min)
Edr Normalization	300 (5 min)
Http Edr Normalization	300 (5 min)
CF Edr Normalization	300 (5 min)
Protocol Summary	1800 (30 min)
Port Aggregation	1800 (30 min)
Subscriber Aggregation (minutely)	1800 (30 min)
Subscriber Aggregation (hourly)	3600 (60 min)
Flow Count	1800 (30 min)
Http Host Aggregation (minutely)	7200 (120 min)
Http Content Summary	7200 (120 min)
Http Host Aggregation (hourly)	7200 (120 min)



Important: Please note that the user does not have the privilege to change these timings.

Note that these alarms (Unprocessed, Error Files, Error Records and TaskLag) can be triggered only for EDR, HTTP-EDR and CF-EDR files, and not for the bulkstats files.

In Release 14.0 and later, the Unprocessed alarms have been extended to the RDPs. It monitors the RDP directory from where the files are exported to the Master. Both the Master and RDP share the same threshold value configured through GUI.

Also for the RDP disk alarm, RDP monitors the default starbi installation path and that cannot be modified. Please note RDP alarms are disabled by default same as that of Master.

Master and RDP share the same SNMP configuration window and are enabled at the same time (provided RDPs are configured properly). In the event of SFTP file transfer failure from RDP to Master, SFTP alarms will be generated. This alarm could also be generated due to one of the following reasons:

- Authentication failed for the host
- SSH negotiation failed while connecting to the host
- Remote Host not defined
- Socket timed out while setting up connection
- Directory creation failed on Master
- Removing the file at RDP failed



Important: During a fresh installation of MUR, please note that there will no SNMP configurations available.



Important: Users with administrative privilege can only manage this configuration.



Important: The change in the configuration for enabling / disabling the alarm generation feature does not require a restart of the MUR application.

MUR also supports generation of KPI alarms through the GUI. KPI parser calculates the values of KPIs for which the alarms are configured through the GUI. The KPI parser uses the information stored by bulkstat parser in the database for KPI calculations and for sending alarms. This avoids reparsing of the same file and redundant connections to the DB.

KPI parser generates alarms only when the alarm functionality is enabled for MUR. The details of KPI alarms which are successfully sent can be seen through **KPI Alarms Log** under the **System** menu. For details on the log, see the *Cisco Mobility Unified Reporting System Online Help* documentation.



Important: Prior to configuring KPI alarms, you must ensure that the gateways and bulkstat schemas are configured and the bulkstats data are available.



Important: In the case of HA mode installation, the floating IP will be used for sending SNMP V1 alarms. For V2, there is no way to explicitly set the agent ip-address. Hence, it will still use the local node IP for sending V2 alarms.

For information on configuring the SNMP parameters, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

For information on the SNMP traps and thresholds supported for MUR, see the *Mobility Unified Reporting System MIB* chapter of the *SNMP MIB Reference*.

Chapter 3

Configuring MegaRAID for MUR Applications

CISCO UCS servers support a MegaRAID controller. This allows for configuring the UCS hard disks into hardware RAID arrays. The MegaRAID controller provides the BIOS utility for configuring the RAID. This chapter describes how to configure MegaRAID; this is for Mobility Unified Reporting System (MUR) applications only.

This chapter comprises the following topics:

- [Option Recommendations](#)
- [Sample Configuration](#)
- [Setting the Boot Drive](#)

Option Recommendations

The following options are recommended for MUR applications:

- Separate disk arrays for the Operating System, MUR, and the Postgres (data directory).
- RAID level. This is a combination of five (5) and zero (0) depending on the fault tolerance.
- Stripe size should be 256KB.

The following RAID Controller parameters for use in the Virtual Drive Definition screen are recommended:

- Read Policy - Adaptive Read-Ahead.
- Write Policy - Write Back.
- I/O Policy - Direct I/O.

Sample Configuration


In the following example, the UCS server has 12 disks each of 280GB. The disks are configured into the following three groups:

- Single (1) disk in RAID-0 for the OS.
- RAID-0 array with four (4) disks for the MUR application.
- RAID-5 array with seven (7) disks for the Postgres Database.

Creating Virtual Drives

Use the following steps to create virtual drives.

1. While the system is booting, press <Ctrl> <H> to start the WebBIOS utility.
2. Press **Start** to enter the utility. If the drives are already configured, the screen will show the current RAID configuration with virtual drives for all the RAID disk groups. Otherwise, the drives display as unconfigured.




3. From the left pane, select **Configuration Wizard**.




Important: In the event that there are already any unwanted virtual drives configured, select the **Clear Configuration** radio button in the Wizard to delete the default virtual drives.

4. Select the **New Configuration** radio button to start creating virtual drives. If you did not delete any drives in the previous step and want to configure them now, select the **Add Configuration** radio button.
5. Press **Next** and then select **Manual Configuration** to open the Drive Group Definitions screen.


Sample Configuration



6. Click on the first disk in the list (Slot 1). This disk is the single disk selected to be the virtual drive for the operating system and it will be added to Drive Group0. Select Drive Group0 and press the **Accept DG** button.




7. To configure the next drive group in this example, use the <Ctrl> key to select the next four disks (Slots 2 through 5) and press the **Add to Array** button. This creates Drive Group1.




8. Use the same method to select the remaining four disks and create Drive Group2. You are now ready to configure the spans. Spans are used to take Drive Groups and use them to create a nested RAID array.

1. Press **Next** to open the Span Definition screen. Because this example only considers RAID-0 and RAID-5 arrays, the span will be created with only one group.



2. Select Drive Group0 and press the **Add to Span** button when it activates.
3. Press **Next** to open the Virtual Drive Definition screen. This is where RAID parameters are configured for the Disk Groups created earlier.




4. Complete this screen using the example parameters found in the Option Recommendations section in this document.




Important: In the Select Size drop-down box, enter the disk size in either GB or TB. For information on exact sizes, the screen shows the possible sizes of the drive for the various RAID levels. Select the appropriate size as per the RAID level selected and press the **Accept** button.

5. Pressing the **Accept** button opens a confirmation page where you confirm a Battery Backup Unit (BBU) is installed when using the Write Back with BBU policy option in the Virtual Drive Definition screen in Step 3. Press **Yes** to confirm.
6. Press the **Back** button to return to the Span Definition screen. Repeat the steps to add Disk Group1 and Disk Group2 to separate spans and configure the RAID parameters.


Sample Configuration



7. Press the **Accept** button to confirm the configuration.
8. Save the configuration and the utility will initialize the drives and display a final screen displaying the new virtual drives.



9. If necessary, check the properties for the drives by pressing the **Back** button to go back to the Main menu.



Continue to the next section to set the boot drive.

Setting the Boot Drive

You created a separate drive for the operating system and added it to DriveGroup0 earlier. Use the following steps to set this drive as the boot drive.

1. From the Main menu, select the **Virtual Drives** option in the left pane.
2. Select the drive created for the operating system,
3. Select the **Set Boot Drive** radio button.
4. Press the **Go** button to change the settings.
5. Press the **Exit** button to save the changes and exit the utility.

Chapter 4

Cisco UCS Server Hardware Configuration for MUR Applications

This section describes procedures for configuring a Cisco UCS server running a custom Red Hat Enterprise Linux (RHEL) operating system. These procedures are for Mobility Unified Reporting System (MUR) applications only.

This chapter comprises the following topics:

- [Overview](#)
- [Prerequisites](#)
- [Storage Recommendations](#)
- [Disk Partitioning](#)
- [Creating Volume Groups and Partitions](#)
- [Configuring the XFS File System](#)

Overview



Important: This section is specific to MUR applications. For all other applications, please refer to the official UCS Server documentation.

This section explains two procedures: disk partition and creating an XFS file system. Disk partitioning is carried out as part of the RHEL OS installation. The XFS file system is configured afterwards.

Prerequisites

The following are prerequisites for this configuration:

- Cisco UCS 460-1 server with multiple hard disks.
- MegaRAID Controller.
- ISO for the RHEL OS.

Storage Recommendations

The following are recommendations for this configuration:

- Separate storage (either a single disk or RAID array for the OS (root and swap space partitions).
- Two RAID arrays: RAID-0 for the MUR application and RAID-5 for the Postgres database.
- LVM: separate physical volume and volume groups for the three RAID array disk groups.

For the XFS file system:

- Block size = 4KB.
- S-unit in terms of RAID strip size = 256KB
- S-width in terms of span of disks in the RAID array.




Important: The rest of this section assumes the RAID arrays have already been configured per the above recommendations.

Disk Partitioning

Follow these steps to partition a disk.

1. Start the RHEL installation. The Disk Partitioning screen opens showing the two available drives. These are the RAID arrays configured earlier.
2. Select both drives. Ensure the default option to remove Linux partitions on the selected drive and create the default layout is selected.
3. Select the Review and Modify Partitioning Layout checkbox. Press **Next**.
4. Press **Yes** to confirm the edits are required. Press **Next** to open the Volume Groups and Partitions configuration screen.



Continue to the next section to create volume groups and partitions.

Creating Volume Groups and Partitions

The following steps explain how to create the volume groups and partitions based on the following table:


 **Important:** In actual scenario, partitioning may differ based on the dimensioning and deployment model selection. Please refer to specific notes per deployment plan.

Table 2. Volume Groups and Partitions Data


Partition	Space (GB)	Logical Volume	File-System	Mount Point
VolGroup00 (on drive with RAID-0)				
Root (RHEL OS)	100	LogVol100	ext3	/
Swap Space	32	LogVol01	Swap	-
VolGroup01 (on drive with RAID-0)				
MUR Application	All remaining space on Volume Group	LogVol00	ext3	/apps
VolGroup02 (on drive with RAID-5)				
Database (Postgres data directory)	All the space in the Volume Group	LogVol00	ext3	/db

 **Important:** In the Disk Partition Screen the default configuration combines all the disks into a single volume group (VolGroup00) having two logical volumes LogVol00 and LogVol01 for '/' and **Swap** partitions respectively. These default partitions have to be deleted.

 **Important:** XFS file system is recommended for MUR application (/apps), Postgres (/db) and archive (/archive) partitions.



1. Select VolGroup00 and press the **Edit** button to re-configure the volume group and its logical volumes (partitions).



2. Delete the the default logical volumes (partitions for **/** and **Swap**) from the bottom window.
3. In the Physical Volumes to Use window, select the drive configured with RAID-0 for the OS installation.
4. Press the **Add** button and configure a partition for **Swap** using parameters defined in the table provided earlier.




5. Press the **Add** button and configure a partition for **/** using parameters defined in the table provided earlier.

■ Creating Volume Groups and Partitions



6. Verify the new Logical Volumes are configured as required and press the **OK** button to set the configuration.
7. Click **LVM Volume Groups** to create a new Volume Group for the MUR Application and Postgres Database partitions. Create these with the name *VolGroup01* and *VolGroup02* respectively.



8. In the Physical Volume to Use field, select the drive configured with RAID-0 and create a Volume Group with the MUR Application as detailed in the table and press the **Add** button.
9. Use the last physical volume to create the Volume Group for the Postgres Database as detailed in the table and press the **Add** button.
10. Press **OK** to display the updated Volume Groups screen. Check this detail against the correct information in the table if required.
11. Press **Next** to finish the RHEL installation then reboot the system.

Configuring the XFS File System

The LVM configuration you just performed created all the partitions with an EXT file system. An XFS file system can be fine-tuned to work with a RAID array. To configure the XFS file system parameters for stripe-unit and stripe-width you must know the RAID information (stripe size and span of disks) for the respective drives. These values have to be provided in terms of 512 byte blocks.

We suggest you use **su** (stripe-unit) and **sw** (stripe-width) to avoid confusion when setting the following values:

- **su** = RAID controller stripe size in Bytes or KBytes
- **sw** = The span of the physical data disks excluding the number of duplicate disks for RAID10 and parity disks for RAID5 and RAID6.

Follow these steps to configure the MUR application (**/apps**) and Postgres (**/db**) partitions with an XFS file system. In this example, these settings are for four (4) disks of RAID-0 and seven (7) disks of RAID-5.

1. Un-mount the partitions using the **umount** command:

```
$ umount /apps
```

```
$ umount /db
```

2. Create the XFS file system using the **mkfs.xfs** command:

```
$ mkfs.xfs -f -b size=4096 -d su=256k,sw=4 /dev/mapper/VolGroup01-LogVol100
```

```
$ mkfs.xfs -f -b size=4096 -d su=256k,sw=6 /dev/mapper/VolGroup02-LogVol100
```

3. Re-mount the partitions using the **mount** command:

```
$ mount -t xfs /dev/mapper/VolGroup01-LogVol100 /apps
```

```
$ mount -t xfs /dev/mapper/VolGroup02-LogVol100 /db
```

4. Update the **/etc/fstab** file to reflect the updated configuration.

5. Open the **/etc/fstab** file and change the file system type for the **/apps** and **/db** partitions to XFS. The correct entries should be as follows:

```
/dev/VolGroup01/LogVol100 /apps xfs defaults 1 2
```

```
/dev/VolGroup02/LogVol100 /db xfs defaults 1 2
```


Chapter 5

Mobility Unified Reporting System Clustering Support for High Availability

This chapter describes how to install and configure the Veritas cluster for high availability in MUR.

This chapter describes the following topics:

- [System and Hardware Recommendations for HA Deployment](#)
- [Configuring the External Storage Disk on UCS for HA Deployment](#)
- [Tuning the VxFS File System](#)
- [Configuring Resources for High Availability](#)
- [Recovering MUR in HA](#)

System and Hardware Recommendations for HA Deployment

The MUR can be configured in HA mode with asymmetric or active/passive configuration. The following requirements must be met in order to ensure the HA deployment in the MUR system:

- 2 UCS 460 M1 servers for cluster nodes
- Up to 4 Internal disks (depending upon the fault tolerance for boot disk/OS)
- Qlogic QLE2462 4Gb dual port FC HBA / Qlogic QLE2560 8 Gbps FC 1 port
- FC cables
- External Storage with multiple disks (EMC/Sun Storage)
- Veritas (VxFS) file system

The hardware setup for the Veritas Cluster Server (VCS) consists of two cluster nodes connected with an external shared storage. UCS 460 supports multiple PCI slots for attaching the FC HBA cards. Both the cluster nodes must be connected to the external storage. The cluster nodes must be installed with the Cisco MITG RHEL operating system, Veritas Storage Foundation (Veritas Volume Manager and Veritas File System), and VCS (for High Availability).

The Veritas Volume Manager (VxVM) is used to create a single Disk Group (DG) containing multiple disks. A separate group of disks is allocated for postgres data directory to achieve good performance.

The RAID recommendations are the same as that of standalone installation i.e. RAID-0 for MUR application and RAID-5 for database.

Separate disk/LUN from the shared storage is required for I/O fencing. I/O fencing is part of the VCS administration. It is assumed that I/O fencing is already configured on the Veritas Cluster setup before the MUR application is installed for HA.

Configuring the External Storage Disk on UCS for HA Deployment

This section describes the procedure to configure an external storage disk on a UCS server for HA deployment.

It is recommended to create separate disk groups for:

- MUR Application
- Postgres data directory
- MUR Archive (This can be either part of MUR Application disk group OR created separately)

The number of disks in each group depends on the total throughput required and the size of disks in the external storage. Please contact your system administrator for setting up the external storage to make the required number of disks (Logical Disk Units (LUNs)) accessible from both the cluster nodes. When the external storage disks are made accessible, you can see them connected using multipath command:

```
$ multipath -l

mpath6 (36006016069902d008892bc0dec14e111) dm-7 DGC,RAID 5

[size=300G][features=1 queue_if_no_path][hwhandler=1 emc][rw]

\_ round-robin 0 [prio=0][active]

\_ 8:0:0:1 sdi 8:128 [active][undef]

\_ 8:0:1:1 sdk 8:160 [active][undef]

mpath5 (36006016069902d00563e8a20ec14e111) dm-6 DGC,RAID 5

[size=300G][features=1 queue_if_no_path][hwhandler=1 emc][rw]

\_ round-robin 0 [prio=0][active]

\_ 8:0:0:0 sdh 8:112 [active][undef]

\_ 8:0:1:0 sdj 8:144 [active][undef]
```

To configure an external storage disk into separate DGs:

Step 1 Rebuild the disk lists with the new disks detected by the kernel using the following command:

```
$ vxdctl initdmp

$ vxdctl enable
```

To see the status of the new disk, use the following command:

```
$ vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
disk_0	auto:none	-	-	online invalid

■ Configuring the External Storage Disk on UCS for HA Deployment

disk_1	auto:none	-	-	online invalid
disk_2	auto:none	-	-	online invalid
disk_3	auto:none	-	-	online invalid
emc_clariion0_28	auto	-	-	error
emc_clariion0_29	auto	-	-	error

Step 2 To setup the disks, use the following commands:

```
$ /etc/vx/bin/vxdisksetup -i emc_clariion0_28
$ /etc/vx/bin/vxdisksetup -i emc_clariion0_29
```

To see the status of the new disk, use the command:

```
$ vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
disk_0	auto:none	-	-	online invalid
disk_1	auto:none	-	-	online invalid
disk_2	auto:none	-	-	online invalid
disk_3	auto:none	-	-	online invalid
emc_clariion0_28	auto:cdsdisk	-	-	online
emc_clariion0_29	auto:cdsdisk	-	-	online

Step 3 With the newly initialized disks, create two separate DGs for apps (MUR) and db (postgres data directory) partitions using the following command:

```
vxvg init apps_dg apps_dg01=emc_clariion0_28
$ vxvg init db_dg db_dg01=emc_clariion0_29
```

 **Important:** You can specify more disks (using *Step 2*) and add the disk to disk groups.

After adding the disk into respective DGs, you can verify them using the following command:

```
$ vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
disk_0	auto:none	-	-	online invalid
disk_1	auto:none	-	-	online invalid
disk_2	auto:none	-	-	online invalid


```

disk_3          auto:none      -          -          online invalid
emc_clariion0_28 auto:cdsdisk  apps_dg01  apps_dg01  online
emc_clariion0_29 auto:cdsdisk  db_dg0     db_dg       online

```

VxVM will ensure that the newly created DGs are visible from both the cluster nodes. These disk groups can be used only from one node at a time. You will have to import/deport a disk group from either node to use the disk groups and their volumes.

Step 4 Create volumes in the disk groups using the following commands:

```

$ vxassist -g apps_dg make apps_vol 299g
$ vxassist -g db_dg make db_vol 299g

```

Step 5 Initialize the volumes with the VxFS file system using the following commands:

```

$ mkfs -t vxfs -o bsize=4096,largefiles /dev/vx/rdisk/db_dg/db_vol
$ mkfs -t vxfs -o bsize=4096,largefiles /dev/vx/rdisk/apps_dg/apps_vol

```

For better performance, use a 4 KB block size and enable support for large files (more than 1 TB).

Step 6 Create the mount points and mount the volumes using the following commands:

```

$ mount -t vxfs -o largefiles /dev/vx/dsk/apps_dg/apps_vol /shared_apps
$ mount -t vxfs -o largefiles /dev/vx/dsk/db_dg/db_vol /shared_db

```

Tuning the VxFS File System

The VxFS file system can be tuned for better performance. The **vxtunefs** command can be used to set the tuning parameters. The default values of these parameters are set when the volume is mounted.

The performance of the MUR application can improve when the following tuning parameters are changed:

- **read_pref_io**: The preferred read request size. The filesystem uses this in conjunction with the **read_nstream** value to determine how much data to read ahead. The default value is 64000. The MUR performance can improve when this value is set to 128000.
- **read_nstream**: This is the desired number of parallel read requests of the size specified in the **read_pref_io** parameter to have outstanding at one time. The file system uses the value specified in the **read_nstream** parameter multiplied by the value specified in the **read_pref_io** parameter to determine its read ahead size. The default value for the **read_nstream** parameter is 1. If you know the hardware RAID configuration on the external storage, then set the **read_nstream** parameter value to be the number of columns (disks) in the disk array.
- **write_pref_io**: The preferred write request size. The filesystem uses this in conjunction with the value specified in the **write_nstream** parameter to determine how to flush behind on writes. The default value is 64000. The MUR performance can improve when this value is set to 128000.
- **write_nstream**: This is the desired number of parallel write requests of the size specified in the **write_pref_io** parameter to have outstanding at one time. The file system uses the value specified in the **write_nstream** parameter multiplied by the value specified in the **write_pref_io** parameter to determine when to flush behind on writes. The default value for the **write_nstream** parameter is 1. For disk striping configurations, set the value of the **write_pref_io** and **write_nstream** parameters to the same values as the **read_pref_io** and **read_nstream** parameters.

Step 1 Use the following command to tune Veritas file system:

```
$ /opt/VRTS/bin/vxtunefs -o
read_pref_io=131072,read_unit_io=131072,write_pref_io=131072,write_unit_io=131072
/shared_db
```

Step 2 To ensure that these values are not lost after a reboot, create the file */etc/vx/tunefstab* and add entries using the command:

```
$ cat /etc/vx/tunefstab /dev/vx/dsk/apps_dg/apps_vol
read_pref_io=131072,read_unit_io=131072,write_pref_io=131072,write_unit_io=131072
```

Please create these files on both the nodes.

Now, you can install the MUR application from one of the cluster nodes. Please ensure that the check box to start the MUR processes is not selected during the MUR installation. MUR application will be started after the required resources are configured through VCS. For more information on how to install the MUR application, refer to the *Managing Mobility Unified Reporting System Installation* chapter of this guide.

Prior to installing MUR, MUR Administrator user and Group need to be created and activated. MUR Administrator user is the user who will own the MUR installation. Execute the following steps as **root** user for creating and activating MUR user and group:

1. Create MUR group using the following command:

```
groupadd -g <groupID> <groupname>
```

For example: **groupadd -g 500 murgroup**

2. Create MUR user using the following command:

```
useradd -u <userID> -c "MUR-administrator" -g <groupname> -m <username>
```

For example: `useradd -u 100014 -c "MUR-administrator" -g murgroup -m muradmin`

3. Activate the MUR user using a password:

```
passwd <username>
```

For example: `passwd muradmin`

During installation, provide this user as the Administrator user input.




Important: For HA mode, the user and group creation must be done on both the nodes in the cluster (Active and Passive). Ensure that you provide the same userID, groupID and password on both the nodes. Also, note that HA mode installation is only supported on RHEL platform.

Configuring Resources for High Availability

After installation of MUR application, the following resources need to be configured with the Veritas cluster:

- NIC — To monitor a Network Interface Card (NIC)
- IP — To monitor an IP address
- Disk Group, Volume, and Mount — for shared storage
- MUR Application — comprising of all the MUR related processes



Assume the following configuration resources as:

- Two disk groups — apps_dg and db_dg
- Two volumes — apps_vol and db_vol mounted on */shared_apps* and */shared_db* respectively
- MUR installation directory — */shared_apps/starbi*
- MUR postgres data directory — */shared_db/data*
- MUR Administrator — muradmin
- Shared/Floating IP address — 10.4.83.151 (on NIC eth0). The floating/shared IP used should be a public IP in DNS, pingable from the client machine.

To configure these resources:

 **Important:** The following configurations should be performed only on the node where the MUR application is installed.

Step 1 Log on as **super user** (root).

Step 2 Make the Veritas config file (main.cf) writable using the following command:

```
$ haconf -makerw
```

Step 3 Create resource group using the following commands:

```
$ hagrps -add mur-ha
$ hagrps -modify mur-ha SystemList <Node1> 0 <Node2> 1
$ hagrps -modify mur-ha NumRetries 1
```

Where, Node1 and Node2 are the hostnames of the active and passive nodes.

Step 4 Create DG resource for the MUR partition using the following commands:

```
$ hares -add mur-apps-dg DiskGroup mur-ha
$ hares -modify mur-apps-dg DiskGroup apps_dg
$ hares -modify mur-apps-dg Enabled 1
```

Step 5 Create DG resource for postgres data partition using the following commands:

```
$ hares -add mur-db-dg DiskGroup mur-ha
$ hares -modify mur-db-dg DiskGroup db_dg
$ hares -modify mur-db-dg Enabled 1
```

Step 6 Create Volume resource for the MUR partition using the following commands:

```
$ hares -add mur-apps-vol Volume mur-ha
$ hares -modify mur-apps-vol DiskGroup apps_dg
$ hares -modify mur-apps-vol Volume apps_vol
$ hares -modify mur-apps-vol Enabled 1
```

Step 7 Create Volume resource for the postgres data partition using the following commands:

```
$ hares -add mur-db-vol Volume mur-ha
$ hares -modify mur-db-vol DiskGroup db_dg
$ hares -modify mur-db-vol Volume db_vol
$ hares -modify mur-apps-vol Enabled 1
```

Step 8 Create Mount resource for the MUR partition using the following commands:

```
$ hares -add mur-apps-mnt Mount mur-ha
$ hares -modify mur-apps-mnt MountPoint /shared_apps
$ hares -modify mur-apps-mnt BlockDevice /dev/vx/dsk/apps_dg/apps_vol
$ hares -modify mur-apps-mnt FSType vxfs
$ hares -modify mur-apps-mnt FsckOpt %y
```

```
$ hares -modify mur-apps-mnt MountOpt largefiles
$ hares -modify mur-apps-mnt Enabled 1
```

Step 9 Create Mount resource for the postgres data partition using the following commands:

```
$ hares -add mur-db-mnt Mount mur-ha
$ hares -modify mur-db-mnt MountPoint /shared_db
$ hares -modify mur-db-mnt BlockDevice /dev/vx/dsk/db_dg/db_vol
$ hares -modify mur-db-mnt FSType vxfs
$ hares -modify mur-db-mnt FsckOpt %-y
$ hares -modify mur-db-mnt MountOpt largefiles
$ hares -modify mur-db-mnt Enabled 1
```

Step 10 Create Application resource for the MUR processes using the following commands:

```
$ hares -add mur-app Application mur-ha
$ hares -modify mur-app User muradmin
$ hares -modify mur-app StartProgram "/shared_apps/starbi/starbi/bin/serv start"
$ hares -modify mur-app StopProgram "/shared_apps/starbi/starbi/bin/serv
forcestop"
$ hares -modify mur-app PidFiles
"/shared_apps/starbi/starbi/server/sysmon/psmon.pid
$ hares -modify mur-app Enabled 1
```

Step 11 Create the NIC resource using the following commands:

```
$ hares -add mur-nic NIC mur-ha
$ hares -modify mur-nic Device eth0
$ hares -modify mur-nic Enabled 1
```

Step 12 Create the IP resource using the following commands:

```
$ hares -add mur-ip IP mur-ha
$ hares -modify mur-ip Device eth0
$ hares -modify mur-ip Address <ip-address>
```



Important: The floating or shared IP address should be a public IP in the DNS to which the client machine can successfully ping.

```
$ hares -modify mur-ip NetMask 255.255.255.0

$ hares -modify mur-ip Enabled 1
```

Step 13 Set the resource dependencies using the following commands:

```
$ hares -link mur-app mur-apps-mnt

$ hares -link mur-app mur-db-mnt

$ hares -link mur-apps-mnt mur-apps-vol

$ hares -link mur-db-mnt mur-db-vol

$ hares -link mur-apps-vol mur-apps-dg

$ hares -link mur-db-vol mur-db-dg

$ hares -link mur-app mur-ip

$ hares -link mur-ip mur-nic
```

Step 14 Dump the configuration to the Veritas config file using the following command:

```
$ haconf -dump -makero
```

Step 15 Bring the MUR HA application online on Node 1 using the following command:

```
$ hagrps -online mur-ha -sys <Node1>
```

Once the above steps are performed, the MUR HA application will start running and you can access GUI using shared IP specified earlier.

Recovering MUR in HA

For recovering MUR in HA mode:

Step 1 Make the Veritas config file writable using the following command:

```
$ haconf -makerw
```

Step 2 Disable MUR HA application on the standby node using the following command:

```
$ hagrps -disable mur-ha -sys <standby-node>
```

Step 3 Make MUR Application (mur-app) resource offline on the active node using the following command:

```
$ hares -offline mur-app -sys <active-node>
```

Step 4 Disable MUR-APP resource using the following command:

```
$ hares -modify mur-app Enabled 0
```

Step 5 Perform usual MUR recovery process on the active-node (Install same MUR version, start postgres and run recovery script).

Step 6 Once recovery is successful, enable MUR-APP resource using the following command:

```
$ hares -modify mur-app Enabled 1
```

Step 7 Bring the MUR-APP resource online on the active node using the following command:

```
$ hares -online mur-app -sys <active-node>
```

Step 8 Enable MUR HA application on the standby node using the following command:

```
$ hagrps -enable mur-ha -sys <standby-node>
```

Step 9 Dump the configuration to Veritas config file using the following command:

```
$ haconf -dump -makero
```


Chapter 6

Managing Mobility Unified Reporting System Installation

This chapter describes how to install, upgrade, and uninstall the MUR application.

This chapter describes the following topics:

- [Installing MUR](#)
- [Upgrading MUR](#)
- [Uninstalling MUR](#)



Important: The procedures for installation, upgrade, and uninstallation of MUR and RDP remain the same.



Important: Please note that the terminologies “starbi”, “inPilot” and “mur” used throughout this guide mean the same.

Installing MUR

This section provides instructions on how to install the MUR application.



Important: Make sure that your system meets the minimum requirements as indicated in the *MUR System Requirements* section in the *MUR Overview* chapter of this guide.

The following MUR components are installed by MUR installer.

- For Solaris platform
 - Apache v2.2.25 with mod_python v3.3.1
 - Python v2.6.4
 - Postgres v 8.2.0
 - Django v1.0.2
 - JRE v1.6.0_12
 - Quartz Scheduler v1.6.4
- For RHEL platform
 - Apache v2.2.25 with mod_python v3.3.1
 - Python v2.6.1
 - Postgres v 8.3.4
 - Django v1.0.2
 - JRE v1.5.0_11
 - Quartz Scheduler v1.6.4



Important: In RHEL-based deployment of MUR, L-ESS is NOT required as the ECS module can be configured to push the xDRs directly to the MUR reporting server. Push from ASR chassis is the Cisco recommended deployment model. Currently L-ESS is supported only on Solaris platforms. For information on the L-ESS installation instructions, refer to the *ESS Installation and Administration Guide*. Existing deployments where L-ESS is installed, to pull EDRs from the chassis, may continue with their deployment model in the 12.0 version of MUR Software Release and later.



Important: It is recommended that you first install the master MUR before proceeding with the RDP installation.

Prior to installing MUR, MUR Administrator user and Group need to be created and activated. MUR Administrator user is the user who will own the MUR installation. Execute the following steps as **root** user for creating and activating MUR user and group:

1. Create MUR group using the following command:

```
groupadd -g <groupID> <groupname>
```

For example: `groupadd -g 500 murgroup`

2. Create MUR user using the following command:

```
useradd -u <userID> -c "MUR-administrator" -g <groupname> -m <username>
```

For example: `useradd -u 100014 -c "MUR-administrator" -g murgroup -m muradmin`

3. Activate the MUR user using a password:

```
passwd <username>
```

For example: `passwd muradmin`

During installation, provide this user as the Administrator user input.



Important: For High-Availability mode, the user and group creation must be done on both the nodes in the cluster (Active and Passive). Ensure that you provide the same userID, groupID and password on both the nodes. Also, note that High-Availability mode installation is only supported on RHEL platform.

Setting the Database Environment Strings

Prior to installing the MUR components onto the server hardware, there are numerous system environment configuration settings that should be configured. While PostgreSQL will be installed during the installation procedure, these settings must be configured manually.



WARNING: Failure to configure these settings may cause data loss and will minimally cause errors in the operation.

Settings for Solaris

Add the following values to system file in the `/etc/system` directory if they are not present and restart the system before continuing with the installation of MUR components.

```
set msgsys:msginfo_msgmb=65536
set msgsys:msginfo_msgtql=1024
set shmsys:shminfo_shmmax=10737418240
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmni=256
set shmsys:shminfo_shmseg=256
set semsys:seminfo_semmap=256
set semsys:seminfo_semmni=512
set semsys:seminfo_semmns=512
set semsys:seminfo_semmsl=270
```

Settings for RHEL

Add the following values to system file in the `/etc/sysctl.conf` if they are not present and restart the system before continuing with the installation of MUR components.

```
kernel.shmmax=10737418240
```

```
kernel.shmall=4294967296
```

Pre-installation Checks

Ensure the following checks are made before installing the MUR application.

 **Important:** Please note that L-ESS is required ONLY for a Solaris-based deployment of MUR. In the case of RHEL-based deployment of MUR, the ECS module is configured to push the xDRs directly from the chassis to the MUR reporting server via SFTP.

Step 1 The recommended filesystem for installation is ZFS. If Solaris-based installation is performed on any other filesystem, a warning message appears indicating the recommended filesystem.

 **Important:** Please note that the ZFS related recommendations mentioned throughout this guide are specific to SOLARIS ONLY and NOT for RHEL.

Step 2 MUR must be installed as a **root** user on the system. Installation with other user privileges is not recommended.

Step 3 Make sure no other Apache web server is running on the port being used for installation (default port is 8080). If it is, stop it before proceeding with the installation or provide a different port for Apache server. Check if an application is running on a given port by entering the following command:

```
netstat -an | grep <port number>
```

Step 4 Make sure no other Postgres server is running on the port being used for installation (default port is 5432). If it is, stop it before proceeding with the installation or provide a different port for Postgres server. Check if an application is running on a given port by entering the following command:

```
netstat -an | grep <port number>
```

Step 5 Make sure no other application/process is running on the port being used for pgBouncer (default 'Postgres port + 1'). If it is, stop it before proceeding with the installation or provide a different port for Postgres server so that the installer finds two consecutive ports free (one for Postgres and the other one for pgBouncer). Check if an application is running on a given port by entering the following command:

```
netstat -an | grep <port number>
```

Step 6 Make sure no other server is running on the port being used for installation for XML-RPC (default port is 9999). If it is, stop it before proceeding with the installation or provide a different port for XML-RPC server. Check if an application is running on a given port by entering the following command:

```
netstat -an | grep <port number>
```

- Step 7** MUR installation will ask for the Administrator login and Administrator Primary Group. Administrator login is the OS level administrator of MUR who will own the MUR installation. Administrator Primary Group is the user group of MUR to allow the interaction with external entities like L-ESS.
- Step 8** If the Administrator login provided during MUR installation/upgrade already exists, ensure that it is not an already logged in user.
- Step 9** L-ESS must be stopped before starting MUR installation / upgrade.
- Step 10** If the L-ESS is installed as a **root** user, the ownership of L-ESS installation should be changed from **root** to **non-root** user. This new user must be added to MUR Group. For example, if L-ESS is initially running as **root** and new user created is *essadmin*, then perform the following sequence of operations.
- Step a** Stop L-ESS.
 - Step b** Add the user *essadmin* to MUR group by entering the following command as **root** user - **usermod -G <MUR Group> essadmin**
 - Step c** Verify whether the user is added correctly to MUR group using the command **groups essadmin**
 - Step d** Change the ownership of L-ESS installation to this new user using the following command - **chown -R essadmin <LESS installation directory>**
 - Step e** Login as *essadmin* with the command **su essadmin**
 - Step f** Start L-ESS again.
- Step 11** If the L-ESS is installed as a **non-root** user say *essadmin*, this user should be added to MUR Group.
- Step a** Stop L-ESS
 - Step b** Add the user *essadmin* to MUR group by running the following command as **root** - **usermod -G <MUR Group> essadmin**
 - Step c** Log off and relogin again as *essadmin* for the group addition to come into effect.
 - Step d** Start the L-ESS application to continue pulling the EDR files from chassis and forwarding it to MUR.
- Step 12** Perform the following steps only if the user wants to push EDR/UDR files from gateway to MUR server using SFTP mechanism. Otherwise, skip this step.
- Step a** Change to the */etc/ssh* directory.
 - Step b** Open *sshd_config* file from the directory using *vi* editor (or any other editor) and observe the default values for the following variables:
- PasswordAuthentication**
- PAMAuthenticationViaKBDInt** (Applicable ONLY for SOLARIS)
- UsePAM** (Applicable ONLY for RHEL)
- Step c** Change the default values for the following variables as indicated here.
- PasswordAuthentication** = *yes*
- PAMAuthenticationViaKBDInt** = *no* (Applicable ONLY for SOLARIS)

UsePAM = *no* (Applicable ONLY for RHEL)

Step d After updating restart SSH daemon using the following command:

In the case of SOLARIS:

```
svcadm restart ssh
```

 **Important:** Please note that the above command can be executed only in Solaris 10 environment.

In the case of RHEL:

```
service sshd restart
```

Step 13 The recommended user/group settings for MUR are:

- NIS-USER<->NIS-GROUP
- NON-NIS-USER<->NON-NIS-GROUP

The NIS users should always be associated with NIS Groups. The non NIS users should be associated with Non NIS groups. Also, it is recommended to have separate non NIS users for MUR installation.

Step 14 For High Availability (HA) mode, please follow the steps described in the section *Configuring External Storage Disk on UCS for HA deployment* in the *Mobility Unified Reporting System Clustering Support for High Availability* chapter of this guide to initialize the external storage and create the Disk Groups, Volumes and Mounts required for installing MUR application. MUR application installation can be run from any of the cluster nodes.

MUR Installation

The MUR installation files are distributed as a single compressed file.

 **Important:** In the MUR Software Releases prior to 11.0.100 build, this installation file is distributed with a **.tar.gz** extension. In the MUR Software Release 11.0.100 and later, this file is distributed in zip format.

 **Important:** The MUR application currently supports UCS Linux platform and Solaris-Sparc/Solaris-x86 platform. The installable tar file names help in identifying the platform. For example, `mur.x.x.xx_rhel_x86.zip` indicates that this file is for RHEL platform. Similarly, `mur.x.x.xx_solaris_sparc.zip` indicates that this file is for Solaris-Sparc platform.

For information on downloading the appropriate MUR package for your requirements, contact your Cisco account representative.

The MUR application and its components can be installed using one of the following two methods.

- [Installing MUR Using GUIConsole based Installer](#)

 **Important:** Please note that the terminologies “starbi”, “inPilot” and “mur” used throughout this guide mean the same.

Installing MUR Using GUI/Console based Installer



Important: To perform the installation procedure explained in this section, you must be logged into the server as a **root** user.



Important: Fresh installation for backup recovery purpose should be installed on the same path where last backup is stored and also should have the same IP address and port configuration if the MUR is deployed in distributed mode. Make sure that the existing older installation is either removed or moved to a different directory because the metadata recovered from previously installed MUR will have all references as per older installation e.g. archive path, SFTP details, etc.



Important: In the MUR Software Releases prior to 11.0.100 build, this installation file is distributed with a **.tar.gz** extension. In the MUR Software Release 11.0.100 and later, this file is distributed in zip format.

Follow the instructions below to install MUR using the GUI/Console based installation wizard.

Step 1 Change to the directory in which the file is stored.

Step 2 Unzip the file by entering the following command:

```
unzip mur.x.x.xx_os_arch.zip
```

x.x.xx is the version of the MUR installation file.

os indicates the Operating System on which the MUR application is running. It can be either RHEL or Solaris.

arch indicates the architecture either Sparc or x86.



Important: To unzip the **.gz** package file, use **tar -xvf <file_name>** command.

Decompressing the installation file results in the following files:

- *inst*: A GUI/Console based installer to install the MUR application.
- *setup.bin*: The executable used by *inst* to install MUR application.

Step 3 Execute the script by entering the following command:

```
./inst [MODE]
```

where [MODE] is optional.

Two installation modes are supported namely:

- gui
- console

The command '*inst/uninst -help*' provides usage of the scripts. This script installs the Apache, Postgres and Scheduling servers functionality. The display must be set for running in GUI mode, else the installation will run in Console mode.

The MUR Installer dialog appears displaying the MUR version getting installed.

Step 4 Click **Next** to proceed.

■ Installing MUR

Step 5 Respond to the on-screen prompts with the help of inputs given in the following table and configure various parameters as required.

Parameter	Description	Default Value
PostgreSQL System Settings		
	This dialog asks the user to check the variable values in system file. If one or more entries are missing, click Cancel to update the system file and restart the system to re-run installer. For more information, refer to the Setting the Database Environment Strings section.	N/A
MUR Clustering Support		
Veritas Cluster Failover	If Veritas clustering environment is found, user will be asked whether to install MUR application with clustering support (High Availability mode) or not.	Yes
Floating IP	The Floating IP to be used in Veritas clustering environment.	N/A
MUR Installation Directory		
Enter Installation Path	Enter the base directory path where MUR is to be installed. Click Browse to change the installation path.  Important: For HA mode installation, the path must be on shared disk.	<current_directory>
A Component Type screen appears showing the components for installation. This screen allows you to select either Master MUR or RDP for installation.  Important: Make sure that you first install the master MUR and then proceed with the RDP installation.		
MUR Administrator and Group Configuration		
Administrator Login	Enter an administrator name for the Operating System (OS) level administrator of MUR.  Important: Please note that you should not login as a root user.  Important: The Administrator user created should be manually activated with a password once the MUR installation is complete. This can be done by entering the following command as root user: passwd <adminusername> Upon executing this command, the user will be asked to enter a suitable administrator password.	muradmin

Parameter	Description	Default Value
Administrator User ID	Type the Administrator User ID for the MUR Administrator login.  Important: This input will be asked only if the Administrator login name provided does not exist.	100014
Administrator Primary Group	Type the Primary Group name for the Administrator.  Important: If the Administrator login name provided already exists, the Primary Group of this login will be considered as the MUR User Group . Otherwise, the user will be asked to enter the Primary Group information.	murgroup
PostgreSQL Server Configuration		
Postgres Login	This is a read-only parameter. The Postgres login name will be the same as the Administrator login name provided earlier.	muradmin
Postgres password	Enter the password for the Postgres database administration.	N/A
Postgres Port	Enter the port number on which PostgreSQL communication will be running.  Important: Ensure that no other application/process is running on configured port as well as on next consecutive port (as this port will be used for pgBouncer).	5432
Enter data directory path	Enter the data directory path of postgres being used. Click Browse to change the installation path.  Important: For HA mode installation, the path must be on shared disk.	<MUR_install_dir>/starbi/postgres/data
MUR Port Configuration		

Parameter	Description	Default Value
Apache Protocol Type	Select the Apache protocol type from the available options: • HTTP • HTTPS • Both HTTP and HTTPS If both HTTP and HTTPS are required, select the check boxes Apache HTTP Support and Apache HTTPS Support. For RDP to MASTER synchronization to take place over RPC, either HTTP or HTTPS will be used. This depends on the user's selection in this screen.	HTTP
Apache HTTP Port Number	Type the port number over which Apache web server communication will occur with MUR. Apache HTTP port should be greater than 1025 and lesser than 65535.  Important: This parameter will be enabled only when HTTP option or Both HTTP and HTTPS option is selected.  Important: Ensure that no other Apache web server is running on the port being used for installation. If the port is being used, abort the installation. For RHEL: For RHEL, Apache port provided should be > 1024. RHEL does not allow port 80 to be used by non-root users. However, Apache Web server requests made on port 80 can be redirected to a port >1024 defined by the operator, with the following two commands. For example, to redirect requests made on port 80 to port 8080: <pre>iptables -t nat -A PREROUTING -p tcp --dport 80 -i eth0 -j REDIRECT --to-port 8080 iptables -t nat -A OUTPUT -p tcp -d 127.0.0.1 --dport 80 -j REDIRECT --to-port 8080</pre> For Solaris: For using the Apache port < 1024, run the following command as root user once the installation is complete, and restart the Apache server. <pre>usermod -K defaultpriv=basic,net_privaddr <MUR admin user></pre> For example: <pre>usermod -K defaultpriv=basic,net_privaddr muradmin</pre>  Important: This poses a major security concern as it will allow <i>muradmin</i> to use all standard ports < 1024.	8080

Parameter	Description	Default Value
Apache HTTPS Port Number	Type the port number over which Apache web server communication will occur with MUR. Apache HTTPS port should be greater than 1025 and lesser than 65535.  Important: This parameter will be enabled only when HTTPS option or Both HTTP and HTTPS option is selected.	8443
SSL Certificate Type	Select the SSL certificate type from the available options: • Cisco • Custom If you select Custom as the option, click Browse to select the path where the SSL certificate file is present. The SSL certificate will be copied to <MUR_install_dir>/starbi directory. The installer will allow only .crt file for certificate selection.  Important: This parameter will be enabled only when HTTPS option or Both HTTP and HTTPS option is selected.	Cisco
SSL Key	Click Browse to select the path where the SSL key is present. The SSL key will be copied to <MUR_install_dir>/starbi directory. The installer will allow only .key file for key selection.  Important: This parameter will be enabled only when Custom is selected as the option for the SSL Certificate Type .	N/A
Available Port Range for MUR Components (200 Ports)	Type the port number in the start port text area. Note that the end port text area is a READ-ONLY field. End port number will be populated automatically based on the start port number. For example, if the start port number entered is 9001 the end port number will be 9200 (startPort + 199). Start port range should be between 1026 and 65535 so that the end port has a range of 1226 to 65535. If any of the port/ports in the specified range is/are not available, then the installer throws error and prompts the user to enter a new start port number.	9001 - 9200
MUR Archive Directory Configuration		

■ Installing MUR

Parameter	Description	Default Value
Enter archive directory path	Enter the directory path for archiving parsed files. Click Browse to change the installation path.  Important: For HA mode installation, the path must be on shared disk.	<MUR_install_dir>/archive
Pre-installation Summary screen		
The pre-installation screen displays the product name, install location, other product configurations, and disk space information before installing the product. Click Cancel to stop installation or Install to continue installation.		
Installing MUR		
The screen shows all the contents being loaded on the machine during installation. Click Cancel to stop installation.		
MUR Server Startup		
Start All Servers After Installation	Select the option to start all servers after installation. Click Next to proceed.  Important: For HA mode, please ensure that you do not start the MUR processes during the MUR installation. MUR application will be started after the required resources are configured through Veritas Cluster Server (VCS). For information on configuring resources, refer to the <i>Configuring Resources for High Availability</i> section in the <i>Mobility Unified Reporting System Clustering Support for High Availability</i> chapter of this guide.	Yes
Install Complete		
	The screen shows whether installation is successful or failed. Click Done to quit the installer.	N/A

When the MUR installation is complete, see the *Cisco Mobility Unified Reporting System Online Help* documentation for information on how to access and use the GUI.

Once the MUR installation is complete, the following configurations are recommended to be performed through the GUI for proper functioning of MUR.

- Configure EDR file name format by selecting **Admin>Gateways**, according to the EDR format received from chassis.
- Enable backup through **Admin>Backup**.
- Enable both Database Purging and Archived Files Purging through **Admin>Purging**.
- Enable SNMP alarms through **System>SNMP Configurations**, and update the SNMP-Manager IP-Address accordingly.

- Add email configurations through **System>Email servers**, and users through **Admin>Users** to receive notifications.

For HA mode, please ensure that you configure the required resources through VCS after the installation of MUR application. For information on configuring resources, refer to the *Configuring Resources for High Availability* section in the *Mobility Unified Reporting System Clustering Support for High Availability* chapter of this guide.

Confirming Successful Installation

Verify that the MUR application is running and accessible by entering the following URL in your Web browser:

```
http://<MUR_installation server name or IP address>:<apache port>
```

– or –

```
https://<MUR_installation server name or IP address>:<apache port>
```

For information on logon details, refer to the *Launching the MUR GUI* section in the *Mobility Unified Reporting System Administration and Management* chapter of this guide.

For information on using the MUR GUI, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

For HA mode installation, verify that the MUR application is running and accessible by entering the following URL in your Web browser:

```
http://<MUR floatingIP address>:<apache port>
```

– or –

```
https://<MUR floatingIP address>:<apache port>
```

Upgrading MUR

This section provides instructions on how to upgrade the installed MUR application.



Caution: Please contact your Cisco account representative to ensure compatibility prior to upgrading.



Important: In RHEL-based deployments, L-ESS is NOT required as the Enhanced Charging Services (ECS) module can be configured to push the xDRs directly to the MUR reporting server. Push from ASR chassis is the Cisco recommended deployment model. Currently L-ESS is supported only on Solaris platforms. For information on the L-ESS installation instructions, refer to the *ESS Installation and Administration Guide*. Existing deployments where L-ESS is installed, to pull EDRs from the chassis, may continue with their deployment model in the 12.0 version of MUR Software Release and later.

The upgrade procedure ensures that the database content is retained in the new installation. It also ensures that if there are any pending files to be processed in the old installation, then those files are also made available in the new installation.



Important: If MUR is being upgraded from a version in which backup and purging features are not available, to a version in which backup and purging features are supported, then it is recommended that you enable backup feature and take one complete successful snapshot of backup before enabling purging feature. If the backup feature is disabled then enabling purging will cause removal of data without waiting for it to be backed up. If the backup is being taken for the first time after upgrade, then it may take considerable time for first backup.



Important: Upgrade from existing non HA installations to HA installation is not supported.



Important: If the previous installation is MUR then the installation script will cause upgrading the software to MUR and if the previously installed component is RDP then the script will cause upgrading to RDP.



Important: Before performing the upgrade process, ensure that the browser cache is cleared.



Important: To perform the upgrade procedure explained in this section, you must be logged into the server as a **root** user.



Important: If the Apache configuration has been modified after the upgrade process, the gateway cannot be added/modified. Hence it is important to maintain the same configuration for the gateway to be added/updated successfully.

When upgrading to MUR software version 12.2 and above, file parsing configurations are NOT synced automatically at all RDPs. As a result, EDR file parsing does not happen at RDPs.

To overcome this, perform the following steps in the MUR GUI:

1. After the upgrade, manually attach appropriate RDPs to all relevant gateways.
 - Create appropriate RDP regions through the **System** menu.

- Attach all RDPs to their respective regions through **Edit RDP** page viewed by clicking **RDP** from **Admin** tab.
 - Attach appropriate gateways to their corresponding RDPs and regions through **Edit Gateway** page viewed by clicking **GATEWAYS** from **Admin** tab.
2. Navigate to **System > File-Parsing Configs** menu and then manually save the file parsing configurations for all the gateways which are attached to RDPs.
- The MUR upgrade process is carried out in two steps:

1. Online Upgrade
2. Offline Upgrade

The online upgrade is the conventional upgrade process. It will upgrade only last 7 days of available data i.e. it will get the latest date for which data is available and upgrade the last 7 days data only from that date.

Once the online upgrade is complete, offline upgrade starts in the background and it will upgrade all the remaining data older than last 7 days.

During the offline upgrade, there is a possibility of data outage. So, the reports older than last 7 days might be inaccessible from GUI during this period. Once the offline upgrade is over, these reports will be visible again.

Please note the following key points:

- Once MUR is upgraded and if any schemas support additional counter then you should reconfigure schema for that gateway.
- After upgrade is over, the previous data and the schemas displayed earlier for the gateway will be shown on the GUI.
- If you want to perform schema configuration for the gateways which were added prior to upgradation, then you should configure the schemas through the GUI by accessing Bulkstat Schema Configuration screen and disable the earlier file format used on the gateway.

 **Important:** After the MUR upgrade, it is highly possible that the Schedule and Email feature for DPI and HTTP might not work properly. To solve this issue, you should clear the cookies and recreate the favorites for these reports.

The following steps describe how to upgrade the MUR application:

Step 1 Stop the L-ESS by running the following command from the `<LESS_install_dir>/ess` directory:

```
./serv stop
```

Step 2 Stop the MUR application using the following command from the `<MUR_install_dir>/starbi/bin` directory:

```
./serv stop
```

 **Important:** For all MUR software versions 9.0.16 and later, use the `serv stop` command.

or

```
./shutdown.sh
```

 **Important:** For all MUR software versions 9.0.15 and lower, use the `shutdown` command.

Then, check the status of processes using the following command:

```
./serv status
```



Important: For all MUR software versions 9.0.16 and later, use the **serv status** command.

or

```
./status.sh
```



Important: For all MUR software versions 9.0.15 and lower, use the **status** command.



Important: Make sure that none of the processes is running.

Step 3 Install the new release of MUR.

MUR is upgradable from:

- Earlier script installer based version to newer script installer based version
- Earlier script installer based version to GUI/Console installer based version
- Earlier GUI/Console installer based version to subsequent GUI/Console installer based version

For instructions on different MUR installers, refer to the [MUR Installation](#) section.

In case of the first two upgrade options mentioned above, make sure that you enter the old installation path (<install_dir>) for upgrade when prompted for the 'MUR Installation directory'. In case of the third upgrade option, it automatically detects the old installation path through registry information. The installation automatically detects earlier setup and reads required configuration for Apache, Postgres and RPC port, etc. You will be prompted with a confirmation message before proceeding with the upgrade process.

After upgrade, the log files are generated at `/starbi/logs/` directory.



Important: The installation script will check if the Administrator user and Primary Group information is already present in database. If it does not exist, it will ask the user to enter this information and then continue with the upgrade.

Step 4 After the installation is done, start all the MUR related processes using the following command from the <MUR_install_dir>/star/bin directory:

```
./serv start
```

Then, start the L-ESS using the following command from the <LESS_install_dir>/ess directory:

```
./serv start
```

Step 5 Modify the L-ESS configuration or HDD configuration to reflect the changes in the MUR installation path.

Step 6 Restart the EDR file generation or HDD file push as needed.



Important: The RDP should be upgraded manually. If the version of the RDP is not compatible with the MUR, then MUR may ignore the data sent by RDP. Thus, RDP should always be upgraded if it is not in sync with the MUR. For change in mode from RDP to MUR or vice-versa, re-installation is required.

Uninstalling MUR

This section provides instructions on how to uninstall the MUR application.

 **Important:** It is recommended that you manually perform a backup of all critical and historical data files before proceeding with this procedure. Failure to do this causes removal of all the directories, files and database. However, in the case of scalability setup, the shared directory will remain intact and can be removed manually using the command `rm -rf <gateway-directory-path>` only when all the RDPs configured for a gateway are uninstalled.

The MUR application and its components can be uninstalled using one of the following two methods:

- [Uninstallation Using GUIConsolebased Uninstaller](#)

 **Important:** The Administrator user and Primary Group configured during installation / upgrade will not be deleted during uninstallation. These have to be deleted manually by entering the following commands as **root** user: `userdel <ADMINUSER>` and `groupdel <ADMINGROUP>`

 **Important:** Before deleting the Group, ensure that NO other users are attached to the same group by entering the following command: `logins -g <ADMINGROUP>`

 **Important:** During MUR uninstallation, users will be prompted to verify if any of the gateways are mapped and also to take the backup of unprocessed EDR files in the input directory path. If the user does not take appropriate actions to the prompt, then the EDR files will be deleted permanently.

Uninstallation Using GUI/Console-based Uninstaller

This method must be used if installation has been done using GUI/Console based installer (using **inst**).

 **Important:** To perform the uninstallation procedure explained in this section, you must be logged into the server as a **root** user.

Step 1 Change to the `<MUR_install_dir>/starbi` directory and enter the following command:

```
./uninst [MODE]
```

where [MODE] is optional.

Two modes are supported namely:

- gui
- console

The display must be set for running in GUI mode, else the uninstallation will run in Console mode.

The MUR Uninstaller dialog appears.

Step 2 Click **Uninstall** to proceed.

This uninstall script stops all the servers if it is running and all the data is wiped off.



Important: The uninstall script does not cleanup the archive directory. You must take a backup of archive directory and manually delete the user `<muradmin>` and group `<murgroup>` using the commands `<userdel muradmin>` and `<groupdel murgroup>`.

Chapter 7

Mobility Unified Reporting System Administration and Management

This chapter provides information on administering and managing the MUR application.

This chapter describes the following topics:

- [Launching the MUR GUI](#)
- [Administration](#)
- [Operations and Management](#)



Important: Please note that the terminologies “starbi”, “inPilot” and “mur” used throughout this guide mean the same.

Launching the MUR GUI

It is recommended to use either Internet Explorer (v 7.0+) or Mozilla Firefox (v 3.0.10+) browser for launching the MUR interface.

Note that:

- No additional plug-in is required.
- The javascript is enabled by default on the intended browser.
- Suggested screen resolution is 1024 x 768 and above.

To launch the MUR interface:

1. In a Web browser, enter the following URL:

```
http://<MUR-server-hostname or IP address>:<apache port>
```

– or –

```
https://<MUR-server-hostname or IP address>:<apache port>
```

Where, the MUR server host name or the IP address is the input that you entered while installing the MUR software application. For example, `http://10.4.5.2:8080`

In Release 14.0, MUR is SSL enabled by enabling the HTTPS apache server. If, during the installation, the user has selected HTTPS support then the MUR application can be accessed through the https URL.

2. Enter your user name and password, and then click **Log In**. The user name must be an alpha and/or numeric string of 3 through 16 characters in length. The only special character that a user name can include is underscore (_).



Important: At first log on, the users are expected to enter *admin* as the input for the **Username** and **Password** fields.

The password must meet the following criteria:

- Must be a minimum of 8 characters long and a maximum of 32 characters long
- Must not be a repeat or reverse of the associated user name
- Must not be more than 3 of the same characters used consecutively
- Must contain at least 3 of the following combinations:
 - English upper case characters (A through Z)
 - English lower case characters (a through z)
 - Numerical (0 through 9)
 - Special characters (such as _, ., !, @, \$, *, =, -, ?, etc)

The only account created after the initial set-up is *admin / admin* and it has Administrator privileges.

Once logged-in, the user's Dashboard will be displayed with reports if already configured (the displayed reports are specific to each user account).



Important: At first log on, the users will see an empty Dashboard. The necessary data should be populated and required parameters should be configured for report generation.

The user name is always displayed on the right-up corner of the page until the user logs out of the application.

Administration

This section provides information on how to administer and manage the MUR application.

Managing User Accounts

The MUR application provides two levels of access privileges:

- Administrator: Users in this group have the following privileges:
 - Create, edit, and delete other user accounts
 - Edit configuration settings
 - Activate, deactivate, and reset password for operator users
 - Generate and view reports
- Operator: Users in this group can:
 - Generate reports
 - View module-level reports available to them



Important: Only Administrator user can create other administrator and operator user accounts.

User authentication can be interfaced with enterprise authentication mechanism i.e. LDAP Active Directory.

Please note the following limitations with respect to user permissions and privileges:

- All MUR administrators have access to **USERS** and **GROUPS** menu in the **Admin** tab available on the MUR GUI.
- Administrator will have the rights to modify and Administrator with admin user name will have the rights to delete all the MUR users' accounts. Only Administrator can modify its own password. Only admin user will be able to delete any administrator or operator user accounts.
- Only admin user will have rights to delete the MUR users except admin user and administrator user can modify user accounts including their passwords.
- After modifying user role from Administrator to Operator and vice-versa, the user should alter the configuration on the GUI to lock/unlock the user account accordingly.

For more details, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Managing Gateways

The MUR application supports configuring multiple gateways for which reports can be customized and generated. Gateways are the chassis from which EDR and bulkstat files are fetched to the reporting server.



Important: Users with administrative privilege can only add and manage gateways.

When a gateway is added through the GUI, a directory by the name of the gateway is created in the `<MUR_install_dir>/starbi/server/data` directory.

The gateway directory structure looks like the following:

```
<data directory>
|
|--> <Gateway name>
|
|--> edr
```

The MUR application expects the EDR files in the directories that are created when adding the gateway.

The MUR application supports the distributed model to allow the deployment which enables network wide view or work load balancing. Newly introduced component, Remote Data Processor (RDP), plays the role of pre-processing the input files from gateways. One or more RDPs, installed separately on remote machines can be registered to a master MUR system and one RDP can process files from one or more gateways. The role of MUR system in such deployments is mostly for report generation, report viewing, RDP management and optionally data processing.

The RDP parses the raw data or EDR files from GGSNs and periodically forwards it to the registered master MUR application through SFTP for report generation. For information on how to configure the RDPs, see the *Mobility Unified Reporting System Online Help* documentation.

 **Important:** The gateways can be added on Remote Data Processor (RDP). For adding gateway on particular RDP corresponding RDP's region should be selected. RDP region is available when RDP is added. For information on configuring the RDPs, see *Mobility Unified Reporting System Online Help* documentation.

In MUR releases prior to 12.2, each of the registered RDPs were considered to be a new region, hence all reports were based on RDP. Whenever an RDP is configured, internally MUR used to create corresponding region for the same. However, with the introduction and need of scalable MUR, one gateway's files would be processed by two or multiple RDPs. In that case, RDP does not stand as a region. So, reports will be required across all the RDPs under one particular region. Hence, the reports will be available per region.

 **Important:** Whenever an MUR is upgraded from an older version to 12.2, the logical regions for the MUR gets void and the user will not be able to see any gateways under the **DPI/Bulkstats/CF/KPI** tab. In this scenario, the user should add that gateway under NOC.

 **Important:** For HA mode, ensure that you specify the Flow-EDR, HTTP-EDR, CF-EDR, BS Filename path on the shared disk through the MUR GUI by accessing ADMIN > GATEWAYS menu. Also while adding RDP, shared IP address of the RDP node must be specified.

 **Important:** If the password of Master/RDP gets expired, then the Master/RDP server will stop communicating with GGSN. Hence, EDR will not be transferred from GGSN to MUR via SFTP.

Modifying Gateway in GUI in Hierarchical Setup

In the hierarchical setup, before updating/deleting a gateway in GUI, RDP and Master communication should be proper. RDP and Master communication should be checked by updating RDP in GUI.

To update the RDP in GUI, follow these steps:

1. Log on to the MUR GUI.

2. Click **ADMIN -> RDP**. Then, click **Update**.

Managing Archive Directory

To use the Offline Subscriber Search feature seamlessly, you must organize the archive directory date-wise. For information on this feature, see the *Mobility Unified Reporting Online Help* documentation.

MUR organizes the archive directory such that the directory structure looks like the following:

```
<Archive Directory>
  <Gateway Name>
    <Reporting Name>
      <Date YYYYMMDD>
        <Archived Files>
          <Other>
            <Archived Files>
```

For the files that are not satisfying the required EDR file name format, MUR stores the files in the *Other* directory.

Configuring Logging

The MUR application facilitates logging to trace and debug problems identified within the reporting system.



Important: Users with administrative privilege can only manage logging.

Configuring Purging Feature

The MUR application supports purging any kind of aggregated data like half-hourly, daily, weekly, monthly, etc. This also supports purging of weekly summary table, monthly top N table, audit logs, etc. While configuring the purging feature, the MUR provides the flexibility to end-user to configure half-hourly, daily, weekly and monthly report viewing duration so that the historical reports can be viewed even at a lowest granularity level.



Important: It is recommended that you enable backup feature and take one complete successful snapshot of backup before enabling purging feature. If the backup feature is disabled then enabling purging will cause removal of data without waiting for it to be backed up.



Important: The backup snapshot can be identified as complete only if the snapshot directory of format *snapshot_<date>_<time stamp>_<version>*, created under configured backup path does not have the prefix *prog.backup*. The *prog.backup* prefix indicates that the backup is in progress.

MUR uses a python script, *purge_db.py*, to accomplish this task. For more information on the script, refer to the [Using the Purging Script](#) section.



Important: Users with administrative privileges can only manage the purging configuration.



Important: The purging configurations are recommended to be a one-time process and should not be changed frequently.

To configure data / file purging through the GUI, see the *Mobility Unified Reporting System Online Help* documentation.



Important: In case of distributed model of MUR, data purging can be done only at the master MUR and file purging can be performed at per RDP level.



Important: For the MUR software with version 10.0.72 and lesser, you must manually purge the archived files. For the MUR software version 10.0.72 and later, you can use the purging script to automate the process.

Configuring Backup Functionality

To avoid data loss due to hardware failure and/or software crash, MUR supports periodical backup and recovery of its database. The backup is actually the snapshot of data tables or metadata on the day when the backup is taken.

In case of hierarchical deployment, when backup feature is enabled, backup of RDP is also taken as per user configured period.

Backup of RDP contains only the metadata i.e. the configuration information and does not consume high disk space. Also this backup snapshot which is taken on RDP node is immediately transferred to Master MUR using SFTP.

The Master MUR moves the snapshot of RDP backup to the backup path configured on RDP. In short, RDP backup snapshot is also stored on user configured backup path.



Important: Please note that the backup and recovery processes are applicable only for MUR database and not for files that are archived.

Please consider the following points while taking backup of the master MUR database.

- Backup related configuration is available under **ADMIN** tab in Web-based MUR GUI.
- Configure the backup path on a separate disk than using the path where MUR is installed. This can be NFS or any other storage path. The File system on storage disk should support creating hard links to the files for performance benefits. For example: UFS, ZFS and NFS.
- MUR takes the snapshot as per the configured period. If the previous snapshot is available on backup path, data which was not modified between last backup and the current backup is copied directly from the previous snapshot. Thus it is recommended that the backup disk path should hold at least the recent snapshot. The disk size for backup path should be selected accordingly. Older snapshots can be archived or deleted regularly.
- If the backup flag is disabled, the purging of data will continue even if some data tables are pending for backup.
- The backup snapshot can be identified as complete only if the snapshot directory of format `snapshot <date>_<time stamp> <version>`, created under configured backup path does not have the prefix `prog.backup`. The `prog.backup` prefix indicates that the backup is in progress.
- If MUR is being upgraded from a version in which backup and purging features are not available, to a version in which backup and purging features are supported, then it is recommended that you enable backup feature and take one complete successful snapshot of backup before enabling purging feature. If the backup feature is

disabled then enabling purging will cause removal of data without waiting for it to be backed up. If the backup is being taken for the first time after upgrade, then it may take considerable time for first backup.

Configuring Recovery Functionality

To recover the backed up data, use the snapshot recovery script that finds the latest available snapshot amongst all the snapshots under configured path. If you want to recover specific snapshot then move only that snapshot to some other path and provide this new path as a parameter to this script.



Important: In case of hierarchical deployment, recovery of master MUR and RDP should be done separately. After recovering and starting RDP, it will start serving the master to which it is attached.

Please note the following key points while recovering the database.

- The recovery of backed up tables is possible only during a fresh installation of MUR software. The fresh installation version should be same as the version for which snapshot is backed up.
- After fresh installation of server, stop all services except postgres. Postgres server should be running while carrying out recovery steps.
- Fresh installation for recovery purpose should be installed on the same path where last backup is stored and also should have the same IP address and port configuration if the MUR is deployed in distributed mode. Make sure that the existing older installation is either removed or moved to a different directory because the metadata recovered from previously installed MUR will have all references as per older installation e.g. archive path, SFTP details, etc.
- The recovered data contains all the configurations as per older setup. Thus, any changes in the configuration of recovered setup, such as backup interval, etc requires reconfiguration explicitly.

The recovery script provides an option to specify if recovery of data is required for RDP.

For recovery of RDP, backed up snapshot should be copied on a local path or it should be available on path that is accessible to RDP.

For the master MUR, use the following command to recover the data:

```
./recover.sh -path <directory path containing data snapshots>
```

For RDP, use the following command to recover the data:

```
./recover.sh -path <directory path containing data snapshots> -r <RDP_Name>
```

The option `-r <RDP_Name>` denotes the name of RDP for which data should be recovered.

The snapshot of RDP is a tar file and not a directory as in the case of master MUR. The following is a sample of RDP backup tar file naming convention.

```
RDP_<RDP_Name>_snapshot_<Date>_<timestamp>_<version>.tar
```

Configuring Offline Mode

MUR can be used in “Offline Subscriber Search Only” mode. In this mode, MUR will not parse incoming EDR files, so no online reports will be generated. It will move the incoming EDR files to archive directory, so that you can avail “Offline Subscriber Search Reports” only. By default, this mode is turned OFF. To run MUR in Offline Mode, please see the configuration parameters `OFFLINE_MODE` and `OFFLINE_SEARCH_PROCESS_COUNT` in **Managing System Configurations > Config Parameters** in the *Mobility Unified Reporting System Online Help* documentation.

Operations and Management

This section provides information on the scripts that can be used to manage the MUR components and the reports.

Generating Reports in Excel Format

To generate the reports in excel format based on the NOC, execute the following script from the `<MUR_install_dir>/starbi/bin` directory.

```
./get_excel_report.sh -day <date for report generation> -f <path where report should be stored> -filter [1 | 2] <filter for the report>
```

The script takes three parameters, the date for which report is to be generated, the path where generated report is to be stored, and the filter for the reports. The date must be in mm-dd-yyyy format only, and the filter can be based on Device Group or Access Point Name (APN). The script accepts filter 1 for Device Group and filter 2 for APN.



Important: In StarOS release 14.0 and later, the excel reports can be generated and fetched per gateway/RDP level.

To generate Top N subscribers report, the `get_excel_report.sh` script must be used without including the *filter* parameter in the script command syntax.



Important: Active Subscriber and Unique Subscriber reports are available only at NOC level.

To generate the gateway level reports in excel format, execute the following script from the `<MUR_install_dir>/starbi/bin` directory.

```
./get_excel_report.sh -day <date for report generation> -f <path where report should be stored> -filter [1 | 2] <filter for the report> -gatewayname <name of the gateway>
```

To generate the region level reports in excel format, execute the following script from the `<MUR_install_dir>/starbi/bin` directory.

```
./get_excel_report.sh -day <date for report generation> -f <path where report should be stored> -filter [1 | 2] <filter for the report> -regionname <name of the region>
```

Generating Unknown URL Files

For CF reporting, MUR should parse CF-EDRs and generate the unknown/unrated URL database. This database will be pulled periodically by WEM and subsequently deliver to Rulespace. The unknown URL files can either be time based or count based.

To generate the unknown URL files, execute the following script from the `<MUR_install_dir>/server/scripts/cfedr` directory:

```
./gen_unknown_url.py
```



Important: Please note that up to a maximum of 85,000 Unknown URLs can be present in each file.

Loading Blacklist and Whitelist File for Tethering Detection

For seamless tethering analysis and detection, users are recommended to create and maintain blacklist file and whitelist file for UA and OS signatures. This can be accomplished either through the MUR GUI or with the use of script `load_tethering_data.sh` located at `<MUR_Install_Dir>/starbi/bin` directory.

NOTE: Entries from UA/OS whitelist will always be added into final UA-DB/OS-DB. Entries from UA/OS blacklist will always be excluded from the final UA-DB/OS-DB.



Important: Please note that the users require administrator credentials given during MUR installation to access this shell script utility.

Usage of `load_tethering_data.sh` Script:

To load the blacklist / whitelist file:

```
./load_tethering_data.sh [ -b | -w ] -l [ os | ua ] -f <path>
```

To clean the blacklist / whitelist file:

```
./load_tethering_data.sh [ -b | -w ] -c [ os | ua ]
```

Option	Meaning
-b	Blacklist data
-w	Whitelist data
-l	Load blacklist / whitelist data
-c	Clean blacklist / whitelist data
os	OS signature
ua	User Agent
-f <path>	Absolute path for OS signature or user agent file. Note that the OS signature file should contain one os-signature per line, and User Agent file should contain one user agent string per line.
-h, -- help	Prints this help

Loading HTTP Group Data Using `csv` File

MUR uses the shell script utility, `load_httpgroup_data.sh`, to load the metadata associated with HTTP group configuration. This script is located in the `<MUR_install_dir>/starbi/bin` directory.

 **Important:** Please note that the users require administrator credentials given during MUR installation to access this shell script utility.

Hence, after installation/upgrade of MUR, it becomes easy for the operators to do a bulk load of HTTP groups using the *csv* file. The *csv* file has HTTP group names in column 1 and the regular expression in column 2.

The following are a few examples of *csv* file:

Search,.*google.*

Yahoo,.*yahoo.*

Facebook,.*facebook.*

News,.*news.*

facebook,.*facebooknew.*

yahoo,.*yahoo2.*

The HTTP groups and the regular expressions in the *csv* file can be loaded to MUR using the script **load_httpgroup_data** with *csv* file name as an argument.

Once successfully loaded, the HTTP group details can be accessed through the MUR GUI under **System** tab in **Edr http group regexs** reports. If the HTTP group name in *csv* file is already configured in MUR, the regular expression is appended to the existing expressions.

Usage of **load_httpgroup_data.sh** Script:

To load the HTTP group information to MUR GUI:

```
load_httpgroup_data.sh -p xxx.csv
```

Where, *xxx* denotes the *csv* file name.

Purging Tethering Database

The shell script, *purge_tethering_data.sh*, handles purging of OS signature, UA and TAC databases. This script is packaged with MUR in the *<MUR_install_dir>/starbi/bin* directory.

This script purges database tables (UA/OS/TAC) and final OS/UA/TAC files based on provided configuration.

 **Important:** Please note that the users require administrator credentials given during MUR installation to access this shell script utility.

Usage of **purge_tethering_data.sh** Script:

To purge the tethering database:

```
./purge_tethering_data.sh db [ os | ua | tac ]
```

To purge tethering database files:

```
./purge_tethering_data.sh files [ version ] [ os | ua | tac ]
```

Option	Meaning
db	Purge database tables
os	Purge OS Signature table / files

Option	Meaning
ua	Purge User Agent table / files
tac	Purge TAC table / files
files	Purge database files
version	Purge files with version less than the provided version
-h, --help	Prints this help

Resetting GUI Administrator User Password

In case the Administrator user forgets the password, a *set_admin_password.py* script is used to reset the password. This script is located at the *<MUR_install_dir>/starbi/server/scripts* directory.

To reset the Administrator user's password, perform the following steps.

Step 1 Execute the following commands to set the environment variables.

```
#source server/env.properties
#export PYTHONPATH
#export LD_LIBRARY_PATH
```

Step 2 Execute the following script.

```
#python2.6.1/bin/python server/scripts/set_admin_password.py
```

The script will update MUR database with Administrator user password as *admin*.

Using the generate_dns_mapp_sql.sh Script

To generate the DNS mapping for the specified list of IP addresses, execute the following script from the *<MUR_install_dir>/starbi/bin* directory:

```
./generate_dns_mapp_sql.sh <input file for IP> <output file where mapping should be stored>
```

Keyword/Variable	Description
input file for IP	A file containing IP addresses. Each IP address must be present in a new line.
output file where mapping should be stored	An output file for storing the DNS mappings in SQL format.

This script is used to perform Internet DNS lookup of the specified IP addresses. It uses the 'nslookup' system administration command to find the DNS name of the specified IP. Please note that the machine must be connected to Internet for successful execution.

Using the getSupportDetails Script

In the event additional troubleshooting assistance is required, debugging information can be collected using a script called *getSupportDetails.pl*. This script collects different log files and captures the output of certain system commands that aid in troubleshooting issues. This script is packaged with MUR in the `<MUR_install_dir>/starbi/tools/supportdetails/` directory.

This script refers to an XML file to get the list of logs. This XML file resides in the same directory as the script. Once executed, the script retrieves the contents of logs, files, folders, and output of certain commands and prepares a zipped file (*MURsupportDetails.tar.gz*), by default it is placed in `/tmp/log` directory.

Requirements

Perl 5.8.5 and above is required for running the script.

Apart from standard Perl modules (which are included in default installation of Perl), some additional modules are required for running the script. The list is as follows:

- expat version 1.95.8
- XML::Parser version 2.34
- XML-Parser-EasyTree
- Devel-CoreStack version 1.3

These modules are installed by default by the product. Please ensure that the above mentioned modules are installed when using a different installation of Perl.

To run the script, change to the directory path where the script is present and type:

```
./getSupportDetails.pl [--level=...] [--xmlfile=...] [--help]
```

Keyword/Variable	Description
--level	Specifies the level of debug to run. It can have a maximum of 4 levels. The level 4 provides the most detailed information. Default: 1
--xmlfile	Specifies the xml file name to be used for collecting the log. Default: getSupportDetails.xml
--onlyrecentlogs	Collects only recent logs and skips detailed logs. Default: Collects detailed logs
--collectFor	Collects problem specific logs and information which is not collected under normal levels. This can be combined with --level option. Default: Collects logs covered under '--level' option.
--help	Displays the supported keywords/variables.



Important: If the change in the log level is not reflected in the GUI, be sure to restart the Notif Server.

For example, `./getSupportDetails.pl --level=4 --xmlfile=/tmp/getSupportDetails.xml`

Follow this procedure to run the *getSupportDetails* script:

1. Change to the directory path `<MUR_install_dir>/starbi/tools/supportdetails/`.

2. Login as *muradmin* user.
3. Set the environment variables.
4. Run the *getSupportDetails* script.

Note that if the Steps 2 and 3 are interchanged, then an exception message is thrown as shown below:

```
bash-3.00$ ./getSupportDetails.pl

Enviornment Variable is set ...Collecting logs for level 1...Collecting Important
statistics

ld.so.1: python: fatal: libpython2.6.so.1.0: open failed: No such file or
directory

/tmp/log/supportDetails/collected-data/Extracting information from Log file
/export/home/14.0.314/starbi/logs/server/starbi_server_develld.so.1: python:
fatal: libpython2.6.so.1.0: open failed: No such file or directory

mv: cannot access *.txt

/tmp/log/inPilotsupportDetails.tar.gz Generated.
```

Supported Levels

The logs that can be collected for different levels are as follows:

- Level 1:
 - Recent Log files
 - Current status (running / not running) of the product
 - Current Config files of the product
 - If database is up and running, important database statistics like configuration and parsing configuration is collected
 - Space occupied by EDR and its subfolders
- Level 2:
 - Logs from Level 1
 - Installation Logs
 - Database Logs (if available)
 - Web Server logs (if available)
 - Information of Solaris version and current patch installed
 - Output of the following commands:


```
netstat -an
ifconfig -a
df -k
etc..
```
- Level 3:

- Logs from level 2
- Syslog Configuration and log files
- Level 4:
 - Logs from level 3
 - All Log files (including old logs)
 - Crontab entries
 - Information of packages installed
 - Stack trace of any crash files (if debugger is installed on local machine)
 - System Libraries only if any core file present in crash directory
 - Level of Solaris installed
 - Output of the following commands:


```
ipcs
ps -eaf
etc..
```

Using the Maintenance Utility

A shell script utility called *serv* is included with MUR in the *<MUR_install_dir>/starbi/bin* directory.

This *serv* script can be used to manage the following MUR processes:

- Process Monitor (PSMON) Application
- Scheduling Server
- Postgres Server
- Apache Server
- Notif Server
- Parser Server
- Cache Server

This utility can report the status of the MUR processes on the system or it can be used to stop the MUR process.

Following are the options available with the *serv* script:

```
./serv { psmonitor | scheduler | postgres | apache | notif_server | parser | cache_server
} [ start | stop | status ]
```

Keyword	Description
psmonitor	This is an optional keyword used with the serv script. This represents the PSMON application. It starts/stops/checks the following MUR processes. • Postgres server • Apache server • Scheduling server • Notif server • Parser server • Cache server
scheduler	This is an optional keyword representing the scheduling server.
postgres	This is an optional keyword representing the postgres server.
apache	This is an optional keyword representing the apache server.
notif_server	This is an optional keyword representing the notif server.
parser	This is an optional keyword representing the parser server.
cache	This is an optional keyword representing the cache server.
start	Starts each MUR process.
stop	Kills or stops the running MUR process.
status	Displays the status of each MUR process. By default, it will show the status of all the MUR processes.

For example, if you want to start only the PSMON, then enter the following command:

```
./serv start psmonitor
```

or

```
./serv psmonitor start
```



Important: If you stop the MUR process, make sure that PSMON is not running. Otherwise PSMON will restart the MUR application.

The following is a sample output of the **serv status** command:

```
-----
----- MUR Process Status -----
PID           Process           Status
-----
4245          Process Monitor    Running
4256          Scheduling server   Running
4267          Postgres Server     Running
```

4289	Apache Server	Running
3249	Notif Server	Running
3243	Parser Server	Running
2430	Cache Server	Running

Using the PSMON Script

PSMON is a perl script that is used to monitor the Scheduling Server, Postgres Server, and Apache Server processes. This script can start or stop the processes based on certain thresholds specified in the MUR configuration file. The PSMON respawns any dead processes using the set of rules defined in the configuration file.

This script can also optionally send notifications to users via e-mail.

Using the Purging Script

The python script, *purge_db.py*, handles both data purging and archived file purging. This script is packaged with MUR in the `<MUR_install_dir>/starbi/server/scripts/utls/` directory.

This script runs daily at the end of the day, picks up the relevant tables, and then purges either data or archived files based on the configurations.

In case of data purging, the script picks up the relevant tables and purges them.

In case of file purging, the script purges the files only if the archived files are organized date-wise for each of the reportings like FLOW-EDR, HTTP-EDR, CF-EDR, and BS. For example, EDR file for 24th September, 2010 is stored in the *archive/<gw>/flowedr/20100924* directory.

Using the unanonymize_msisdn.sh Script

MUR reports the subscribers data like Mobile Station Integrated Network (MSISDN) in the encrypted format both in the GUI and Excel file. Decryption functionality is ONLY supported on CLI through the use of `unanonymize_msisdn.sh` script.

This shell script utility will check for user's privilege before decrypting the MSISDNs. It will prompt for the GUI administrator password.

Usage of unanonymize_msisdn.sh Script:



Important: Please note that this script must be run ONLY in bash shell.

```
./unanonymize_msisdn.sh -u <username> -f <input file> -d <output path>
```

Option	Meaning
-u	Used to specify GUI Administrator user name.

Option	Meaning
-f	Used to specify the absolute input file path of a file containing list of anonymized MSISDNs. Each anonymized MSISDN will be separated by a new line.
-d	Used to specify the output directory path where the decrypted MSISDNs file will be stored.
-h, --help	Prints this help

To decrypt the MSISDN(s), perform the following steps:

1. Get an excel report for the day for which you want to see the clear text top subscribers or MSISDN(s). For information on how to get the excel report, refer to the [Generating Reports in Excel Format](#) section in this chapter.
2. Copy all lines from "Anonymized MSISDN" column of work sheet "Top 1000 Subscribers Traffic" and paste them in to a separate text file.
3. Provide this text file as an input to the `unanonymize_msisdn.sh` script.



Important: Please note that the users require administrator credentials given during MUR installation to access this utility.

Server Script Parameters

The number of files being processed during each parsing interval for HTTP and non-HTTP EDRs can be controlled using the following parameters:

- EDR_TOTAL_NO_OF_FILES = 25
- EDR_MAX_NO_OF_PROCESSES = 5
- HTTP_TOTAL_NO_OF_FILES = 25
- HTTP_MAX_NO_OF_PROCESSES = 5

These parameters are defined in **System** menu under **File parsing configs** option available on the GUI.

With the above default configuration, if the number of files being accumulated are less than 25 and not in multiples of 5, then MUR spawns one more process to parse the remaining files.

Chapter 8

Scalable Solution for MUR

In case of achieving high throughput (more than 10 Gbps), one RDP is not sufficient to process data. Scalability feature enables load distribution, so that data processing can be performed at multiple RDPs to support high data rate.

This chapter comprises the following topics:

- [Scalable MUR Overview](#)
- [System Requirements for Scalable MUR](#)
- [Deployment of Scalable MUR](#)
- [Hardware Configurations](#)
- [Installation Upgrade to Scalable Model](#)

Scalable MUR Overview

Scalability is achieved with the help of clustering of RDP nodes which share a common storage. The gateway pushes the input files to the shared storage through one of the RDPs, and all the RDPs then pick up those files as per certain rules defined for parsing. The parsing of files is based on the filename pattern that is specified while configuring the gateway. The MUR parsers match regular expressions given in the file name pattern and accordingly pick up the files. For example, users can specify `*flow_*1` to match all flow EDR files ending with one to be picked up for that gateway. This way, multiple nodes parse the files in parallel. In such a cluster, maximum 16 nodes can be operating in parallel.

RDPs pick up the files based on the gateways that are attached to them. In the scalability model, one main instance of the gateway is configured on one of the RDPs. For the rest of the RDPs, a 'pseudo' instance of the same gateway is configured. This means, while adding another instance of the same gateway on another RDP, this gateway is marked as pseudo (through the MUR GUI) of an existing gateway (gateway configuration on GUI asks for name of a gateway to which gateway being configured is a pseudo). This configuration informs MUR that while processing data for pseudo gateway, this will not be considered as a separate gateway, but as the original gateway only.

In order to enable the RDPs to pick up files distinctly from each other, the configuration 'File Name Pattern' on the gateway configuration GUI is used. User can specify certain patterns for the file names for gateway (e.g. `edr_*[1,3,5,7,9]` to pick up only odd files and `edr_*[0,2,4,6,8]` to pick up only even numbered files) during such configuration. So, in the case of two RDPs, the user can specify pattern matching odd numbered files for first gateway that is configured on first RDP, and pattern matching even numbered files for second (pseudo) gateway that is configured on second RDP.

While viewing reports, in order to make sense to users, the pseudo gateway is NOT shown in the gateway hierarchical tree. The gateway tree is typically available under DPI, HTTP, CF, and Bulkstats tabs. However, the pseudo gateway information will be available through the **Admin** tab.


Basic Scalability Model

This section highlights the basic system requirements for the scalability in master MUR and RDPs.

- Master — Internal disk: OS, MUR (RAID-0), Postgres (RAID-5)
- RDP(s) — Internal disk: OS, MUR (RAID-0), Postgres (RAID-5), Archive (archive can be shifted to shared disk based on the requirement and sizing)
- Shared-disk — Only incoming files (RAID-0)

The following figure outlines the basic model of the scalability implementation in MUR.

Figure 6. Scalable MUR



All cluster nodes need to be interconnected using switch and need to be in same VLAN/ shared disk. After the cluster installation, the shared disk needs to be configured in all RDP nodes. Check if the nodes are connected using `multipath -ll` command.

System Requirements for Scalable MUR

Hardware Requirements

- UCS 460 M2 machine with full internal disks capacity for each node in scalable deployment identified as RDP
- UCS 460 M2 machine with internal SAS disks capacity for Master in scalable model
- Dual port Fiber Channel HBA card for each RDP node.
- Sun StorageTek 2540 External Storage Array with all four controllers
- Fiber Channel Cables

Software Requirements

- **Operating System:** Cisco_MITG_RHEL_5.5_12.2.10
- Veritas Storage Foundation Enterprise Cluster File System Ha/Dr 5.1 S64Lnx
- MUR version 14.0 onwards
- CAM software (to be downloaded from Oracle site) for managing StorageTek

Deployment of Scalable MUR

Following is the sequence of steps to be carried during the fresh deployment of Scalable MUR.

- Step 1** Arrange UCS machines as RDP hosts and master.
- Step 2** Configure network link and assign IP addresses.
- Step 3** Arrange two separate private Ethernet heartbeat links between RDP hosts and connect them via switch if needed.
- Step 4** Install and configure RHEL on RDP hosts.
- Step 5** Do disk partitioning of local disks on RDP hosts while installing RHEL.
- Step 6** Connect Fiber channel cables between RDP hosts and StorageTek.
- Step 7** Arrange and install CAM software.
- Step 8** Configure volumes on StorageTek.
- Step 9** Map volumes to all the RDP hosts.
- Step 10** Configure multipath service on each of the RDP hosts.
- Step 11** Reboot and test if all redundant paths are in ready state and are available on each of the RDP hosts.
- Step 12** Install and configure Veritas.
- Step 13** Test if shared path is available on all the nodes.
- Step 14** Reboot the nodes after Veritas installation and check if all the paths are available again after rebooting all the hosts.
- Step 15** Check if the output of this command **gabconfig** is not in “jeopardy” state after reboot.
- Step 16** Install RDPs on each RDP hosts.
- Step 17** Add the RDPs on Master.
- Step 18** Add Gateways on Master and do all the necessary configurations.
- Step 19** Activate EDR generation on ASR node and check if files are successfully transferred to shared location.
- Step 20** Check if each RDP has started processing files and files are not pending in input directory.
- Step 21** Check if the reports are seen on Master.
- Step 22** Do the post installation recommended configurations on MUR.
- Step 23** Keep monitoring the setup for a week and test if disk partitions are in control and reports are satisfactory and there are no pending files on the disks.

Hardware Configurations

Heartbeat Link between RDP Hosts

Veritas Cluster file system recommends having two private dedicated heartbeat links interconnection between all the nodes which are accessing cluster file system. Hence for all the RDP nodes, interconnect each dedicated heartbeat link separately through switch. VCS does not require these two links to have IP configured and it automatically detects these links during installation.



Important: The link having IP addresses to access these RDP hosts is separate than the heartbeat links.

Installing RHEL on RDP Hosts

Install the recommended version of the RHEL on all the RDP hosts. The current recommended RHEL version for UCS 460 machines is *Cisco_MITG_RHEL_5.5_12.2.10*.

While installing the RHEL, use disk partitioning wizard to do all the required partitioning on the RDP hosts as specified below. For more information, see the *Configuring MegaRAID for MUR Applications* and *Cisco UCS Server Hardware Configuration for MUR Applications* chapters in this guide.

Configurations Post RHEL Installation

The following section describes the configurations that should be done post RHEL installation.

When UCS reboots, press F8 (CMIC config) and configure the following:

- NIC Mode: Dedicated
- NIC Redundancy: None
- DHCP: Disabled, IPs configured
- VLAN ID: Disabled

While installing the MITG RHEL5.5, execute the following command at boot prompt:

```
boot: linux ks=cdrom:/ks.cfg
```

This will result in the following:

1. Skipping the installation key
2. Disabling the firewall
3. Staring FTP server
4. Permitting selinux policy by default

While partitioning the disks on UCS, see the *Configuring MegaRAID for MUR Applications* and *Cisco UCS Server Hardware Configuration for MUR Applications* chapters in this guide.

After RHEL has been installed, perform the following steps:

1. Set the hostnames (hostname for each RDP node) in */etc/hosts* of all the nodes and restart the networking by executing the following command:

```
/etc/init.d/network restart
```

2. If the interfaces have not been configured while installing the OS, please configure them now and then follow the below step:

- Restart the network interfaces by executing the following command:

```
/etc/init.d/network restart
```

3. Check if the firewall is disabled on platform and SELinux is set to permissive by executing the following command:

```
#system-config-security-level-tui
```

4. Check that sshd demon is running, if not then start it by executing the following commands:

```
/etc /etc/init.d/sshd start
```

```
/etc/init.d/sshd status
```

Disk Partitioning for RDP

The disk partitioning should be performed as described in the *Configuring MegaRAID for MUR Applications* and *Cisco UCS Server Hardware Configuration for MUR Applications* chapters of this guide.

The actual sizing of the disk partitions should be chosen by following the MUR disk sizing sheet. For more information on the sizing sheet, contact your local Cisco account representative.

However in the case of scalable model, the input directory on the shared storage is required. Hence, a smaller disk partition can be chosen for the MUR application, and a large disk partition for the MUR database.

Below is a sample disk partitioning for RDP serving throughput up to 10 Gbps. This is considering 600 GB disk size each.

- Single (1) disk in RAID-0 for the OS.
- RAID-0 array with (1) disk for the MUR application. This partition can also be placed on the disk on which operating system is installed but with separate disk partition.
- RAID-5 array with all remaining (10) disks for the Postgres Database.



Important: The above partitions refer to partitions on local disk. Separate partition on external storage has to be created for incoming files and archive files as explained in the later section.

Disk Partitioning for MUR Master

Disk partitioning for MUR Master should be the same as that for RDP. However, database partition can be smaller depending on the sizing.

Additionally there should be separate partition to be configured for archiving. MUR sizing excel tool describes sizing for the Master node. This archiving would be used only to store the bulkstat files.

Below is a sample disk partitioning for Master serving throughput up to 40 Gbps. This is considering 600 GB disk size each.

- Single (1) disk in RAID-0 for the OS.
- RAID-0 array with (2) disks for the MUR application. This partition can also be placed on the disk on which operating system is installed but with separate disk partition.
- RAID-5 array with (6) disks at least for the Postgres Database.
- RAID-0/1 array with (1) disks for the archive partition.

Please note down the following commands output after the local disk partitioning is done so that it would be useful for further steps before connecting or mapping the StorageTek volumes.

- **mount**
- **df -kh**
- **fdisk -l**

The following is the console trace.

```
[root@pnextappsucs460-1 ~]# mount
/dev/sda1 on / type ext3 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw)
/dev/sda3 on /apps type xfs (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
[root@pnextappsucs460-1 ~]#
[root@pnextappsucs460-1 ~]# df -kh
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda1	95G	3.8G	86G	5%	/
tmpfs	63G	0	63G	0%	/dev/shm
/dev/sda3	399G	5.3G	393G	2%	/apps

```
[root@pnextappsucs460-1 ~]#
[root@pnextappsucs460-1 ~]#
[root@pnextappsucs460-1 ~]# fdisk -ll
```

Disk /dev/sda: 599.5 GB, 599584145408 bytes
 255 heads, 63 sectors/track, 72895 cylinders
 Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	1	12748	102398278+	83	Linux

```

/dev/sda2          12749          20907      65537167+   82   Linux swap / Solaris
/dev/sda3          20908          72895      417593610   83   Linux

```

WARNING: GPT (GUID Partition Table) detected on '/dev/sdb'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.0 TB (5995841454080 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

```

Disk /dev/sdb: 5995.8 GB, 5995841454080 bytes
255 heads, 63 sectors/track, 728952 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

```

```

      Device Boot      Start         End      Blocks   Id  System
/dev/sdb1                1       267350    2147483647+   ee  EFI GPT

```

WARNING: GPT (GUID Partition Table) detected on '/dev/sdk'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.6 TB (6595489038336 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

```
Disk /dev/sdk: 6595.4 GB, 6595489038336 bytes
255 heads, 63 sectors/track, 801855 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdk1		1	267350	2147483647+	ee	EFI GPT

WARNING: GPT (GUID Partition Table) detected on '/dev/sdl'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.0 TB (5995899125760 bytes).
 DOS partition table format can not be used on drives for volumes
 larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID
 partition table format (GPT).

```
Disk /dev/sdl: 5995.8 GB, 5995899125760 bytes
255 heads, 63 sectors/track, 728959 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdl1		1	267350	2147483647+	ee	EFI GPT

```
Disk /dev/sdm (Sun disk label): 255 heads, 189 sectors, 24296 cylinders
Units = cylinders of 48195 * 512 bytes
```

Device	Flag	Start	End	Blocks	Id	System
/dev/sdm3	u	0	24296	585472860	5	Whole disk

```
/dev/sdm8  u          0      24296 585472860      f  Unknown
```

Disk /dev/sdm3 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders

Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdm3p3	u	0	24296	585472860	5	Whole disk
/dev/sdm3p8	u	0	24296	585472860	f	Unknown

Disk /dev/sdm8 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders

Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdm8p3	u	0	24296	585472860	5	Whole disk
/dev/sdm8p8	u	0	24296	585472860	f	Unknown

WARNING: GPT (GUID Partition Table) detected on '/dev/sdo'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.6 TB (6595489038336 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/sdo: 6595.4 GB, 6595489038336 bytes

255 heads, 63 sectors/track, 801855 cylinders

Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/sdo1		1	267350	2147483647+	ee	EFI GPT

WARNING: GPT (GUID Partition Table) detected on '/dev/sdp'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.0 TB (5995899125760 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/sdp: 5995.8 GB, 5995899125760 bytes

255 heads, 63 sectors/track, 728959 cylinders

Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/sdp1		1	267350	2147483647+	ee	EFI GPT

Disk /dev/sdq (Sun disk label): 255 heads, 189 sectors, 24296 cylinders

Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdq3	u	0	24296	585472860	5	Whole disk
/dev/sdq8	u	0	24296	585472860	f	Unknown

Disk /dev/sdq3 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders

Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdq3p3	u	0	24296	585472860	5	Whole disk
/dev/sdq3p8	u	0	24296	585472860	f	Unknown

Disk /dev/sdq8 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders

Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdq8p3	u	0	24296	585472860	5	Whole disk
/dev/sdq8p8	u	0	24296	585472860	f	Unknown

WARNING: GPT (GUID Partition Table) detected on '/dev/sdw'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.6 TB (6595489038336 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/sdw: 6595.4 GB, 6595489038336 bytes

255 heads, 63 sectors/track, 801855 cylinders

Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/sdw1		1	267350	2147483647+	ee	EFI GPT

WARNING: GPT (GUID Partition Table) detected on '/dev/sdx'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.0 TB (5995899125760 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/sdx: 5995.8 GB, 5995899125760 bytes
 255 heads, 63 sectors/track, 728959 cylinders
 Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/sdx1		1	267350	2147483647+	ee	EFI GPT

Disk /dev/sdy (Sun disk label): 255 heads, 189 sectors, 24296 cylinders
 Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdy3	u	0	24296	585472860	5	Whole disk
/dev/sdy8	u	0	24296	585472860	f	Unknown

Disk /dev/sdy3 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders
 Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
--------	------	-------	-----	--------	----	--------

```

/dev/sdy3p3  u          0      24296 585472860      5  Whole disk
/dev/sdy3p8  u          0      24296 585472860      f  Unknown

```

Disk /dev/sdy8 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders
 Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdy8p3	u	0	24296	585472860	5	Whole disk
/dev/sdy8p8	u	0	24296	585472860	f	Unknown

WARNING: GPT (GUID Partition Table) detected on '/dev/sdz'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.6 TB (6595489038336 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/sdz: 6595.4 GB, 6595489038336 bytes
 255 heads, 63 sectors/track, 801855 cylinders
 Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/sdz1		1	267350	2147483647+	ee	EFI GPT

WARNING: GPT (GUID Partition Table) detected on '/dev/sdaa'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.0 TB (5995899125760 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/sdaa: 5995.8 GB, 5995899125760 bytes
 255 heads, 63 sectors/track, 728959 cylinders
 Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/sdaa1		1	267350	2147483647+	ee	EFI GPT

Disk /dev/sdab (Sun disk label): 255 heads, 189 sectors, 24296 cylinders
 Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdab3	u	0	24296	585472860	5	Whole disk
/dev/sdab8	u	0	24296	585472860	f	Unknown

Disk /dev/sdab3 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders
 Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdab3p3	u	0	24296	585472860	5	Whole disk
/dev/sdab3p8	u	0	24296	585472860	f	Unknown

Disk /dev/sdab8 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders

Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/sdab8p3	u	0	24296	585472860	5	Whole disk
/dev/sdab8p8	u	0	24296	585472860	f	Unknown

WARNING: GPT (GUID Partition Table) detected on '/dev/dm-0'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.0 TB (5995899125760 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/dm-0: 5995.8 GB, 5995899125760 bytes

255 heads, 63 sectors/track, 728959 cylinders

Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/dm-0p1		1	267350	2147483647+	ee	EFI GPT

Disk /dev/dm-1 (Sun disk label): 255 heads, 189 sectors, 24296 cylinders

Units = cylinders of 48195 * 512 bytes

Device	Flag	Start	End	Blocks	Id	System
/dev/dm-1p3	u	0	24296	585472860	5	Whole disk
/dev/dm-1p8	u	0	24296	585472860	f	Unknown

WARNING: GPT (GUID Partition Table) detected on '/dev/dm-2'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.0 TB (5995841454080 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/dm-2: 5995.8 GB, 5995841454080 bytes

255 heads, 63 sectors/track, 728952 cylinders

Units = cylinders of 16065 * 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/dm-2p1		1	267350	2147483647+	ee	EFI GPT

WARNING: GPT (GUID Partition Table) detected on '/dev/dm-3'! The util fdisk doesn't support GPT. Use GNU Parted.

WARNING: The size of this disk is 6.6 TB (6595489038336 bytes).

DOS partition table format can not be used on drives for volumes larger than 2.2 TB (2199023255040 bytes). Use parted(1) and GUID partition table format (GPT).

Disk /dev/dm-3: 6595.4 GB, 6595489038336 bytes

255 heads, 63 sectors/track, 801855 cylinders

```
Units = cylinders of 16065 * 512 = 8225280 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/dm-3p1		1	267350	2147483647+	ee	EFI GPT

```
Disk /dev/dm-4: 5995.8 GB, 5995841419776 bytes
```

```
255 heads, 63 sectors/track, 728952 cylinders
```

```
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Disk /dev/dm-4 doesn't contain a valid partition table
```

```
Disk /dev/dm-5: 5995.8 GB, 5995899084288 bytes
```

```
255 heads, 63 sectors/track, 728959 cylinders
```

```
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Disk /dev/dm-5 doesn't contain a valid partition table
```

```
Disk /dev/dm-6: 6595.4 GB, 6595488996864 bytes
```

```
255 heads, 63 sectors/track, 801855 cylinders
```

```
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Disk /dev/dm-6 doesn't contain a valid partition table
```

```
[root@pnextappsucs460-1 ~]#
```

Cabling between StorageTek and RDP

The Fiber channel cable needs to be connected between Sun StorageTek, controllers and HBA ports of the RDP nodes.

It is required to have cabling connections so that each RDP have redundant access to StorageTek for fault tolerance and performance gain. If the RDP host is using dual channel HBA card then one port should get connected to one controller and other port should get connected to other controller. Below are some guidelines. Two hosts can be directly connected to StorageTek however more than two hosts must be connected using Fiber Switch.

Please make sure that all the redundant paths are available and so it is recommended to have all the controllers connected.

Figure 7. Connecting two hosts to StorageTek using Direct Connection





Figure 8. Connecting more than two hosts to StorageTek using Fiber Switch



Test the connections between StorageTek and the RDP hosts using the procedure described in the [Installing the Management Software \(CAM\)](#) section of this guide.

Configuring StorageTek

MUR application needs to have two volumes mapped onto the StorageTek for incoming files storage and for archival purpose. Each would be shared between all the connected RDP machines.

The method to configure the single volume and mapping it to all the hosts from StorageTek through CAM software and how to install the CAM software is explained in the following section.

Installing the Management Software (CAM)

This section describes the method to test the connections between StorageTek and the RDP hosts.

Once you copy the storage software on a machine, please make sure that following directories and files have execute permissions.

- linux/util/*
- linux/bin/tools/*
- linux/components/lockhartLinux/*
- linux/RunMe.bin

```
[root@intracerR CAM_linux]# ./RunMe.bin -c
```

```
Initializing Wizard.....
```

```
Launching InstallShield Wizard.....
```

```
-----
Sun StorageTek(TM) Common Array Manager 6.2
```

```
The InstallShield Wizard will install Sun StorageTek(TM)
Common Array Manager on your computer.
```

```
To continue, choose Next.
```

```
Sun StorageTek(TM) Common Array Manager 6.2
```

```
Sun Microsystems, Inc.
```

```
http://www.sun.com
```

```
Press 1 for Next, 3 to Cancel or 5 to Redisplay [1]
```

Sun StorageTek(TM) Common Array Manager 6.2

Please read the following license agreement carefully.

Sun StorageTek(TM) Common Array Manager

Copyright 2008 Sun Microsystems, Inc. All rights reserved. Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries. U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements. Use is subject to license terms. This distribution may include materials developed by third parties. Portions may be derived from Berkeley BSD systems, licensed from U. of CA. Sun, Sun Microsystems, the Sun logo, Java, Solaris and Sun StorageTek Common Array Manager are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries.

Please choose from the following options:

☐ 1 - I accept the terms of the license agreement.

☒ 2 - I do not accept the terms of the license agreement.

To select an item enter its number, or 0 when you are finished: [0] 1

[X] 1 - I accept the terms of the license agreement.

☐ 2 - I do not accept the terms of the license agreement.

```
To select an item enter its number, or 0 when you are finished: [0]
```

Press 1 for Next, 2 for Previous, 3 to Cancel or 5 to Redisplay [1]

Sun StorageTek(TM) Common Array Manager 6.2

Choose the installation type that best suits your needs.

[X] 1 - Typical

The program will be installed with the suggested configuration.

Recommended for most users.

[] 2 - Custom

The program will be installed with the features you choose.

Recommended for advanced users.

```
Select the number corresponding to the type of install you would like: [0]
```

Press 1 for Next, 2 for Previous, 3 to Cancel or 5 to Redisplay [1]

```
Checking current system ...
```

Sun StorageTek(TM) Common Array Manager 6.2

Software To Be Installed:

Full Install

* Browser User Interface (BUI)

* Local and Remote CLI

* Array Firmware

Press 1 for Next, 2 for Previous, 3 to Cancel or 5 to Redisplay [1] Preparing for installation ...

Pre Uninstall Old Action ...

Removing old features ...

Sun StorageTek(TM) Common Array Manager 6.2

Installing Sun StorageTek(TM) Common Array Manager 6.2. Please wait...

```
|-----|-----|-----|-----|
0%          25%          50%          75%          100%
|||||
```

Installing Java 2 Standard Edition

Sun StorageTek(TM) Common Array Manager 6.2

```

|-----|-----|-----|-----|
0%          25%          50%          75%          100%

```

```

||||||||||||||||||||||||||||||||||||||||||||||||||||||||

```

```

-----

```

Sun StorageTek(TM) Host Software Installation Summary

View results:

Info:

Installation success.

The following have been installed: Browser User Interface (BUI), Local and Remote CLI, and Array Firmware.

To access the Browser User Interface point a browser at:

`https://installation_host:6789`

The logs may be found in `/var/opt/cam/`

Press 3 to Finish or 5 to Redisplay [3]

Accessing the Storage Management GUI





Follow these steps to access the storage management GUI.

1. Put Eth0 of both nodes and Port-1 on both storage controllers in a private VLAN. Configure the following IP on Eth0 of the node managing the storage:
`ifconfig eth0:array 192.168.128.110 netmask 255.255.255.0 up`
Default Controller A IP: 192.168.128.101
Default Controller B IP: 192.168.128.102
2. Access the Management GUI using a browser on PC (65.198.111.26 is the public IP of the node which management software was installed)
`https://65.198.111.26:6789`
3. The first time login to the CAM software is always through the admin user of the operating system. For example, Administrator on Windows and root on the unix/Linux.

Configuring the Storage System

Perform the following steps to configure the storage system.


1. Discover the storage system by clicking **Storage Systems** -> **Register** -> **Scan the local network**.





2. Create a storage pool for co-ordinator disks (to be used for I/O fencing) by clicking **Storage Systems -> (discovered storage) -> Pools -> New**.

- Create profile for RAID 0 disks. (This can also be RAID 1)




- Create a pool for CPDisks.

Hardware Configurations



381289



381290


3. Create a volume and map it by selecting this option **Storage Systems -> (discovered storage system) -> Volume -> New.**



381291




381291



381293



381292



381295

Hardware Configurations

Sun StorageTek™ Common Array Manager

Create New Volume

Steps: Help Step 3: Review Selections

1. Select Storage Characteristics

1.1 Select Storage - Create Pool

1.2 Select Storage - Populate Pool

2. Configure Volume

2.1 Select Mapping

2.2 Select a Host or Host Group

3. Review Selections

Review and confirm your selections.

No Hot-Spares
There are currently no hot-spares.

Volume Name: Create 3 volumes with 'CPVol' as prefix
Volume Capacity: 50 MB each
Pool Name: Create Storage Pool 'CPPool'
Storage Profile: Create Storage Profile 'CPProfile'
Storage Characteristics: RAID 0, Segment Size: 256 KB, Read-Ahead: false, Number Of Disks: Variable, Disk Type: ANY
Virtual Disk Name: Chosen automatically by CAM
Controller: Any
Mapping To: Default Storage Domain - LUN/Next Available

Previous Finish Cancel

381296

APPLICATIONS VERSION

User: Administrator Server: PNOAHP0V1

Sun StorageTek™ Common Array Manager

REFRESH SEARCH SERVICE ADVISOR LOG OUT HELP

Current Logins: 1
Last Update: Jun 28, 2013 1:09:43 PM IST
Current Alarms: 0 17 0

Alarms

Storage Systems

clusters5_6

pnDevIntracsr-STek2540

pnQAIntracsr-STek2540

Alarms

Volumes

Mappings

Snapshots

Virtual Disks

Pools

Profiles

Jobs

Physical Devices

Troubleshooting

Administration

unlabeled

unlabeled

VFIScalableTest

General Configuration

Virtual Disk Summary on Storage System pnQAIntracsr-STek2540

To view details for a virtual disk, click on its name below.

Virtual Disks (1)

New Delete

Name	Status	State	Disk Type	RAID Level	Used Percent	Total Capacity	Configured Capacity	Maximum Volume Size
1	Optimal	Ready	SAS	0	0%	136,227 GB	100,000 MB	136,129 GB

New Delete

381297

- Create volumes for input and archive.
- Create Storage Pool for these volumes and then actual volumes.

APPLICATIONS VERSION

User: Administrator Server: PNOAHP0V1

Sun StorageTek™ Common Array Manager

REFRESH SEARCH SERVICE ADVISOR LOG OUT HELP

Current Logins: 1
Last Update: Jun 28, 2013 6:59:43 PM IST
Current Alarms: 0 17 0

Alarms

Storage Systems

clusters5_6

pnDevIntracsr-STek2540

pnQAIntracsr-STek2540

Alarms

Volumes

Mappings

Snapshots

Virtual Disks

Pools

Profiles

Jobs

Physical Device

Troubleshooting

Administration

unlabeled

unlabeled

VFIScalableTest

General Configuration

Storage Profile Summary - Storage Profile Details

Storage Profile Details - VFIPProfile on Storage System pnQAIntracsr-STek2540

Details Related Information

Details

* Name: MURProfile

Description: MUR Profile

Status: Not In Use

Type: User

* RAID Level: RAID 5

* Segment Size: 256 KB

Read Ahead Enabled: ☒

* No. of Disks: Variable

* Disk Type: ANY

Back to top

Related Information

Storage Pools (0)

Volumes (0)

Back to top

Save Reset

* Indicates required field

381299

Sun StorageTek™ Common Array Manager

Create New Volume

Steps: Help

Step 1: Select Storage Characteristics

Storage characteristics are defined by the storage pool. Select from the available pools or create a new one.

☒ Create New Storage Pool
☐ Use Existing Storage Pool

Name	Maximum Volume Size	Storage Profile
Default	15.814 TB	Default
Default1	15.814 TB	Default1
Default2	5.998 TB	Default2
MUR_Pool	1.635 TB	Vof8
RAID0-512KB-ReadAhead	19.086 TB	RAID0-512KB-ReadAhead
RAID5-512KB-ReadAhead	15.814 TB	RAID5-512KB-ReadAhead
SP1	1.635 TB	SP1Prof

Previous Next Cancel

APPLICATIONS | VERSION

User: Administrator Server: PNQAHP0V1

Sun StorageTek™ Common Array Manager

Storage Pool Summary - New Storage Pool

Create New Storage Pool on Storage System pnQAIntracsr-STek2540

New Storage Pool

* Name: MURPool

Description: Storage Pool for MUR

* Storage Profile: MURProfile

OK Cancel

* Indicates required field

APPLICATIONS | VERSION

User: Administrator Server: PNQAHP0V1

Sun StorageTek™ Common Array Manager

Storage Pool Summary on Storage System pnQAIntracsr-STek2540

The selected objects were deleted.

Name	Storage Profile	State	Used Percent	Total Capacity	Configured Capacity	Available Capacity
CPPool	CPPProfile	In Use	0%	136.227 GB	150.000 MB	19.219 TB
Default	Default	Not In Use	0%	0.000 MB	0.000 MB	15.814 TB
Default1	Default1	Not In Use	0%	0.000 MB	0.000 MB	15.814 TB
Default2	Default2	Not In Use	0%	0.000 MB	0.000 MB	5.998 TB
MURPool	MURProfile	Not In Use	0%	0.000 MB	0.000 MB	15.814 TB
RAID0-512KB-ReadAhead	RAID0-512KB-ReadAhead	Not In Use	0%	136.227 GB	0.000 MB	19.219 TB
RAID5-512KB-ReadAhead	RAID5-512KB-ReadAhead	Not In Use	0%	0.000 MB	0.000 MB	15.814 TB
SP1	SP1Prof	Not In Use	0%	0.000 MB	0.000 MB	1.635 TB

Sun StorageTek™ Common Array Manager

Create New Volume

Steps Help

Step 1: Select Storage - Populate Pool

1. Select Storage Characteristics

➔ 1.1 Select Storage - Populate Pool

2. Configure Volume

3. Review Selections

Select the physical storage from which the new volume will be allocated.

Create new volume on:

- ☒ Storage Selected Automatically by CAM
- ☐ An Existing Virtual Disk with Available Capacity
- ☐ Currently Unassigned Disks (Create a New Virtual Disk)

Previous Next Cancel

381303

Sun StorageTek™ Common Array Manager

Create New Volume

Steps Help

Step 1: Select Storage Characteristics

➔ 1. Select Storage Characteristics

2. Configure Volume

3. Review Selections

Storage characteristics are defined by the storage pool. Select from the available pools or create a new one.

☐ Create New Storage Pool

☒ Use Existing Storage Pool

Name	Maximum Volume Size	Storage Profile
CPPool	13.088 TB	CPPProfile
Default	15.814 TB	Default
Default1	15.814 TB	Default1
Default2	5.998 TB	Default2
MURPool	15.814 TB	MURProfile
RAID0-512K-ReadAhead	19.088 TB	RAID0-512K-ReadAhead
RAID5-512K-ReadAhead	15.814 TB	RAID5-512K-ReadAhead
SP1	1.635 TB	SP1Prof

Previous Next Cancel

381304

Sun StorageTek™ Common Array Manager

Create New Volume

Steps Help

Step 2: Configure Volume

➔ 2. Configure Volume

1. Select Storage Characteristics

1.1 Select Storage - Populate Pool

3. Review Selections

Specify the quantity, name and size of the volume(s) to create.

* Volume Name: InputVol

Name can be up to 30 characters long and contain the following characters: "A-Z", "a-z", "*" and ".". Spaces are not allowed.

Number to Create: 1

If creating more than 1 volume, a unique number will be appended to the specified volume name.

Size: ☒ Fill One Virtual Disk 15.814 TB

☒ Specify Size

5 TB

Controller: Any

Previous Next Cancel

381305

Sun StorageTek™ Common Array Manager

Create New Volume

Steps Help

Step 3: Review Selections

➔ 3. Review Selections

1. Select Storage Characteristics

1.1 Select Storage - Populate Pool

2. Configure Volume

2.1 Select Mapping

2.2 Select a Host or Host Group

Review and confirm your selections.

No Hot-Spares
There are currently no hot-spares.

Volume Name: InputVol

Volume Capacity: 5 TB

Pool Name: MURPool

Storage Profile: MURProfile

Storage Characteristics: RAID-5, Segment Size: 256 KB, Read-Ahead: true, Number Of Disks: Variable, Disk Type: ANY

Virtual Disk Name: Chosen automatically by CAM

Controller: Any

Mapping To: Default Storage Domain - LUN:3

Previous Finish Cancel

381306

Sun StorageTek™ Common Array Manager

Create New Volume

Steps Help Step 2.1: Select Mapping

1. Select Storage Characteristics

1.1 Select Storage - Populate Pool

2. Configure Volume

2.1 Select Mapping

3. Review Selections

Select whether you want to map this volume now or later.

☐ Do Not Map at This Time

☒ Map to an Existing Host/Host Group or the Default Storage Domain

☐ Map to a New Host

Host Name:

Member of Host Group:

Initiator (Fill in All Fields or None)

Initiator Name:

Unique Identifier: ☒ Enter New Unique Identifier:

☐ Select Discovered Unique Identifier:

Host Type:

Previous Next Cancel

381307

Sun StorageTek™ Common Array Manager

Create New Volume

Steps Help Step 3: Review Selections

1. Select Storage Characteristics

1.1 Select Storage - Populate Pool

2. Configure Volume

2.1 Select Mapping

2.2 Select a Host or Host Group

3. Review Selections

Review and confirm your selections.

No Hot-Spares
There are currently no hot-spares.

Volume Name: ArchiveVol

Volume Capacity: 6 TB

Pool Name: MURPool

Storage Profile: MURProfile

Storage Characteristics: RAID 5, Segment Size: 256 KB, Read-Ahead: true, Number Of Disks: Variable, Disk Type: ANY

Virtual Disk Name: Chosen automatically by CAM

Controller: Any

Mapping To: Default Storage Domain - LUN4

Previous Finish Cancel

381308

APPLICATIONS VERSION

User: Administrator Server: PNOAH-POV1

Sun StorageTek™ Common Array Manager

REFRESH SEARCH SERVICE ADVISOR LOGS OUT HELP

Current License: 1

Last Update: Jun 30, 2013 7:23:26 PM IST

Current Alarms: 0 32 0 0

java

Alarms

Storage Systems

clusters5_6

pnDevIntracn-STe

pnQAIntracn-STe

Alarms

Volumes

Mappings

Snapshots

Virtual Disks

Pools

Profiles

Jobs

Physical Device

Troubleshooting

Administration

unlabeled

unlabeled

VFIScaleTest

General Configuration

Volume Summary on Storage System pnQAIntracn-STeK2540

Volumes (8)

New... Map... Delete... View Performance Statistics... Filter: All Items

Name	State	Condition	Type	Virtual Disk	Pool	Capacity	WWN
ArchiveVol	Mapped	Optimal	Standard	3	MURPool	6,000 TB	60 0A 0B 80 00 67 49 E3 00 00 5B 93 51 CF E2 17
CPVol_01	Mapped	Optimal	Standard	1	CPPool	50,000 MB	60 0A 0B 80 00 67 49 E3 00 00 5B 9F 51 CE 39 3E
CPVol_02	Mapped	Optimal	Standard	1	CPPool	50,000 MB	60 0A 0B 80 00 67 49 E7 00 00 35 C7 51 CE 37 F5
CPVol_03	Mapped	Optimal	Standard	1	CPPool	50,000 MB	60 0A 0B 80 00 67 49 E3 00 00 5B 91 51 CE 39 72
InqurVol	Mapped	Optimal	Standard	2	MURPool	6,000 TB	60 0A 0B 80 00 67 49 E7 00 00 35 C9 51 CF E0 17

New... Map... Delete... View Performance Statistics...

381309

Virtual Disk Summary on Storage System pnQAIntracsr-STek2540

To view details for a virtual disk, click on its name below.

Name	Status	State	Disk Type	RAID Level	Used Percent	Total Capacity	Configured Capacity	Maximum Volume Size
1	Optimal	Ready	SAS	0	0%	136,227 GB	150,000 MB	136,081 GB
2	Optimal	Initializing...	SAS	5	91%	5,453 TB	5,000 TB	464,116 GB
3	Optimal	Initializing...	SAS	5	91%	5,453 TB	5,000 TB	556,939 GB

381310

Testing the Port Connections between StorageTek and UCS Machines

After fiber channel cabling is done, go to each UCS node and check if the port name of the HBA port appears as an initiator on the StorageTek.

Important: If any of the volumes is not to be mapped to specific UCS host or group of hosts, there is no need to create and map these initiators. Hence, in case where volumes have to be shared across all the nodes, do not map the initiators and the volumes to specific hosts or group of hosts.

1. Enlist the port names on the UCS nodes as shown below. This sample has all the two HBA ports connected.

```
cat /sys/class/fc_host/host5/port_name 0x21000024ff31fe86
```

```
cat /sys/class/fc_host/host6/port_name 0x21000024ff31fe87
```

2. Check if these ports appear on the StorageTek as initiator by clicking **Initiators -> New**.

Initiator Summary on Storage System pnQAIntracsr-STek2540

Name	Host	Host Type	Unique Identifier
pnqaextapps460-2-p1	pnqaextapps460-2	Linux	21:00:00:24:FF:31:FE:64
pnqaextapps460-2-p2	pnqaextapps460-2-p2	Linux	21:00:00:24:FF:31:FE:65
pnqaextapps460-4-p1	pnqaextapps460-4	Linux	21:00:00:24:FF:31:FE:6C
pnqaextapps460-5-p1	pnqaextapps460-5	Linux	21:00:00:24:FF:3E:9E:7A
pnqaextapps460-7-p1	pnqaextapps460-7	Linux	21:00:00:24:FF:3E:A0:4A
pnqaextapps460-7-p2	pnqaextapps460-7-p22	Linux	21:00:00:24:FF:3E:A0:4B

381317

3. Check if the port is available.

Create New Initiator

Steps: 1. Specify Initiator Properties, 2. Review, 3. Results

Step 1: Specify Initiator Properties

Type Name, Unique Identifier, select the Host Type. Select the host in which you want the initiator to be created or select 'Create New Host' to create a new host.

* Initiator Name:

Up to 30 characters, can contain 'A-Z', 'a-z', '0-9', '-' and '_'. Spaces are not allowed.

* Unique Identifier: ☐ Enter New Unique Identifier:
☒ Select Discovered Unique Identifier:

21:00:00:24:FF:31:B1:F6
21:00:00:24:FF:31:91:F7
21:00:00:24:FF:31:FE:98
21:00:00:24:FF:31:FE:87
21:00:00:24:FF:31:FE:A0

* Host Type:

* Host: ☐ Select Existing Host:

21:00:00:24:FF:31:FE:98
21:00:00:24:FF:31:FE:87
21:00:00:24:FF:31:FE:A0

☐ Create New Host:

Previous Next Cancel

If the ports do not appear as initiator, please check the fiber cables and see if the green LEDs are blinking on the connected ends of the fiber cable.

Configuring Alerts on StorageTek

It is advised to configure the monitoring events and alerts on the StorageTek for quicker fault alert and analysis.

Configuring SNMP Notifications

To configure SNMP notifications, click **SNMP -> Add SNMP Notification**.

APPLICATIONS VERSION

User: Administrator Server: PNQAHP0V1

REFRESH SEARCH SERVICE ADVISOR LOG OUT HELP

Current Logins: 1
Last Update: May 24, 2013 7:33:32 AM MDT
Current Alarms: 0 18 5 0

Alarms

- Storage Systems
 - clustgss5_6
 - pnDevIntracsr-STek2540
 - pnQAIntracsr-STek2540
 - unlabeled
 - unlabeled
 - unlabeled
- General Configuration
 - Auto Service Request
 - General Health Monitoring
- Notification
 - Email
 - Email Filters
 - SNMP**
 - User Management
 - Activity Log

SNMP Notification > Add SNMP Notification

Add SNMP Notification

OK Reset Cancel

* Indicates required field

SNMP Properties

* IP Name/Address:

* Port:


Minimum Alert Level:

Send Configuration Change Events: ☐ Yes ☒ No

OK Reset Cancel

Configuring Email Alerts

To configure Email notifications, click **Email -> Add Email Notification**.



Incoming Data Partition

The incoming data partition should hold a data of at least 3 days so that in worst cases, pending data can stay on the disk. Size this disk partition as per the sizing sheet. Below examples consider the total throughput from single ASR node.

For example:

- As per the sizing sheet, for 40 Gbps throughput total disk size required would be equal to 3 days of archival storage and so 15+1 disks (considering RAID 5) of 600 GB and size of 9.3 TB.
- As per the sizing sheet, for 20 Gbps throughput total disk size required would be equal to 3 days of archival storage and so 8+1 disks (considering RAID 5) of 600 GB and size of 4.8 TB.

Archival Data Partition

Decide the duration of the files retention period for archiving and size the archival volume accordingly.

For example:

- As per the sizing sheet, for 20 Gbps of total throughput total disk size required for the 7 days of archival period are 18 + 1 disks (considering RAID 5) of 600 GB and so the size of 10.5 TB.
- As per the sizing sheet, for 40 Gbps of total throughput total disk size required for the 7 days of archival period are 35 + 1 disks (considering RAID 5) of 600 GB and so the size of 20.5 TB.

Important: The CAM software can be downloaded from Oracle Support site or it can be ordered along with the StorageTek hardware. Please use the CAM software recommended by Oracle compatible to the purchased StorageTek.

Important: Configure the alerts as Email and SNMP events of the StorageTek as described in the [Installing the Management Software \(CAM\)](#) section of this guide.

Configuring Multipaths on RDP Hosts

The volumes of the StorageTek shared appear on the RDP hosts through the paths. There could be multiple paths available on the machine. Even if the paths are redundant they contribute for a better load distribution and performance.

Configuring Multipath Service

Device Mapper Multipath (DM-MP) allows nodes to route I/O over multiple paths to a storage controller. A path refers to the connection from an HBA port to a storage controller port. As paths fail and new paths come up, DM-MP reroutes the I/O over the available paths.

- The devices in `/dev/mapper` are created early in the boot process. Use these devices to access the multipathed devices, for example when creating logical volumes.
- The devices in `/dev/mpath` are provided as a convenience so that all multipathed devices can be seen in one directory. These devices are created by the udev device manager and may not be available on startup when the system needs to access them. Do not use these devices for creating logical volumes or file systems.
- Any devices of the form `/dev/dm -n` are for internal use only and should never be used.

The following section explains how to configure the multipath on the system after mapping volume of StorageTek. Follow the procedure for each node of the RDP on each UCS.

1. Check if the device multipath rpm is installed.

```
[root@murl ~]# rpm -qa | grep multi
device-mapper-multipath-0.4.7-34.el5
[root@murl ~]#
```

2. The blacklist section of the multipath configuration file specifies the devices that will not be used when the system configures multipath devices. Devices that are blacklisted will not be grouped into a multipath device. By default, all devices are blacklisted, since the following lines appear in the initial configuration file.

```
blacklist {
    devnode "*"
}
```

To enable multipathing on all of the devices that are supported by default, comment out above lines from the *multipath.conf* file in the */etc/* directory.

3. Run the command “**fdisk -l**” to see all the disk partitions and to list which disks are shown as SAN connected disk.
4. For each of the SAN connected block devices, run the following command and note which block devices show the same value of UUID.
5. To identify the local disks and SAN disks compare the snapshot already taken with the command “**fdisk -l**”.

```
[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdb
3600605b002e6aeb0190918a72dc76fad
[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdc
```

```

3600a0b80006749870000322f5154598b

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdd
3600a0b80006749870000322f5154598b

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sde
3600a0b80006749e30000588e515d14a0

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdf
3600a0b80006749e30000588e515d14a0

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdk
3600a0b80006749e30000588e515d14a0

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdl
3600a0b80006749e30000588e515d14a0

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdm
3600a0b80006749870000322f5154598b

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdn
3600a0b80006749870000322f5154598b

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdq
3600a0b80006749870000322f5154598b

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdr
3600a0b80006749870000322f5154598b

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sds
3600a0b80006749e30000588e515d14a0

[root@pnextappsucs460-1 ~]# scsi_id -g -u -s /block/sdt
3600a0b80006749e30000588e515d14a0

```

6. Edit the `/etc/multipath.conf` file and add the aliases for the SAN connected disk.

```

multipaths {
    multipath {
        wwid                3600a0b80006749870000322f5154598b
        alias                archive
    }
}

```

```

multipath {
    wwid                      3600a0b80006749e30000588e515d14a0
    alias                     input
}

```

7. Execute the following commands:

```

#chkconfig multipathd on
#service multipathd restart

```

8. Run the command “**fdisk -l**” again and see if DM shows up. If not, then reboot the system and check again.

```

Disk /dev/dm-0 (Sun disk label): 255 heads, 189 sectors, 21759 cylinders
Units = cylinders of 48195 * 512 bytes

```

Device	Flag	Start	End	Blocks	Id	System
/dev/dm-0p3	u	0	21759	524337502+	5	Whole disk
/dev/dm-0p8	u	0	21759	524337502+	f	Unknown

```

Disk /dev/dm-1: 20 MB, 20971520 bytes
255 heads, 63 sectors/track, 2 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

```

9. Execute the command “**ls -al /dev/mpath**”. The defined alias should show up here. If not, then reboot the system and check again. Also, you should see dm-0 and dm-1.

10. Execute the command “**multipath -v3 -d**” to verify if DM is working fine.

11. Run **multipath -ll** command to check if all the paths are seen and all of them are in active/ready condition. Note that the paths should not be in fault state.

```

mpath1 (3600605b002e6aeb0190918a72dc76fad) dm-0 LSI,MR9260-8i
[size=5.5T][features=0][hwhandler=0][rw]
\_ round-robin 0 [prio=1][active]
\_ 0:2:1:0 sdb 8:16 [active][ready]
input (3600a0b80006749e30000588e515d14a0) dm-2 SUN,LCSM100_F
[size=5.5T][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
\_ round-robin 0 [prio=400][active]

```

```

\_ 5:0:1:1 sdj 8:144 [active][ready]
\_ 6:0:1:1 sdm 8:192 [active][ready]
\_ 5:0:3:1 sdr 65:16 [active][ready]
\_ 6:0:3:1 sdt 65:48 [active][ready]

```

Configuring Veritas

Veritas Cluster File system allows having disk partition shared across multiple hosts. Below are the steps involved.

Hardware Setup for Veritas Cluster File System

All the RDP hosts which are part of the cluster should be in the same VLAN. Additionally as per recommendation of VCS there should be interconnected network links to be created for heartbeat mechanism as described in the [Heartbeat Link between RDP Hosts](#) section of this guide.

Installing Veritas Cluster File System

Make available Veritas installer on either of the RDP nodes. Please note that from the available packages, install **Veritas Storage Foundation Cluster File System (SFCFS)**.

Follow the instructions as per VCS installer for this package.

Below is a sample trace file:

```

Storage Foundation and High Availability Solutions 5.1 SP1 Install Program

Copyright (c) 2010 Symantec Corporation. All rights reserved. Symantec, the
Symantec Logo are trademarks or registered trademarks of Symantec Corporation or
its affiliates in the U.S. and

other countries. Other names may be trademarks of their respective owners.

The Licensed Software and Documentation are deemed to be "commercial computer
software" and "commercial computer software documentation" as defined in FAR
Sections 12.212 and DFARS Section

227.7202.

Logs are being written to /var/tmp/installer-201306302040lvT while installer is
in progress.

Storage Foundation
and High Availability Solutions 5.1 SP1 Install
Program

```

```

Symantec Product                                Version Installed  Licensed
=====
Symantec Licensing Utilities (VRTSvlic) are not installed due to which products
and licenses are not discovered.

Use the menu below to continue.

```

Task Menu:

```

P) Perform a Pre-Installation Check      I) Install a Product
C) Configure an Installed Product        G) Upgrade a Product
O) Perform a Post-Installation Check      U) Uninstall a Product
L) License a Product                    S) Start a Product
D) View Product Descriptions             X) Stop a Product
R) View Product Requirements             ?) Help

```

Enter a Task: [P,I,C,G,O,U,L,S,D,X,R,?] I

```

Storage Foundation
and High Availability Solutions 5.1 SP1 Install
Program

```

- 1) Veritas Dynamic Multi-Pathing (DMP)
- 2) Veritas Cluster Server (VCS)
- 3) Veritas Storage Foundation (SF)
- 4) Veritas Storage Foundation and High Availability (SFHA)
- 5) Veritas Storage Foundation Cluster File System (SFCFS)
- 6) Veritas Storage Foundation Cluster File System/HA (SFCFSHA)
- 7) Veritas Storage Foundation for Oracle RAC (SFRAC)

- 8) Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFS RAC)
- 9) Symantec VirtualStore (SVS)
- 10) Symantec Product Authentication Services (AT)
- b) Back to previous menu

Select a product to install: [1-10,b,q] 5

Do you agree with the terms of the End User License Agreement as specified in the storage_foundation_cluster_file_system/EULA/en/EULA_CFS_Ux_5.1SP1.pdf file present on media? [y,n,q,?] y

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install Program

- 1) Install minimal required Veritas Storage Foundation Cluster File System rpms - 379 MB required
- 2) Install recommended Veritas Storage Foundation Cluster File System rpms - 584 MB required
- 3) Install all Veritas Storage Foundation Cluster File System rpms - 620 MB required
- 4) Display rpms to be installed for each option

Select the rpms to be installed on all systems? [1-4,q,?] (2) 3

Enter the 64 bit RHEL5 system names separated by spaces: [q,?] (pnqaextapps460-7 pnextappsucs460-1)

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install Program

pnqaextapps460-7 pnextappsucs460-1

Logs are being written to /var/tmp/installer-201306302040lvT while installer is in progress

Verifying systems: 0%

Estimated time
remaining:

0 of 8

Checking system communication -\|/-\|/ Verifying systems:
12%

Estimated time remaining:
0:55

8

1 of

Checking system communication
.....
..... Done

Verifying systems: 12%

Estimated time remaining:
0:55

8

1 of

```

Checking system communication
.....
..... Done

```

```

Checking release compatibility -\|    Verifying systems:
25%

```

```

Estimated time remaining:

```

```

0:30

```

```

2 of

```

```

8

```

```

Checking system communication
.....
..... Done

```

```

Checking release compatibility
.....
..... Done

```

```

Verifying systems: 25%

```

```

Estimated time remaining:

```

```

0:35

```

```

2 of

```

```

8

```

```

Checking system communication
.....
..... Done

```

```

Checking release compatibility
.....
..... Done

```

```

Checking installed product /-\\/-\\/-\    Verifying systems:
37%

```

Estimated time remaining:

0:35

3 of

8

Checking system communication

.....
..... Done

Checking release compatibility

.....
..... Done

Checking installed product

.....
..... Done

Verifying systems:

37%

.....
.....

Estimated time remaining:

0:35

3 of

8

Checking system communication

.....
..... Done

Checking release compatibility

.....
..... Done

Checking installed product

.....
..... Done

Checking prerequisite patches and rpms || Verifying systems:

50%

.....
.....

Estimated time remaining:
0:25

4 of

8

Checking system communication

.....
..... Done

Checking release compatibility

.....
..... Done

Checking installed product

.....
..... Done

Checking prerequisite patches and rpms

.....
..... Done

Verifying systems:
50%

.....
.....

Estimated time remaining:
0:25

4 of

8

Checking system communication

.....
..... Done

Checking release compatibility

.....
..... Done

Checking installed product

.....
..... Done

```
Checking prerequisite patches and rpms
.....
..... Done
```

```
Checking platform version -\    Verifying systems:
62%
```

```
_____
```

```
Estimated time remaining:
```

```
0:15
```

```
5 of
```

```
8
```

```
Checking system communication
.....
..... Done
```

```
Checking release compatibility
.....
..... Done
```

```
Checking installed product
.....
..... Done
```

```
Checking prerequisite patches and rpms
.....
..... Done
```

```
Checking platform version
.....
.....
Done
```

```
Verifying systems:
```

```
62%
```

```
_____
```

```
Estimated time remaining:
```

```
0:15
```

```
5 of
```

```
8
```

```
Checking system communication
..... Done

Checking release compatibility
..... Done

Checking installed product
..... Done

Checking prerequisite patches and rpms
..... Done

Checking platform version
.....
Done

Checking file space |/-\    Verifying systems:
75%
```

```
-----

Estimated time remaining:
0:10
8                                     6 of
```

```
Checking system communication
..... Done

Checking release compatibility
..... Done

Checking installed product
..... Done

Checking prerequisite patches and rpms
..... Done
```

```
Checking platform version
.....
.....
Done
```

```
Checking file space
.....
.....
.. Done
```

```
Verifying systems:
75%
```

```
Estimated time remaining:
0:10
```

6 of

8

```
Checking system communication
.....
..... Done
```

```
Checking release compatibility
.....
..... Done
```

```
Checking installed product
.....
..... Done
```

```
Checking prerequisite patches and rpms
.....
..... Done
```

```
Checking platform version
.....
.....
Done
```

```
Checking file space
.....
.....
.. Done
```

```
Performing product license checks |/ Verifying systems:
87%
```

Estimated time remaining:
0:04

7 of

8

Checking system communication
.....
..... Done

Checking release compatibility
.....
..... Done

Checking installed product
.....
..... Done

Checking prerequisite patches and rpms
.....
..... Done

Checking platform version
.....
.....
Done

Checking file space
.....
.....
.. Done

Performing product license checks
.....
..... Done

Verifying systems:
87%

Estimated time remaining:
0:04

7 of

8

```

Checking system communication
.....
..... Done

Checking release compatibility
.....
..... Done

Checking installed product
.....
..... Done

Checking prerequisite patches and rpms
.....
..... Done

Checking platform version
.....
.....
Done

Checking file space
.....
.....
.. Done

Performing product license checks
.....
..... Done

Performing product prechecks -\|    Verifying systems:
100%
```

```

Estimated time remaining:
0:00
```

8 of

8

```

Checking system communication
.....
..... Done
```

```

    Checking release compatibility
    .....
    ..... Done

    Checking installed product
    .....
    ..... Done

    Checking prerequisite patches and rpms
    .....
    ..... Done

    Checking platform version
    .....
    .....
Done

    Checking file space
    .....
    .....
.. Done

    Performing product license checks
    .....
    ..... Done

    Performing product prechecks
    .....
    ..... Done

System verification checks completed successfully

The following warnings were discovered on the systems:

Nodes have difference in clock by more than 5 sec

Do you want to continue? [y,n,q] (y) y

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

```


The following Veritas Storage Foundation Cluster File System rpms will be installed on all systems:

Rpm	Rpm Description
VRTSvlic	Veritas Licensing
VRTSperl	Veritas Perl 5.10.0 Redistribution
VRTSspt	Veritas Software Support Tools by Symantec
VRTSvxvm	Veritas Volume Manager Binaries
VRTSaslapm	Volume Manager - ASL/APM
VRTSob	Veritas Enterprise Administrator Service by Symantec
VRTSlvconv	Veritas Linux LVM to VxVM Converter
VRTSsfmh	Veritas Storage Foundation Managed Host by Symantec
VRTSvxfs	Veritas File System
VRTSfssdk	Veritas File System Software Developer Kit
VRTSatClient	Symantec Product Authentication Service Client
VRTSatServer	Symantec Product Authentication Service
VRTSvxfen	Veritas I/O Fencing by Symantec
VRTSamf	Veritas Asynchronous Monitoring Framework by Symantec
VRTSvc	Veritas Cluster Server
VRTScps	Veritas Cluster Server - Coordinated Point Server
VRTSvcasg	Veritas Cluster Server Bundled Agents by Symantec
VRTSvcldr	Veritas Cluster Server Disk Reservation Modules
VRTSvcsea	Veritas Cluster Server Enterprise Agents by Symantec
VRTSdbed	Veritas Storage Foundation Databases
VRTSglm	Veritas Group Lock Manager
VRTScavf	Veritas Cluster Server Agents for Cluster File System
VRTSgms	Veritas Group Messaging Services
VRTSodm	Veritas Oracle Disk Manager

The following Veritas Storage Foundation Cluster File System rpms will be installed on pnqaextapps460-7:

Rpm	Rpm Description
VRTSl1t	Veritas Low Latency Transport
VRTSgab	Veritas Group Membership and Atomic Broadcast

Press [Enter] to continue:

Veritas Storage
Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

Logs are being written to /var/tmp/installer-201306302040lvT while installer is
in progress

Installing SFCFS: 0%

Estimated time
remaining:

0 of 29

Performing SFCFS preinstall tasks /-\\| Installing SFCFS: 3%

```

Estimated time remaining:
1:50
29
1 of

```

```

Performing SFCFS preinstall tasks
.....
..... Done

Installing SFCFS: 3%
_____
_____

```

```

Estimated time remaining:
1:50
29
1 of

```

```

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm /-\\ Installing SFCFS:
6%
_____
_____
_____

```

```

Estimated time remaining:
1:30
29
2 of

```

```

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....

```

```
.....
Done

    Installing SFCFS:
6%      _____
_____
_____
```

```

    Estimated time remaining:
1:30
29
2 of
```

```

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

    Installing VRTSperl rpm |/-\|/-\    Installing SFCFS:
10%      _____
_____
_____
```

```

    Estimated time remaining:
2:10
29
3 of
```

```

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

    Installing VRTSperl rpm
.....
```

```
.....
Done

    Installing SFCFS:
10%      _____
_____
```

```

    Estimated time remaining:
2:10
                                                    3 of
29
```

```

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

    Installing VRTSperl rpm
.....
.....
Done

    Installing VRTSspt rpm |/-    Installing SFCFS: 13%
_____
_____
```

```

    Estimated time remaining:
1:50
                                                    4 of
29
```

```

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
```

```

.....
Done

    Installing VRTSperl rpm
.....
.....
Done

    Installing VRTSspt rpm
.....
.....
Done

    Installing SFCFS: 13%
.....
.....

    Estimated time remaining:
1:50

29
4 of

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

    Installing VRTSperl rpm
.....
.....
Done

    Installing VRTSspt rpm
.....
.....
Done

    Installing VRTSvxvm rpm \|-\\|-\\|-\\|-\\| Installing SFCFS:
17%
.....
.....

```

Estimated time remaining:
3:02

5 of

29

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing SFCFS:
17%
.....
.....
.....

Estimated time remaining:
3:02

5 of

29

Performing SFCFS preinstall tasks
.....
..... Done

```

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm -\|/      Installing SFCFS:
20%
_____
_____

```

```

Estimated time remaining:
2:40

```

6 of

29

```

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

```



```
Installing VRTSvxvm rpm
.....
.....
Done
```

```
Installing VRTSaslapm rpm
.....
.....
Done
```

```
Installing SFCFS: 20%
```

```
Estimated time remaining:
2:40
29
```

6 of

```
Performing SFCFS preinstall tasks
.....
..... Done
```

```
Installing VRTSvlic rpm
.....
.....
Done
```

```
Installing VRTSperl rpm
.....
.....
Done
```

```
Installing VRTSspt rpm
.....
.....
Done
```

```
Installing VRTSvxvm rpm
.....
.....
Done
```

```
Installing VRTSaslapm rpm
.....
.....
Done
```

```

Installing VRTSob rpm -\|/-      Installing SFCFS:
24%

```

```

Estimated time remaining:
2:25

```

7 of

29

```

Performing SFCFS preinstall tasks
.....
..... Done

```

```

Installing VRTSvlic rpm
.....
..... Done

```

```

Installing VRTSperl rpm
.....
..... Done

```

```

Installing VRTSspt rpm
.....
..... Done

```

```

Installing VRTSvxvm rpm
.....
..... Done

```

```

Installing VRTSaslapm rpm
.....
..... Done

```

```

Installing VRTSob rpm
.....
..... Done

```

```

Installing SFCFS: 24%

```

Estimated time remaining:
2:25
29

7 of

```

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSsvmconv rpm \\\ Installing SFCFS:
27%
_____
_____
_____

```

```
Estimated time remaining:
2:10
8 of
29

Performing SFCFS preinstall tasks
..... Done

Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSsvmconv rpm
..... Done

Installing SFCFS:
27%
_____
_____
_____
```

Estimated time remaining:
2:10

8 of

29

```

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSsvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm -\|/-\|/-\|/- Installing SFCFS:
31%

```

Estimated time remaining:

2:20

9 of

29

```

    Performing SFCFS preinstall tasks
    .....
    ..... Done

```

```

    Installing VRTSvlic rpm
    .....
    .....
    Done

```

```

    Installing VRTSperl rpm
    .....
    .....
    Done

```

```

    Installing VRTSspt rpm
    .....
    .....
    Done

```

```

    Installing VRTSvxvm rpm
    .....
    .....
    Done

```

```

    Installing VRTSaslapm rpm
    .....
    .....
    Done

```

```

    Installing VRTSob rpm
    .....
    .....
    Done

```

```

    Installing VRTSslmconv rpm
    .....
    ..... Done

```

```

    Installing VRTSvxfs rpm
    .....
    .....
    Done

```

Installing SFCFS: 31%

Estimated time remaining:

2:20

9 of

29

Performing SFCFS preinstall tasks

.....
..... Done

Installing VRTSvlic rpm

.....
.....
Done

Installing VRTSperl rpm

.....
.....
Done

Installing VRTSspt rpm

.....
.....
Done

Installing VRTSvxvm rpm

.....
.....
Done

Installing VRTSaslapm rpm

.....
.....
Done

Installing VRTSob rpm

.....
.....
Done

Installing VRTSsvmconv rpm

.....
..... Done

Installing VRTSvxfs rpm

.....
.....
Done

```

Installing VRTSfssdk rpm \|- Installing SFCFS:
34%
_____
_____

Estimated time remaining:
2:05
_____
_____

10 of
29

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
..... Done

Installing VRTSperl rpm
.....
..... Done

Installing VRTSspt rpm
.....
..... Done

Installing VRTSvxvm rpm
.....
..... Done

Installing VRTSaslapm rpm
.....
..... Done

Installing VRTSob rpm
.....
..... Done

Installing VRTSsvmconv rpm
.....
..... Done

```



```
Installing VRTSvxfs rpm
.....
.....
Done
```

```
Installing VRTSfssdk rpm
.....
.....
Done
```

```
Installing SFCFS:
34%
_____
_____
_____
```

```
Estimated time remaining:
2:05
29
10 of
```

```
Performing SFCFS preinstall tasks
.....
..... Done
```

```
Installing VRTSvlic rpm
.....
.....
Done
```

```
Installing VRTSperl rpm
.....
.....
Done
```

```
Installing VRTSspt rpm
.....
.....
Done
```

```
Installing VRTSvxvm rpm
.....
.....
Done
```

```
Installing VRTSaslapm rpm
.....
.....
Done
```

```

Installing VRTSob rpm
.....
Done

Installing VRTSsvmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm \||/ Installing SFCFS:
37%
_____
_____

```

Estimated time remaining:

1:50

11 of

29

```

Performing SFCFS preinstall tasks
..... Done

Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

```

```

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSsvmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing SFCFS:
37%

```

```

Estimated time remaining:
1:50

29

```

11 of

```

Performing SFCFS preinstall tasks
..... Done

```

```

Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm -\|/-\|/-\ Installing SFCFS:
41%
_____
_____
_____

```

Estimated time remaining:
1:50

12 of

29

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSlvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....

```

.....
Done

    Installing VRTSatClient rpm
.....
..... Done

    Installing VRTSatServer rpm
.....
..... Done

    Installing SFCFS:
41%
_____
_____

    Estimated time remaining:
1:50

29
12 of

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

    Installing VRTSperl rpm
.....
.....
Done

    Installing VRTSspt rpm
.....
.....
Done

    Installing VRTSvxvm rpm
.....
.....
Done

    Installing VRTSaslapm rpm
.....

```

```

.....
Done

    Installing VRTSob rpm
.....
.....
Done

    Installing VRTSsvmconv rpm
.....
..... Done

    Installing VRTSvxfs rpm
.....
.....
Done

    Installing VRTSfssdk rpm
.....
.....
Done

    Installing VRTSatClient rpm
.....
..... Done

    Installing VRTSatServer rpm
.....
..... Done

    Installing VRTSl1t rpm |/-    Installing SFCFS:
44%

```

```

.....

    Estimated time remaining:
1:40

29
13 of

```

```

.....

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

```

```
Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing SFCFS:
44%
```

Estimated time remaining:
1:40
29

13 of

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSsvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....

```
.....
Done

    Installing VRTSfssdk rpm
    .....
    .....
Done

    Installing VRTSatClient rpm
    .....
    ..... Done

    Installing VRTSatServer rpm
    .....
    ..... Done

    Installing VRTSl1t rpm
    .....
    .....
Done

    Installing VRTSgab rpm \\\      Installing SFCFS:
48%
```

```
    Estimated time remaining:
1:30

29
```

14 of

```
    Performing SFCFS preinstall tasks
    .....
    ..... Done

    Installing VRTSvlic rpm
    .....
    .....
Done

    Installing VRTSperl rpm
    .....
    .....
Done

    Installing VRTSspt rpm
    .....
```

```
.....
Done

    Installing VRTSvxvm rpm
    .....
    .....
Done

    Installing VRTSaslapm rpm
    .....
    .....
Done

    Installing VRTSob rpm
    .....
    .....
Done

    Installing VRTSsvmconv rpm
    .....
    ..... Done

    Installing VRTSvxfs rpm
    .....
    .....
Done

    Installing VRTSfssdk rpm
    .....
    .....
Done

    Installing VRTSatClient rpm
    .....
    ..... Done

    Installing VRTSatServer rpm
    .....
    ..... Done

    Installing VRTSllt rpm
    .....
    .....
Done

    Installing VRTSgab rpm
    .....
    .....
Done

    Installing SFCFS:
48%
```

Estimated time remaining:
1:30

29

14 of

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....

```

.....
Done

    Installing VRTSfssdk rpm
    .....
    .....
Done

    Installing VRTSatClient rpm
    .....
    ..... Done

    Installing VRTSatServer rpm
    .....
    ..... Done

    Installing VRTSl1t rpm
    .....
    .....
Done

    Installing VRTSgab rpm
    .....
    .....
Done

    Installing VRTSvxfen rpm -\|    Installing SFCFS:
51%
_____
_____

    Estimated time remaining:
1:20

29
15 of

    Performing SFCFS preinstall tasks
    .....
    ..... Done

    Installing VRTSvlic rpm
    .....
    .....
Done

    Installing VRTSperl rpm
    .....
```

```

.....
Done

    Installing VRTSspt rpm
.....
Done

    Installing VRTSvxvm rpm
.....
Done

    Installing VRTSaslapm rpm
.....
Done

    Installing VRTSob rpm
.....
Done

    Installing VRTSslmconv rpm
..... Done

    Installing VRTSvxfs rpm
.....
Done

    Installing VRTSfssdk rpm
.....
Done

    Installing VRTSatClient rpm
..... Done

    Installing VRTSatServer rpm
..... Done

    Installing VRTSl1t rpm
.....
Done

    Installing VRTSgab rpm
.....
Done

```

```
Installing VRTSvxfen rpm
.....
.....
Done
```

```
Installing SFCFS:
51%
_____
_____
```

```
Estimated time remaining:
1:20
15 of
29
```

```
Performing SFCFS preinstall tasks
.....
..... Done
```

```
Installing VRTSvlic rpm
.....
.....
Done
```

```
Installing VRTSperl rpm
.....
.....
Done
```

```
Installing VRTSspt rpm
.....
.....
Done
```

```
Installing VRTSvxvm rpm
.....
.....
Done
```

```
Installing VRTSaslapm rpm
.....
.....
Done
```

```
Installing VRTSob rpm
.....
.....
Done
```

```
Installing VRTSsvmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
Done

Installing VRTSamf rpm /-\ Installing SFCFS:
55%
```

Estimated time remaining:
1:10

29


```
Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSlvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done
```

```

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSllt rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing SFCFS:
55%

```

```

Estimated time remaining:
1:10

29

```

16 of

```

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

```

```
Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSsvmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
```

```

.....
Done

    Installing VRTSamf rpm
.....
Done

    Installing VRTSvcS rpm |/-\|/-    Installing SFCFS:
58%
_____

Estimated time remaining:
1:05

29                                                    17 of

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
Done

    Installing VRTSperl rpm
.....
Done

    Installing VRTSspt rpm
.....
Done

    Installing VRTSvxvm rpm
.....
Done

    Installing VRTSaslapm rpm
.....
Done

```

```
Installing VRTSob rpm
.....
.....
Done

Installing VRTSsvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSllt rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcS rpm
.....
.....
Done

Installing SFCFS:
58%
```

Estimated time remaining:

1:05

17 of

29

Performing SFCFS preinstall tasks

.....
 Done

Installing VRTSvlic rpm

.....

 Done

Installing VRTSperl rpm

.....

 Done

Installing VRTSspt rpm

.....

 Done

Installing VRTSvxvm rpm

.....

 Done

Installing VRTSaslapm rpm

.....

 Done

Installing VRTSob rpm

.....

 Done

Installing VRTSslmconv rpm

.....
 Done

Installing VRTSvxfs rpm

.....

```

.....
Done

    Installing VRTSfssdk rpm
.....
.....
Done

    Installing VRTSatClient rpm
.....
..... Done

    Installing VRTSatServer rpm
.....
..... Done

    Installing VRTSl1t rpm
.....
.....
Done

    Installing VRTSgab rpm
.....
.....
Done

    Installing VRTSvxfen rpm
.....
.....
Done

    Installing VRTSamf rpm
.....
.....
Done

    Installing VRTSvcs rpm
.....
.....
Done

    Installing VRTScps rpm \|- Installing SFCFS:
62%

```

```

    Estimated time remaining:
1:02

```

29

18 of

```

    Performing SFCFS preinstall tasks
    .....
    ..... Done

    Installing VRTSvlic rpm
    .....
    .....
    Done

    Installing VRTSperl rpm
    .....
    .....
    Done

    Installing VRTSspt rpm
    .....
    .....
    Done

    Installing VRTSvxvm rpm
    .....
    .....
    Done

    Installing VRTSaslapm rpm
    .....
    .....
    Done

    Installing VRTSob rpm
    .....
    .....
    Done

    Installing VRTSslmconv rpm
    .....
    ..... Done

    Installing VRTSvxfs rpm
    .....
    .....
    Done

    Installing VRTSfssdk rpm
    .....
    .....
    Done

    Installing VRTSatClient rpm
    .....
    ..... Done

```



```

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSllt rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcS rpm
.....
.....
Done

Installing VRTScps rpm
.....
.....
Done

Installing SFCFS:
62%
_____
_____

```

Estimated time remaining:

1:02

18 of

29

```

Performing SFCFS preinstall tasks
.....
..... Done

```

```
Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
```

```

.....
Done

    Installing VRTSgab rpm
.....
.....
Done

    Installing VRTSvxfen rpm
.....
.....
Done

    Installing VRTSamf rpm
.....
.....
Done

    Installing VRTSvcs rpm
.....
.....
Done

    Installing VRTScps rpm
.....
.....
Done

    Installing VRTSvcsag rpm \|-\\|/-\\|/-      Installing SFCFS:
65%

```

```

.....
    Estimated time remaining:
1:00

```

19 of

29

```

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

```

```
Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
```

```

.....
Done

    Installing VRTSvxfen rpm
.....
.....
Done

    Installing VRTSamf rpm
.....
.....
Done

    Installing VRTSvcs rpm
.....
.....
Done

    Installing VRTScps rpm
.....
.....
Done

    Installing VRTSvcsag rpm
.....
.....
Done

    Installing SFCFS:
65%

```

```

    Estimated time remaining:
1:00

```

19 of

29

```

    Performing SFCFS preinstall tasks
.....
..... Done

    Installing VRTSvlic rpm
.....
.....
Done

```

```
Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
```

```

.....
Done

    Installing VRTSvxfen rpm
.....
Done

    Installing VRTSamf rpm
.....
Done

    Installing VRTSvcs rpm
.....
Done

    Installing VRTScps rpm
.....
Done

    Installing VRTSvcsag rpm
.....
Done

    Installing VRTSvcsdr rpm \\\    Installing SFCFS:
68%
_____

Estimated time remaining:
0:50

29                                     20 of

    Performing SFCFS preinstall tasks
..... Done

    Installing VRTSvlic rpm
.....
Done

```

```
Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
```



```

.....
Done

    Installing VRTSvxfen rpm
.....
Done

    Installing VRTSamf rpm
.....
Done

    Installing VRTSvcS rpm
.....
Done

    Installing VRTScps rpm
.....
Done

    Installing VRTSvcSag rpm
.....
Done

    Installing VRTSvcSdr rpm
.....
Done

    Installing SFCFS:
68%
_____
_____

```

```

    Estimated time remaining:
0:50

29

```

20 of

```

    Performing SFCFS preinstall tasks
.....
..... Done

```

```

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
.....

```

```

.....
Done

    Installing VRTSgab rpm
.....
Done

    Installing VRTSvxfen rpm
.....
Done

    Installing VRTSamf rpm
.....
Done

    Installing VRTSvcscs rpm
.....
Done

    Installing VRTScps rpm
.....
Done

    Installing VRTSvcscsag rpm
.....
Done

    Installing VRTSvcscsdr rpm
.....
Done

    Installing VRTSvcsea rpm -\|    Installing SFCFS:
72%

```

```

.....

    Estimated time remaining:
0:45

```

29

21 of

```
Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done
```

```
Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcscs rpm
.....
.....
Done

Installing VRTScps rpm
.....
.....
Done

Installing VRTSvcscsag rpm
.....
.....
Done

Installing VRTSvcscdr rpm
.....
.....
Done

Installing VRTSvcsea rpm
.....
.....
Done

Installing SFCFS:
72%
```

```
Estimated time remaining:
0:45
29
21 of

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done
```

```
Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSllt rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcscs rpm
.....
.....
Done

Installing VRTSvcscsag rpm
.....
.....
Done

Installing VRTSvcscsdr rpm
.....
.....
Done
```

```

Installing VRTSvcsea rpm
.....
Done

Installing VRTSdbed rpm /-\\/-      Installing SFCFS:
75%
_____

Estimated time remaining:
0:40

29                                     22 of

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

```



```
Installing VRTSsvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcsc rpm
.....
.....
Done

Installing VRTScps rpm
.....
.....
Done

Installing VRTSvcscsag rpm
.....
```

```
.....
Done

    Installing VRTSvcldr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
    .....
Done

    Installing VRTSdbed rpm
    .....
    .....
Done

    Installing SFCFS:
75%
_____

Estimated time remaining:
0:40

29                                     22 of

    Performing SFCFS preinstall tasks
    .....
    ..... Done

    Installing VRTSvlic rpm
    .....
    .....
Done

    Installing VRTSperl rpm
    .....
    .....
Done

    Installing VRTSspt rpm
    .....
    .....
Done
```

```
Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
Done

Installing VRTSamf rpm
.....
```

```
.....
Done

    Installing VRTSvcscs rpm
.....
Done

    Installing VRTScps rpm
.....
Done

    Installing VRTSvcscsag rpm
.....
Done

    Installing VRTSvcscsdr rpm
.....
Done

    Installing VRTSvcsea rpm
.....
Done

    Installing VRTSdbed rpm
.....
Done

    Installing VRTSglm rpm \|-      Installing SFCFS:
79%

.....

    Estimated time remaining:
0:30

29                                     23 of

    Performing SFCFS preinstall tasks
.....
..... Done
```

```
Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
```

```
.....
Done

    Installing VRTSgab rpm
    .....
    .....
Done

    Installing VRTSvxfen rpm
    .....
    .....
Done

    Installing VRTSamf rpm
    .....
    .....
Done

    Installing VRTSvcscs rpm
    .....
    .....
Done

    Installing VRTScps rpm
    .....
    .....
Done

    Installing VRTSvcscsag rpm
    .....
    .....
Done

    Installing VRTSvcscsdr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
    .....
Done

    Installing VRTSdbed rpm
    .....
    .....
Done

    Installing VRTSglm rpm
    .....
    .....
Done
```

Installing SFCFS:
79%

Estimated time remaining:
0:30

23 of

29

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSsvmconv rpm
.....
..... Done

```
Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
Done

Installing VRTSamf rpm
.....
Done

Installing VRTSvcs rpm
.....
Done

Installing VRTScps rpm
.....
Done

Installing VRTSvcsag rpm
.....
Done
```



```

Installing VRTSvcldr rpm
.....
Done

Installing VRTSvcsea rpm
.....
Done

Installing VRTSdbed rpm
.....
Done

Installing VRTSglm rpm
.....
Done

Installing VRTScavf rpm \|\| Installing SFCFS:
82%

```

```

Estimated time remaining:
0:25

29

```

24 of

```

Performing SFCFS preinstall tasks
..... Done

Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

```

```

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....

```

```
.....
Done

    Installing VRTSvcs rpm
    .....
    .....
Done

    Installing VRTScps rpm
    .....
    .....
Done

    Installing VRTSvcscag rpm
    .....
    .....
Done

    Installing VRTSvcscdr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
    .....
Done

    Installing VRTSdbed rpm
    .....
    .....
Done

    Installing VRTSglm rpm
    .....
    .....
Done

    Installing VRTScavf rpm
    .....
    .....
Done

    Installing SFCFS:
82%

.....

Estimated time remaining:
0:25
```

```

    Performing SFCFS preinstall tasks
    ..... Done

    Installing VRTSvlic rpm
    .....
    Done

    Installing VRTSperl rpm
    .....
    Done

    Installing VRTSspt rpm
    .....
    Done

    Installing VRTSvxvm rpm
    .....
    Done

    Installing VRTSaslapm rpm
    .....
    Done

    Installing VRTSob rpm
    .....
    Done

    Installing VRTSslmconv rpm
    ..... Done

    Installing VRTSvxfs rpm
    .....
    Done

    Installing VRTSfssdk rpm
    .....
    Done
```

```
Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSllt rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcS rpm
.....
.....
Done

Installing VRTScps rpm
.....
.....
Done

Installing VRTSvcSag rpm
.....
.....
Done

Installing VRTSvcSdr rpm
.....
.....
Done

Installing VRTSvcsea rpm
.....
.....
Done
```

```

Installing VRTSdbed rpm
.....
Done

Installing VRTSglm rpm
.....
Done

Installing VRTScavf rpm
.....
Done

Installing VRTSgms rpm -\|      Installing SFCFS:
86%
_____

Estimated time remaining:
0:20

29                                     25 of

Performing SFCFS preinstall tasks
..... Done

Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done

Installing VRTSvxvm rpm
.....
Done

```

```
Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSslmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
Done

Installing VRTSamf rpm
.....
Done

Installing VRTSvcS rpm
.....
```

```

.....
Done

    Installing VRTScps rpm
.....
Done

    Installing VRTSvcscag rpm
.....
Done

    Installing VRTSvcscdr rpm
.....
Done

    Installing VRTSvcsea rpm
.....
Done

    Installing VRTSdbed rpm
.....
Done

    Installing VRTSglm rpm
.....
Done

    Installing VRTScavf rpm
.....
Done

    Installing VRTSgms rpm
.....
Done

    Installing SFCFS:
86%

.....

Estimated time remaining:
0:20

```


29

```
Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done
```

```

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcS rpm
.....
.....
Done

Installing VRTScps rpm
.....
.....
Done

Installing VRTSvcSag rpm
.....
.....
Done

Installing VRTSvcSdr rpm
.....
.....
Done

Installing VRTSvcSea rpm
.....
.....
Done

```

```
Installing VRTSdbed rpm
.....
Done

Installing VRTSglm rpm
.....
Done

Installing VRTScavf rpm
.....
Done

Installing VRTSgms rpm
.....
Done

Installing VRTSodm rpm /-\ Installing SFCFS:
89%
```

```
Estimated time remaining:
0:15
```

26 of

29

```
Performing SFCFS preinstall tasks
..... Done

Installing VRTSvlic rpm
.....
Done

Installing VRTSperl rpm
.....
Done

Installing VRTSspt rpm
.....
Done
```

```
Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSsvmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
Done

Installing VRTSamf rpm
.....
```

```
.....
Done

    Installing VRTSvcs rpm
    .....
    .....
Done

    Installing VRTScps rpm
    .....
    .....
Done

    Installing VRTSvcscag rpm
    .....
    .....
Done

    Installing VRTSvcscdr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
    .....
Done

    Installing VRTSdbed rpm
    .....
    .....
Done

    Installing VRTSglm rpm
    .....
    .....
Done

    Installing VRTScavf rpm
    .....
    .....
Done

    Installing VRTSgms rpm
    .....
    .....
Done

    Installing VRTSodm rpm
    .....
    .....
Done
```

```
Installing SFCFS:
89%
```

```
Estimated time remaining:
0:15
```

26 of

29

```
Performing SFCFS preinstall tasks
.....
..... Done
```

```
Installing VRTSvlic rpm
.....
.....
Done
```

```
Installing VRTSperl rpm
.....
.....
Done
```

```
Installing VRTSspt rpm
.....
.....
Done
```

```
Installing VRTSvxvm rpm
.....
.....
Done
```

```
Installing VRTSaslapm rpm
.....
.....
Done
```

```
Installing VRTSob rpm
.....
.....
Done
```

```
Installing VRTSslmconv rpm
.....
..... Done
```

```
Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
Done

Installing VRTSamf rpm
.....
Done

Installing VRTSvcsc rpm
.....
Done

Installing VRTScps rpm
.....
Done

Installing VRTSvcscsag rpm
.....
Done
```

```
Installing VRTSvcscr rpm
.....
Done

Installing VRTSvcsea rpm
.....
Done

Installing VRTSdbed rpm
.....
Done

Installing VRTSglm rpm
.....
Done

Installing VRTScavf rpm
.....
Done

Installing VRTSgms rpm
.....
Done

Installing VRTSodm rpm
.....
Done

Installing VRTSsfmh rpm |/-\|/-\|/-\|/-\|/ Installing SFCFS:
93%
```

```
Estimated time remaining:
0:10

29
```

27 of

```
Performing SFCFS preinstall tasks
.....
..... Done
```



```
Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSslmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
```

```
.....
Done

    Installing VRTSgab rpm
    .....
    .....
Done

    Installing VRTSvxfen rpm
    .....
    .....
Done

    Installing VRTSamf rpm
    .....
    .....
Done

    Installing VRTSvcscs rpm
    .....
    .....
Done

    Installing VRTScps rpm
    .....
    .....
Done

    Installing VRTSvcscsag rpm
    .....
    .....
Done

    Installing VRTSvcscsdr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
    .....
Done

    Installing VRTSdbed rpm
    .....
    .....
Done

    Installing VRTSglm rpm
    .....
    .....
Done
```

```
Installing VRTScavf rpm
.....
.....
Done

Installing VRTSgms rpm
.....
.....
Done

Installing VRTSodm rpm
.....
.....
Done

Installing VRTSsfmh rpm
.....
.....
Done

Installing SFCFS:
93%
```

```
Estimated time remaining:
0:10
```

27 of

29

```
Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done
```

```
Installing VRTSvxvm rpm
.....
Done

Installing VRTSaslapm rpm
.....
Done

Installing VRTSob rpm
.....
Done

Installing VRTSsvmconv rpm
..... Done

Installing VRTSvxfs rpm
.....
Done

Installing VRTSfssdk rpm
.....
Done

Installing VRTSatClient rpm
..... Done

Installing VRTSatServer rpm
..... Done

Installing VRTSl1t rpm
.....
Done

Installing VRTSgab rpm
.....
Done

Installing VRTSvxfen rpm
.....
Done

Installing VRTSamf rpm
.....
```

```
.....
Done

    Installing VRTSvcs rpm
    .....
    .....
Done

    Installing VRTScps rpm
    .....
    .....
Done

    Installing VRTSvcscag rpm
    .....
    .....
Done

    Installing VRTSvcscdr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
    .....
Done

    Installing VRTSdbed rpm
    .....
    .....
Done

    Installing VRTSglm rpm
    .....
    .....
Done

    Installing VRTScavf rpm
    .....
    .....
Done

    Installing VRTSgms rpm
    .....
    .....
Done

    Installing VRTSodm rpm
    .....
    .....
Done
```

```

Installing VRTSsfmh rpm
.....
.....
Done

```

```

Performing SFCFS postinstall tasks -\|/-\|/-\|/ Installing SFCFS:
96%

```

```

Estimated time remaining:
0:05

```

28 of

29

```

Performing SFCFS preinstall tasks
.....
..... Done

```

```

Installing VRTSvlic rpm
.....
.....
Done

```

```

Installing VRTSperl rpm
.....
.....
Done

```

```

Installing VRTSspt rpm
.....
.....
Done

```

```

Installing VRTSvxvm rpm
.....
.....
Done

```

```

Installing VRTSaslapm rpm
.....
.....
Done

```

```

Installing VRTSob rpm
.....
.....
Done

```

```
Installing VRTSsvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....
.....
Done

Installing VRTSatClient rpm
.....
..... Done

Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcsc rpm
.....
.....
Done

Installing VRTScps rpm
.....
.....
Done

Installing VRTSvcscsag rpm
.....
```

```

.....
Done

    Installing VRTSvcldr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
    .....
Done

    Installing VRTSdbed rpm
    .....
    .....
Done

    Installing VRTSglm rpm
    .....
    .....
Done

    Installing VRTScavf rpm
    .....
    .....
Done

    Installing VRTSgms rpm
    .....
    .....
Done

    Installing VRTSodm rpm
    .....
    .....
Done

    Installing VRTSsfmh rpm
    .....
    .....
Done

    Performing SFCFS postinstall tasks
    .....
    ..... Done

    Installing SFCFS:
96%

```


Estimated time remaining:
0:05

28 of

29

Performing SFCFS preinstall tasks
.....
..... Done

Installing VRTSvlic rpm
.....
.....
Done

Installing VRTSperl rpm
.....
.....
Done

Installing VRTSspt rpm
.....
.....
Done

Installing VRTSvxvm rpm
.....
.....
Done

Installing VRTSaslapm rpm
.....
.....
Done

Installing VRTSob rpm
.....
.....
Done

Installing VRTSvvmconv rpm
.....
..... Done

Installing VRTSvxfs rpm
.....
.....
Done

Installing VRTSfssdk rpm
.....

```
.....
Done

    Installing VRTSatClient rpm
    .....
    ..... Done

    Installing VRTSatServer rpm
    .....
    ..... Done

    Installing VRTSllt rpm
    .....
    .....
Done

    Installing VRTSgab rpm
    .....
    .....
Done

    Installing VRTSvxfen rpm
    .....
    .....
Done

    Installing VRTSamf rpm
    .....
    .....
Done

    Installing VRTSvcs rpm
    .....
    .....
Done

    Installing VRTScps rpm
    .....
    .....
Done

    Installing VRTSvcsag rpm
    .....
    .....
Done

    Installing VRTSvcsdr rpm
    .....
    .....
Done

    Installing VRTSvcsea rpm
    .....
```

```

.....
Done

    Installing VRTSdbed rpm
.....
.....
Done

    Installing VRTSglm rpm
.....
.....
Done

    Installing VRTScavf rpm
.....
.....
Done

    Installing VRTSgms rpm
.....
.....
Done

    Installing VRTSodm rpm
.....
.....
Done

    Installing VRTSsfmh rpm
.....
.....
Done

    Performing SFCFS postinstall tasks
.....
..... Done

    Copying installer libraries and scripts -\|/-\|/-\    Installing SFCFS:
100%

```

```

    Estimated time remaining:
0:00

```

29

29 of

```

    Performing SFCFS preinstall tasks
    .....
    ..... Done

    Installing VRTSvlic rpm
    .....
    .....
Done

    Installing VRTSperl rpm
    .....
    .....
Done

    Installing VRTSspt rpm
    .....
    .....
Done

    Installing VRTSvxvm rpm
    .....
    .....
Done

    Installing VRTSaslapm rpm
    .....
    .....
Done

    Installing VRTSob rpm
    .....
    .....
Done

    Installing VRTSlvmconv rpm
    .....
    ..... Done

    Installing VRTSvxfs rpm
    .....
    .....
Done

    Installing VRTSfssdk rpm
    .....
    .....
Done

    Installing VRTSatClient rpm
    .....
    ..... Done

```

```
Installing VRTSatServer rpm
.....
..... Done

Installing VRTSl1t rpm
.....
.....
Done

Installing VRTSgab rpm
.....
.....
Done

Installing VRTSvxfen rpm
.....
.....
Done

Installing VRTSamf rpm
.....
.....
Done

Installing VRTSvcsc rpm
.....
.....
Done

Installing VRTScps rpm
.....
.....
Done

Installing VRTSvcscsag rpm
.....
.....
Done

Installing VRTSvcscdr rpm
.....
.....
Done

Installing VRTSvcscsea rpm
.....
.....
Done

Installing VRTSdbed rpm
.....
.....
Done
```

```

Installing VRTSglm rpm
.....
Done

Installing VRTScavf rpm
.....
Done

Installing VRTSgms rpm
.....
Done

Installing VRTSodm rpm
.....
Done

Installing VRTSsfmh rpm
.....
Done

Performing SFCFS postinstall tasks
..... Done

Copying installer libraries and scripts
..... Done

Veritas Storage Foundation Cluster File System Install completed successfully

Veritas Storage
Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

To comply with the terms of Symantec's End User License Agreement, you have 60
days to either:

* Enter a valid license key matching the functionality in use on the systems

```

* Enable keyless licensing and manage the systems with a Management Server. For more details visit <http://go.symantec.com/sfhakeyless>. The product is fully functional during these 60 days.

- 1) Enter a valid license key
- 2) Enable keyless licensing and complete system licensing later

How would you like to license the systems? [1-2,q] (2) y

Invalid input. Please retry.

How would you like to license the systems? [1-2,q] (2) 2

Checking system licensing

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

Would you like to enable the Veritas Volume Replicator? [y,n,q] (n) n\

Registering SFCFS license

Would you like to configure SFCFS on pnqaextapps460-7 pnextappsucs460-1? [y,n,q]
(n) y

I/O Fencing

It needs to be determined at this time if you plan to configure I/O Fencing in enabled or disabled mode, as well as help in determining the number of network interconnects (NICS) required on

your systems. If you configure I/O Fencing in enabled mode, only a single NIC is required, though at least two are recommended.

A split brain can occur if servers within the cluster become unable to communicate for any number of reasons. If I/O Fencing is not enabled, you run the risk of data corruption should a split

brain occur. Therefore, to avoid data corruption due to split brain in CFS environments, I/O Fencing has to be enabled.

If you do not enable I/O Fencing, you do so at your own risk

See the Administrator's Guide for more information on I/O Fencing

#####Below should be opted as no and should be configured later.

Do you want to configure I/O Fencing in enabled mode? [y,n,q,?] (y) y

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

To configure VCS, answer the set of questions on the next screen.

When [b] is presented after a question, 'b' may be entered to go back to the first question of the configuration set.

When [?] is presented after a question, '?' may be entered for help or additional information about the question.

Following each set of questions, the information you have entered will be presented for confirmation. To repeat the set of questions and correct any previous errors, enter 'n' at the

confirmation prompt.

No configuration changes are made to the systems until all configuration questions are completed and confirmed.

Press [Enter] to continue:

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

To configure VCS for SFCFS the following information is required:

A unique cluster name

A unique cluster ID number between 0-65535

One or more NICs per system used for heartbeat links

One or more heartbeat links are configured as private links

You can configure one heartbeat link as a low-priority link

All systems are being configured to create one cluster.

Enter the unique cluster name: [q,?] murcluster

Enter a unique cluster ID number between 0-65535: [b,q,?] 0 1

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

- 1) Configure heartbeat links using LLT over Ethernet
- 2) Configure heartbeat links using LLT over UDP

- 3) Automatically detect configuration for LLT over Ethernet
- b) Back to previous menu

How would you like to configure heartbeat links? [1-3,b,q,?] (1) 3

On Linux systems, only activated NICs could be detected and configured automatically.

Press [Enter] to continue:

Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

Logs are being written to /var/tmp/installer-201306302040lvT while installer is
in progress

Configuring LLT links: 0%

Estimated time
remaining:

0 of 4

Checking system NICs on pnqaextapps460-7 Configuring LLT links:
25%

```

Estimated time remaining:
0:05
1 of
4

```

```

Checking system NICs on pnqaextapps460-7
.....
..... 4 NICs found

```

```

Configuring LLT links:
25%
_____
_____

```

```

Estimated time remaining:
0:05
1 of
4

```

```

Checking system NICs on pnqaextapps460-7
.....
..... 4 NICs found

```

```

Checking system NICs on pnextappsucs460-1    Configuring LLT links:
50%
_____
_____

```

```

Estimated time remaining:
0:15
2 of
4

```

```

Checking system NICs on pnqaextapps460-7
.....
..... 4 NICs found

```

```

Checking system NICs on pnextappsucs460-1
.....
..... 4 NICs found

```

Configuring LLT links:

50%

Estimated time remaining:

0:15

2 of

4

Checking system NICs on pnqaextapps460-7

.....
 4 NICs found

Checking system NICs on pnextappsucs460-1

.....
 4 NICs found

Checking network links Configuring LLT links:

75%

Estimated time remaining:

0:05

3 of

4

Checking system NICs on pnqaextapps460-7

.....
 4 NICs found

Checking system NICs on pnextappsucs460-1

.....
 4 NICs found

Checking network links

.....
 1 link
 found

Configuring LLT links:

75%

```

Estimated time remaining:
0:05
4
3 of

```

```

Checking system NICs on pnqaextapps460-7
.....
..... 4 NICs found

Checking system NICs on pnextappsucs460-1
.....
..... 4 NICs found

Checking network links
.....
..... 1 link
found

Setting link priority    Configuring LLT links:
100%

```

```

Estimated time remaining:
0:00
4
4 of

```

```

Checking system NICs on pnqaextapps460-7
.....
..... 4 NICs found

Checking system NICs on pnextappsucs460-1
.....
..... 4 NICs found

Checking network links
.....
..... 1 link
found

```

```

    Setting link priority
    .....
    .....
Done

                                                                    Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

Cluster information verification:

Cluster Name:      murcluster

Cluster ID Number: 1

Private Heartbeat NICs for pnqaextapps460-7:

link1=eth3

Private Heartbeat NICs for pnextappsucs460-1:

link1=eth3

Is this information correct? [y,n,q,b,?] (y)

All SFCFS processes that are currently running must be stopped

Do you want to stop SFCFS processes now? [y,n,q,?] (y)

                                                                    Veritas Storage

Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

Logs are being written to /var/tmp/installer-201306302040lvT while installer is
in progress

```

Stopping SFCFS: 0%

```

_____
_____
_____

```

Estimated time
remaining:

0 of 11

Performing SFCFS prestop tasks |/-\ Stopping SFCFS:
9%

```

_____
_____
_____

```

Estimated time remaining:

1:40

1 of

11

Performing SFCFS prestop tasks

```

.....
..... Done

```

Stopping SFCFS:
9%

```

_____
_____
_____

```

Estimated time remaining:

1:40

1 of

11

```
Performing SFCFS prestop tasks
.....
..... Done
```

```
Stopping vxgms |/      Stopping SFCFS:
18%
.....
.....
```

```
Estimated time remaining:
0:50
.....
..... 2 of
11
```

```
Performing SFCFS prestop tasks
.....
..... Done
```

```
Stopping vxgms
.....
..... Done
```

```
Stopping SFCFS:
18%
.....
.....
```

```
Estimated time remaining:
0:50
.....
..... 2 of
11
```

```
Performing SFCFS prestop tasks
.....
..... Done
```

```
Stopping vxgms
.....
..... Done
```



```

Stopping vxglm -\|      Stopping SFCFS:
27%

```

```

Estimated time remaining:
0:40

```

3 of

11

```

Performing SFCFS prestop tasks
.....
..... Done

```

```

Stopping vxgms
.....
..... Done

```

```

Stopping vxglm
.....
..... Done

```

```

Stopping SFCFS: 27%

```

```

Estimated time remaining:
0:40

```

3 of

11

```

Performing SFCFS prestop tasks
.....
..... Done

```

```

Stopping vxgms
.....
..... Done

```

```
Stopping vxglm
.....
..... Done

Stopping vxcperv /-      Stopping SFCFS:
36%
_____
_____
_____
```

```
Estimated time remaining:
0:25
11
4 of
```

```
Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
..... Done

Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
..... Done

Stopping SFCFS:
36%
_____
_____
_____
```

```
Estimated time remaining:
0:25
11
4 of
```

```
Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
..... Done

Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
..... Done

Stopping had \ | Stopping SFCFS:
45%
-----
-----
```

```
Estimated time remaining:
0:20
```

5 of

11

```
Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
..... Done

Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
```

```
.....
.... Done

    Stopping had
.....
.....
..... Done

    Stopping SFCFS:
45%
-----

Estimated time remaining:
0:20

11                                                    5 of

    Performing SFCFS prestop tasks
.....
..... Done

    Stopping vxgms
.....
..... Done

    Stopping vxglm
.....
..... Done

    Stopping vxcperv
.....
..... Done

    Stopping had
.....
..... Done

    Stopping hashadow /-    Stopping SFCFS:
54%
-----
```

Estimated time remaining:
0:15

6 of

11

Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
..... Done

Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
.... Done

Stopping had
.....
..... Done

Stopping hashadow
.....
.... Done

Stopping SFCFS:
54%



Estimated time remaining:
0:15

6 of

11

```

    Performing SFCFS prestop tasks
    .....
    ..... Done

    Stopping vxgms
    .....
    ..... Done

    Stopping vxglm
    .....
    ..... Done

    Stopping vxcperv
    .....
    ..... Done

    Stopping had
    .....
    ..... Done

    Stopping hashadow
    .....
    ..... Done

    Stopping CmdServer \|| Stopping SFCFS:
63%

```

```

    Estimated time remaining:
0:10

```

11

7 of

```

    Performing SFCFS prestop tasks
    .....
    ..... Done

    Stopping vxgms
    .....
    ..... Done

```

```
Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
..... Done

Stopping had
.....
..... Done

Stopping hashadow
.....
..... Done

Stopping CmdServer
.....
... Done

Stopping SFCFS:
63%
```

```
Estimated time remaining:
0:10

11
```

7 of

```
Performing SFCFS prestop tasks
..... Done

Stopping vxgms
.....
..... Done

Stopping vxglm
.....
..... Done
```

```
Stopping vxcpserv
.....
.....
.... Done

Stopping had
.....
.....
..... Done

Stopping hashadow
.....
.....
.... Done

Stopping CmdServer
.....
.....
... Done

Stopping amf /-      Stopping SFCFS:
72%
_____
_____
```

Estimated time remaining:
0:05

11

8 of

```
Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
.....
..... Done

Stopping vxglm
.....
.....
..... Done

Stopping vxcpserv
.....
.....
.... Done
```



```

Stopping had
.....
..... Done

Stopping hashadow
.....
..... Done

Stopping CmdServer
.....
..... Done

Stopping amf
.....
..... Done

Stopping SFCFS:
72%

```

```

Estimated time remaining:
0:05

```

8 of

11

```

Performing SFCFS prestop tasks
..... Done

Stopping vxgms
..... Done

Stopping vxglm
..... Done

Stopping vxcperv
..... Done

```

```
Stopping had
.....
..... Done

Stopping hashadow
.....
.... Done

Stopping CmdServer
.....
... Done

Stopping amf
.....
..... Done

Stopping vxfen \\\ Stopping SFCFS:
81%
_____
```

```
Estimated time remaining:
0:05
11 9 of
```

```
Performing SFCFS prestop tasks
..... Done

Stopping vxgms
..... Done

Stopping vxglm
..... Done

Stopping vxcperv
..... Done
```

```

Stopping had
.....
..... Done
    
```

```

Stopping hashadow
.....
..... Done
    
```

```

Stopping CmdServer
.....
..... Done
    
```

```

Stopping amf
.....
..... Done
    
```

```

Stopping vxfen
.....
..... Done
    
```

```

Stopping SFCFS:
81%
_____
_____
    
```

```

Estimated time remaining:
0:05
11
    
```

9 of

```

Performing SFCFS prestop tasks
..... Done
    
```

```

Stopping vxgms
.....
..... Done
    
```

```

Stopping vxglm
.....
..... Done
    
```

```
Stopping vxcpserv
.....
.....
.... Done

Stopping had
.....
.....
..... Done

Stopping hashadow
.....
.....
.... Done

Stopping CmdServer
.....
.....
... Done

Stopping amf
.....
.....
..... Done

Stopping vxfen
.....
.....
..... Done

Stopping gab /-      Stopping SFCFS:
90%
```

Estimated time remaining:
0:03
11

10 of

```
Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
.....
..... Done
```

```
Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
..... Done

Stopping had
.....
..... Done

Stopping hashadow
.....
..... Done

Stopping CmdServer
.....
..... Done

Stopping amf
.....
..... Done

Stopping vxfen
.....
..... Done

Stopping gab
.....
..... Done

Stopping SFCFS:
90%
```

```
Estimated time remaining:
0:03
```

11

10 of

```
Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
..... Done

Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
..... Done

Stopping had
.....
..... Done

Stopping hashadow
.....
..... Done

Stopping CmdServer
.....
... Done

Stopping amf
.....
..... Done

Stopping vxfen
.....
..... Done

Stopping gab
.....
..... Done

Stopping llt \ |      Stopping SFCFS:
100%
```

Estimated time remaining:
0:00

11 of

11

```
Performing SFCFS prestop tasks
.....
..... Done

Stopping vxgms
.....
..... Done

Stopping vxglm
.....
..... Done

Stopping vxcperv
.....
..... Done

Stopping had
.....
..... Done

Stopping hashadow
.....
..... Done

Stopping CmdServer
.....
... Done

Stopping amf
.....
..... Done
```

```

Stopping vxfen
.....
..... Done

Stopping gab
.....
..... Done

Stopping llc
.....
..... Done

Veritas Storage Foundation Cluster File System Shutdown completed successfully

Veritas Storage
Foundation Cluster File System 5.1 SP1 Install
Program

pnqaextapps460-7 pnextappsucs460-
1

Logs are being written to /var/tmp/installer-201306302040lvT while installer is
in progress

Starting SFCFS: 0%

=====
=====
=====

Estimated time
remaining:

0 of 23

Performing SFCFS configuration /-\\/-\\/-\\ Starting SFCFS: 4%

=====
=====

```



```

Estimated time remaining:
4:20
23
1 of

```

```

Performing SFCFS configuration
.....
..... Done

Starting SFCFS: 4%
=====
=====

```

```

Estimated time remaining:
4:20
23
1 of

```

```

Performing SFCFS configuration
.....
..... Done

Starting vxdmp /- Starting SFCFS:
8%
=====
=====
=====

```

```

Estimated time remaining:
2:25
23
2 of

```

```

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....

```

```
.....
..... Done

Starting SFCFS:
8%
_____
_____
_____
```

```
Estimated time remaining:
2:25
23
2 of
```

```
Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio \  Starting SFCFS: 13%
_____
_____
```

```
Estimated time remaining:
1:45
23
3 of
```

```
Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
```

.....
 Done

Starting SFCFS: 13%

.....

Estimated time remaining:

1:45

3 of

23

Performing SFCFS configuration

.....
 Done

Starting vxdmp

.....

 Done

Starting vxio

.....

 Done

Starting vxspec /- Starting SFCFS:
 17%

.....

Estimated time remaining:

1:25

4 of

23

Performing SFCFS configuration

.....
 Done

Starting vxdmp

.....

```

.....
..... Done

    Starting vxio
.....
..... Done

    Starting vxspec
.....
..... Done

    Starting SFCFS:
17%
_____
_____

    Estimated time remaining:
1:25
23
4 of

    Performing SFCFS configuration
.....
..... Done

    Starting vxdmp
.....
..... Done

    Starting vxio
.....
..... Done

    Starting vxspec
.....
..... Done

    Starting vxconfigd    Starting SFCFS:
21%
_____
_____

```

Estimated time remaining:
1:40

5 of

23

```

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting SFCFS: 21%

```

Estimated time remaining:
1:40

5 of

23

```

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....

```

```
.....
..... Done

    Starting vxio
.....
..... Done

    Starting vxspec
.....
..... Done

    Starting vxconfigd
.....
... Done

    Starting vxesd      Starting SFCFS:
26%
.....
.....
```

Estimated time remaining:
1:25

23

6 of

```
    Performing SFCFS configuration
..... Done

    Starting vxdmp
..... Done

    Starting vxio
..... Done

    Starting vxspec
..... Done
```

```

Starting vxconfigd
.....
... Done

```

```

Starting vxesd
.....
..... Done

```

Starting SFCFS: 26%

```

Estimated time remaining:
1:25
23

```

6 of

```

Performing SFCFS configuration
.....
..... Done

```

```

Starting vxdmp
.....
..... Done

```

```

Starting vxio
.....
..... Done

```

```

Starting vxspec
.....
..... Done

```

```

Starting vxconfigd
.....
... Done

```

```

Starting vxesd
.....
..... Done

```

```

Starting vxrelocd      Starting SFCFS:
30%
_____
_____
_____

```

```

Estimated time remaining:
1:15
_____
_____
_____
23
_____
_____
_____

```

7 of

```

Performing SFCFS configuration
.....
..... Done

```

```

Starting vxdmp
.....
..... Done

```

```

Starting vxio
.....
..... Done

```

```

Starting vxspec
.....
..... Done

```

```

Starting vxconfigd
.....
..... Done

```

```

Starting vxesd
.....
..... Done

```

```

Starting vxrelocd
.....
..... Done

```

```

Starting SFCFS:
30%
_____
_____
_____

```


Estimated time remaining:

1:15

7 of

23

Performing SFCFS configuration

.....
..... Done

Starting vxdmp

.....
.....
..... Done

Starting vxio

.....
.....
..... Done

Starting vxspec

.....
.....
..... Done

Starting vxconfigd

.....
.....
... Done

Starting vxesd

.....
.....
..... Done

Starting vxrelocd

.....
.....
.... Done

Starting vxcached Starting SFCFS:

34%

.....
.....

Estimated time remaining:

1:15

8 of

23

Performing SFCFS configuration

 Done

Starting vxdmp

 Done

Starting vxio

 Done

Starting vxspec

 Done

Starting vxconfigd

 ... Done

Starting vxesd

 Done

Starting vxrelocd

 Done

Starting vxcached

 Done

Starting SFCFS: 34%

Estimated time remaining:
1:15

8 of

23

```
Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd    Starting SFCFS:
39%
```

Estimated time remaining:

1:10

9 of

23

Performing SFCFS configuration

.....
..... Done

Starting vxdmp

.....
..... Done

Starting vxio

.....
..... Done

Starting vxspec

.....
..... Done

Starting vxconfigd

.....
... Done

Starting vxesd

.....
..... Done

Starting vxrelocd

.....
.... Done

Starting vxcached

.....
.... Done

Starting vxconfigbackupd

.....
Done

Starting SFCFS:

39%

Estimated time remaining:
1:10

9 of

23

```

Performing SFCFS configuration
..... Done

Starting vxdmp
..... Done

Starting vxio
..... Done

Starting vxspec
..... Done

Starting vxconfigd
..... Done

Starting vxesd
..... Done

Starting vxrelocd
..... Done

Starting vxcached
..... Done
    
```

```

Starting vxconfigbackupd
.....
.....
Done

```

```

Starting vxattachd      Starting SFCFS:
43%
_____
_____
_____

```

```

Estimated time remaining:
1:05
                                                    10 of
23

```

```

Performing SFCFS configuration
.....
..... Done

```

```

Starting vxdmp
.....
..... Done

```

```

Starting vxio
.....
..... Done

```

```

Starting vxspec
.....
..... Done

```

```

Starting vxconfigd
.....
..... Done

```

```

Starting vxesd
.....
..... Done

```

```

Starting vxrelocd
.....
..... Done

```

```

Starting vxcached
.....
.....
.... Done

```

```

Starting vxconfigbackupd
.....
.....
Done

```

```

Starting vxattachd
.....
.....
... Done

```

```

Starting SFCFS:
43%
_____
_____

```

```

Estimated time remaining:
1:05
23
10 of

```

```

Performing SFCFS configuration
.....
..... Done

```

```

Starting vxdmp
.....
.....
..... Done

```

```

Starting vxio
.....
.....
..... Done

```

```

Starting vxspec
.....
.....
..... Done

```

```

Starting vxconfigd
.....
.....
... Done

```

```
Starting vxesd
.....
..... Done

Starting vxrelocd
.....
..... Done

Starting vxcached
.....
..... Done

Starting vxconfigbackupd
.....
..... Done

Starting vxattachd
.....
... Done

Starting vxportal \\\ Starting SFCFS:
47%
_____
_____
```

Estimated time remaining:
0:55
23 11 of

```
Performing SFCFS configuration
..... Done

Starting vxdmp
..... Done

Starting vxio
..... Done
```



```
Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting SFCFS:
47%
_____
_____
```

```
Estimated time remaining:
0:55
```

23

11 of

```

    Performing SFCFS configuration
    .....
    ..... Done

    Starting vxdmp
    .....
    ..... Done

    Starting vxio
    .....
    ..... Done

    Starting vxspec
    .....
    ..... Done

    Starting vxconfigd
    .....
    ... Done

    Starting vxesd
    .....
    ..... Done

    Starting vxrelocd
    .....
    .... Done

    Starting vxcached
    .....
    .... Done

    Starting vxconfigbackupd
    .....
    Done

    Starting vxattachd
    .....
    ... Done

    Starting vxportal
    .....

```

```
.....
.... Done
```

```
Starting fdd /-      Starting SFCFS:
52%
```

```
Estimated time remaining:
0:45
```

12 of

23

```
Performing SFCFS configuration
.....
..... Done
```

```
Starting vxdmp
.....
..... Done
```

```
Starting vxio
.....
..... Done
```

```
Starting vxspec
.....
..... Done
```

```
Starting vxconfigd
.....
... Done
```

```
Starting vxesd
.....
..... Done
```

```
Starting vxrelocd
.....
.... Done
```

```
Starting vxcached
.....
.....
.... Done

Starting vxconfigbackupd
.....
.....
Done

Starting vxattachd
.....
.....
... Done

Starting vxportal
.....
.....
.... Done

Starting fdd
.....
.....
..... Done

Starting SFCFS:
52%
```

Estimated time remaining:
0:45

2312 of

```
Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
.....
..... Done

Starting vxio
.....
.....
..... Done
```

```
Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting llt \|-\\|/-\\| Starting SFCFS:
56%
```

Estimated time remaining:

0:50

13 of

23

```

    Performing SFCFS configuration
    .....
    ..... Done

    Starting vxdmp
    .....
    ..... Done

    Starting vxio
    .....
    ..... Done

    Starting vxspec
    .....
    ..... Done

    Starting vxconfigd
    .....
    ... Done

    Starting vxesd
    .....
    ..... Done

    Starting vxrelocd
    .....
    .... Done

    Starting vxcached
    .....
    .... Done

    Starting vxconfigbackupd
    .....
    .....
    Done

    Starting vxattachd
    .....
```

```
.....
... Done

    Starting vxportal
.....
..... Done

    Starting fdd
.....
..... Done

    Starting llc
.....
..... Done

    Starting SFCFS:
56%
_____
_____

    Estimated time remaining:
0:50

23 13 of

    Performing SFCFS configuration
..... Done

    Starting vxdmp
..... Done

    Starting vxio
..... Done

    Starting vxspec
..... Done
```

```

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting llt
.....
..... Done

Starting gab -\| Starting SFCFS:
60%

```

Estimated time remaining:
0:40

14 of

23

```
Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
.....
Done

Starting vxattachd
.....
```

```
.....
... Done

    Starting vxportal
.....
..... Done

    Starting fdd
.....
..... Done

    Starting llt
.....
..... Done

    Starting gab
.....
..... Done

    Starting SFCFS:
60%
_____
_____

    Estimated time remaining:
0:40

23                                     14 of

    Performing SFCFS configuration
.....
..... Done

    Starting vxdmp
.....
..... Done

    Starting vxio
.....
..... Done
```

```
Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting lltd
.....
..... Done

Starting gab
.....
..... Done
```

```

Starting vxfen /-\\/-\\/-\\/-      Starting SFCFS:
65%

```

```

Estimated time remaining:
0:40

```

15 of

23

```

Performing SFCFS configuration
.....
..... Done

```

```

Starting vxdmp
.....
..... Done

```

```

Starting vxio
.....
..... Done

```

```

Starting vxspec
.....
..... Done

```

```

Starting vxconfigd
.....
... Done

```

```

Starting vxesd
.....
..... Done

```

```

Starting vxrelocd
.....
.... Done

```

```

Starting vxcached
.....
.... Done

```

```
Starting vxconfigbackupd
.....
.....
Done

Starting vxattachd
.....
.....
... Done

Starting vxportal
.....
.....
.... Done

Starting fdd
.....
.....
..... Done

Starting llc
.....
.....
..... Done

Starting gab
.....
.....
..... Done

Starting vxfen
.....
.....
..... Done

Starting SFCFS:
65%
```

```
Estimated time remaining:
0:40

23
```

15 of

```
Performing SFCFS configuration
.....
..... Done
```

```
Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done
```

```

Starting llc
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm \  Starting SFCFS:
69%

```

```

Estimated time remaining:
0:35

```

16 of

23

```

Performing SFCFS configuration
..... Done

Starting vxdmp
..... Done

Starting vxio
..... Done

Starting vxspec
..... Done

Starting vxconfigd
..... Done

```

```
Starting vxesd
.....
..... Done

Starting vxrelocd
.....
..... Done

Starting vxcached
.....
..... Done

Starting vxconfigbackupd
.....
..... Done

Starting vxattachd
.....
..... Done

Starting vxportal
.....
..... Done

Starting fdd
.....
..... Done

Starting llc
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done
```


Starting SFCFS:
69%

Estimated time remaining:
0:35

16 of

23

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

```
Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting llt
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done

Starting had      Starting SFCFS:
73%

-----

Estimated time remaining:
0:35

23
```

```
Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
... Done

Starting vxcached
.....
... Done

Starting vxconfigbackupd
.....
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
```

```
.....
.... Done

    Starting fdd
.....
..... Done

    Starting llt
.....
..... Done

    Starting gab
.....
..... Done

    Starting vxfen
.....
..... Done

    Starting vxglm
.....
..... Done

    Starting had
.....
..... Done

    Starting SFCFS:
73%
_____

Estimated time remaining:
0:35

23

    Performing SFCFS configuration
..... Done
```

```
Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done
```

```
Starting llt
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done

Starting had
.....
..... Done

Starting hashadow /-      Starting SFCFS:
78%

_____
```

Estimated time remaining:
0:30

23

18 of

```
Performing SFCFS configuration
..... Done

Starting vxdmp
..... Done

Starting vxio
..... Done
```

```
Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting lltd
.....
..... Done

Starting gab
.....
..... Done
```

```

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done

Starting had
.....
..... Done

Starting hashadow
.....
.... Done

Starting SFCFS:
78%

```

```

Estimated time remaining:
0:30

```

23

18 of

```

Performing SFCFS configuration
..... Done

Starting vxdmp
..... Done

Starting vxio
..... Done

Starting vxspec
..... Done

```



```
Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting llc
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done
```

```

Starting vxglm
.....
..... Done

Starting had
.....
..... Done

Starting hashadow
.....
..... Done

Starting CmdServer \\\ Starting SFCFS:
82%

```

```

Estimated time remaining:
0:20

```

19 of

23

```

Performing SFCFS configuration
..... Done

Starting vxdmp
..... Done

Starting vxio
..... Done

Starting vxspec
..... Done

Starting vxconfigd
..... Done

```

```
Starting vxesd
.....
..... Done

Starting vxrelocd
.....
..... Done

Starting vxcached
.....
..... Done

Starting vxconfigbackupd
.....
..... Done

Starting vxattachd
.....
..... Done

Starting vxportal
.....
..... Done

Starting fdd
.....
..... Done

Starting llc
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done
```

```

Starting had
.....
..... Done

```

```

Starting hashadow
.....
.... Done

```

```

Starting CmdServer
.....
... Done

```

```

Starting SFCFS:
82%

```

```

Estimated time remaining:

```

```

0:20

```

```

19 of

```

```

23

```

```

Performing SFCFS configuration
.....
..... Done

```

```

Starting vxdmp
.....
..... Done

```

```

Starting vxio
.....
..... Done

```

```

Starting vxspec
.....
..... Done

```

```

Starting vxconfigd
.....
... Done

```

```
Starting vxesd
.....
..... Done

Starting vxrelocd
.....
..... Done

Starting vxcached
.....
..... Done

Starting vxconfigbackupd
.....
..... Done

Starting vxattachd
.....
..... Done

Starting vxportal
.....
..... Done

Starting fdd
.....
..... Done

Starting llc
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done
```

```

Starting had
.....
..... Done

Starting hashadow
.....
.... Done

Starting CmdServer
.....
... Done

Starting vxdbd /-      Starting SFCFS:
86%
_____

Estimated time remaining:
0:15
23
20 of

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

```

```
Starting vxesd
.....
..... Done

Starting vxrelocd
.....
..... Done

Starting vxcached
.....
..... Done

Starting vxconfigbackupd
.....
..... Done

Starting vxattachd
.....
..... Done

Starting vxportal
.....
..... Done

Starting fdd
.....
..... Done

Starting llc
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done
```

```
Starting had
.....
..... Done

Starting hashadow
.....
.... Done

Starting CmdServer
.....
... Done

Starting vxdbd
.....
..... Done

Starting SFCFS:
86%
```

```
Estimated time remaining:
0:15
```

23

20 of

```
Performing SFCFS configuration
..... Done

Starting vxcmp
..... Done

Starting vxio
..... Done

Starting vxspec
..... Done
```



```
Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting llt
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done
```

```

Starting vxglm
.....
..... Done

Starting had
.....
..... Done

Starting hashadow
.....
..... Done

Starting CmdServer
.....
... Done

Starting vxdbd
.....
..... Done

Starting vxgms \\\ Starting SFCFS:
91%

-----

Estimated time remaining:
0:10

23 21 of

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

```

```
Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting lltd
.....
..... Done

Starting gab
.....
..... Done
```

```
Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done

Starting had
.....
..... Done

Starting hashadow
.....
..... Done

Starting CmdServer
.....
... Done

Starting vxdbd
.....
..... Done

Starting vxgms
.....
..... Done

Starting SFCFS:
91%

Estimated time remaining:
0:10

23

Performing SFCFS configuration
.....
..... Done
```

```
Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done
```

```

Starting llt
.....
..... Done

Starting gab
.....
..... Done

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done

Starting had
.....
..... Done

Starting hashadow
.....
.... Done

Starting CmdServer
.....
... Done

Starting vxdbd
.....
..... Done

Starting vxgms
.....
..... Done

Starting vxodm -\\/-\\/-\\/-\\ Starting SFCFS:
95%

```

Estimated time remaining:
0:05

22 of

23

```

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
..... Done

Starting vxio
.....
..... Done

Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
.....
Done

Starting vxattachd
.....

```

```
.....
... Done

    Starting vxportal
.....
..... Done

    Starting fdd
.....
..... Done

    Starting llt
.....
..... Done

    Starting gab
.....
..... Done

    Starting vxfen
.....
..... Done

    Starting vxglm
.....
..... Done

    Starting had
.....
..... Done

    Starting hashadow
.....
..... Done

    Starting CmdServer
.....
... Done

    Starting vxdbd
.....
..... Done
```



```
Starting vxgms
.....
..... Done
```

```
Starting vxodm
.....
..... Done
```

```
Starting SFCFS:
95%
```

```
Estimated time remaining:
0:05
```

22 of

23

```
Performing SFCFS configuration
.....
..... Done
```

```
Starting vxdmp
.....
..... Done
```

```
Starting vxio
.....
..... Done
```

```
Starting vxspec
.....
..... Done
```

```
Starting vxconfigd
.....
... Done
```

```
Starting vxesd
.....
..... Done
```

```
Starting vxrelocd
.....
.....
.... Done

Starting vxcached
.....
.....
.... Done

Starting vxconfigbackupd
.....
.....
Done

Starting vxattachd
.....
.....
... Done

Starting vxportal
.....
.....
.... Done

Starting fdd
.....
.....
..... Done

Starting lltd
.....
.....
..... Done

Starting gab
.....
.....
..... Done

Starting vxfen
.....
.....
..... Done

Starting vxglm
.....
.....
..... Done

Starting had
.....
.....
..... Done
```

```
Starting hashadow
.....
.....
.... Done

Starting CmdServer
.....
.....
... Done

Starting vxdbd
.....
.....
..... Done

Starting vxgms
.....
.....
..... Done

Starting vxodm
.....
.....
..... Done

Performing SFCFS poststart tasks    Starting SFCFS:
100%

Estimated time remaining:
0:00

23                                     23 of

Performing SFCFS configuration
.....
..... Done

Starting vxdmp
.....
.....
..... Done

Starting vxio
.....
.....
..... Done
```

```
Starting vxspec
.....
..... Done

Starting vxconfigd
.....
... Done

Starting vxesd
.....
..... Done

Starting vxrelocd
.....
.... Done

Starting vxcached
.....
.... Done

Starting vxconfigbackupd
.....
Done

Starting vxattachd
.....
... Done

Starting vxportal
.....
.... Done

Starting fdd
.....
..... Done

Starting llc
.....
..... Done

Starting gab
.....
..... Done
```

```

Starting vxfen
.....
..... Done

Starting vxglm
.....
..... Done

Starting had
.....
..... Done

Starting hashadow
.....
..... Done

Starting CmdServer
.....
... Done

Starting vxdbd
.....
..... Done

Starting vxgms
.....
..... Done

Starting vxodm
.....
..... Done

Performing SFCFS poststart tasks
..... Done

Veritas Storage Foundation Cluster File System Startup completed successfully

Veritas Storage
Foundation Cluster File System 5.1 SP1 Install
Program

```

```
pnqaextapps460-7 pnextappsucs460-
1
```

Fencing configuration

- 1) Configure CP client based fencing
- 2) Configure disk based fencing

Select the fencing mechanism to be configured in this Application Cluster: [1-2,q] 2

Do you have SCSI3 PR enabled disks? [y,n,q] (y)

Since you have selected to configure Disk based fencing, you would be asked to give either the Disk group to be used as co-ordinator or asked to create disk group and the mechanism to be used

Select one of the options below for fencing disk group:

- 1) Create a new disk group
- b) Back to previous menu

Enter the choice for a disk group: [1-1,b,q] 1

List of available disks to create a new disk group

A new disk group can not be created as the number of free VxVM CDS disks available is less than three. If there are disks available which are not under VxVM control, use the command vxdisksetup

to initialize them as VxVM disks first.

Do you want to retry fencing configuration? [y,n,q,b] (n) y ^[[B n

Would you like to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] (y) y

Upload completed successfully

installer log files, summary file, and response file are saved at:

```
/opt/VRTS/install/logs/installer-201306302040lvT
```

Configuring Shared Volume

The following section describes the post installation steps for shared volume creation and post installation checks and configurations for Veritas SFCFS.



Important: The command inputs provided in this procedure are only for reference purposes. It is highly recommended to not to use these sample inputs as it might vary in the actual implementation.

1. Post installation of Veritas cluster, perform the following checks:

- LLT Link Checks:

LLT configuration needs to be done on all cluster nodes. This will be automatically done as a part of VCS installation. However, the prerequisite is to have two dedicated private Ethernet cable connections.

The following is a sample of the llttab configuration:

```
[root@extapps460-5 ~]# cat /etc/llttab

set-node extapps460-5

set-cluster 9999

link eth4 eth-a0:36:9f:22:97:4d - ether - -

link eth5 eth-a0:36:9f:22:97:4e - ether - -

link-lowpri eth2 eth-44:d3:ca:da:7d:14 - ether - -
```

After proper LLT configuration on all cluster nodes all nodes need to be rebooted.

After rebooting, execute the command “**gabconfig -a**” to check if LLT is running and if does not have any errors.

- GAB Configuration Checks:

Check if there is required membership between GAB ports of all the hosts. If the membership is shown as "jeopardy", it means there is some fault in VCS setup and needs to be corrected with the VCS documents and support. Generally this fault is due to failure in heart bit cable connection.

The following sample shows the valid membership for the GAB.

```
[root@pnqaextapps460-2 ~]# gabconfig -a
```

GAB Port Memberships

```
=====

Port a gen      9f8b01 membership 0123

Port b gen      9f8b02 membership 0123

Port d gen      9f8b06 membership 0123

Port f gen      9f8b0f membership 0123

Port h gen      9f8b05 membership 0123

Port u gen      9f8b0d membership 0123

Port v gen      9f8b08 membership 0123

Port w gen      9f8b0a membership 0123
```

- Listing and mapping Storage Disks.

Check if all shared disks are listed by Veritas as shown below.

```
[root@pnqaextapps460-1 rhel5_x86_64]# vxdisk -o allids list
DEVICE  TYPE      DISK      GROUP      STATUS
disk_0   auto none      -          -          online invalid
sda      auto none      -          -          online invalid
st2540-0-3 auto      -          -          error
st2540-0-4 auto      -          -          error
st2540-0-5 auto      -          -          error
st2540-0-6 auto      -          -          error
st2540-0-7 auto      -          -          error
```

Map each of the listed disks to configured multipaths as below.

```
[root@pnqaextapps460-2 ~]# vxdm -v8783 bin-vxdiskadm
Volume Manager Support Operations
Menu: VolumeManager/Disk

1Add on initialize one or more disks
2Encapsulate one or more disks
3Remove a disk
4Remove a disk for replacement
5Replace a failed or removed disk
6Mirror volumes on a disk
7Move volumes from a disk
8Enable access to (import) a disk group
9Remove access to (export) a disk group
10Enable (online) a disk device
11Disable (offline) a disk device
12Mark a disk as a spare for a disk group
13Turn off the spare flag on a disk
14Unrelocate subdisks back to a disk
15Exclude a disk from hot-relocation use
16Hide a disk available for hot-relocation use
17Prevent multipathing/Suppress devices from VxVM's view
18Allow multipathing/Unsuppress devices from VxVM's view
19List currently suppressed/non-multipathed devices
20Change the disk naming scheme
21Choose/Display the default disk layouts
22List disk information

?Display help about menu
??Display help about the running system
exit from menu

Select an operation to perform: list

List disk information.
Menu: VolumeManager/Disk/ListDisk

Use this menu operation to display a list of disks. You can
also choose to list detailed information about the disk at
a specific disk device address.

Enter disk device or 'all' (address: all or 0) (default: all):
DEVICE  DISK      GROUP      STATUS
disk_0   -          -          online invalid
sda      -          -          online invalid
st2540-0-3 -          -          error
st2540-0-4 -          -          error
st2540-0-5 -          -          error
st2540-0-6 -          -          error
st2540-0-7 -          -          error
```


81688

```
st2540-0_3 >> CPd2
st2540-0_4 >> CPd1
st2540-0_5 >> CPd3
st2540-0_6 >> InputPath
st2540-0_7 >> ArchivePath
```

```
[root@nexusappsum1001 ~]# ifconfig -dev nappor@CDD1
WARNING: (n for help):
Command (n for help): w
The partition table has been altered.

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 22: Invalid argument
The kernel still uses the old table
The new table will be used at the next reboot
Syncing disks

[root@nexusappsum1001 ~]# ifconfig -dev nappor@CDD1
Command (n for help):
Command (n for help): w
The partition table has been altered.

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 22: Invalid argument
The kernel still uses the old table
The new table will be used at the next reboot
Syncing disks

[root@nexusappsum1001 ~]# ifconfig -dev nappor@CDD1
Command (n for help):
Command (n for help): w
The partition table has been altered.

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 22: Invalid argument
The kernel still uses the old table
The new table will be used at the next reboot
Syncing disks

[root@nexusappsum1001 ~]# ifconfig -dev nappor@ArchivePath
Command (n for help):
Command (n for help): w
The partition table has been altered.

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 22: Invalid argument
The kernel still uses the old table
The new table will be used at the next reboot
Syncing disks

[root@nexusappsum1001 ~]# ifconfig -dev nappor@tmuPath
Command (n for help):
Command (n for help): w
The partition table has been altered.

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 22: Invalid argument
The kernel still uses the old table
The new table will be used at the next reboot
Syncing disks
```

```
root@pnxetaappsucs460-1 ~#]# vxdisksetup -i st2540-0.3
root@pnxetaappsucs460-1 ~#]# vxdisksetup -i st2540-0.4
root@pnxetaappsucs460-1 ~#]# vxdisksetup -i st2540-0.5
root@pnxetaappsucs460-1 ~#]# vxdisksetup -i st2540-0.6
root@pnxetaappsucs460-1 ~#]# vxdisksetup -i st2540-0.7
```

```

[root@pnxtappsucs460-1 ~]# vxdisk -o all -l
SERVICE  TYPE  DISK  GROUP  STATUS
disk_0     auto none  -      -      online invalid
ada        auto none  -      -      online invalid
st2540-0_3 auto cddisk -      -      online
st2540-0_4 auto cddisk -      -      online
st2540-0_5 auto cddisk -      -      online
st2540-0_6 auto cddisk -      -      online
st2540-0_7 auto cddisk -      -      online

```

4. Create shared disk mount and creation of volume. The following steps need to be carried on either of the hosts of the cluster.

- To add shared disk group for input path use the command **vxldg -s init inputDG disk=st2540-0_6**.
 - Create volume using the command **vxassist -g inputDG make inputVol 5120g**.
 - Create file system using the command **mkfs -t vxfs -o bsize=4096,largefiles /dev/vx/rdisk/inputDG/inputVol**.
- To add shared disk group for archive path use the command **vxldg -s init archiveDG disk=st2540-0_7**.
 - Create volume using the command **vxassist -g archiveDG make archiveVol 6144g**.
 - Create file system using the command **mkfs -t vxfs -o bsize=4096,largefiles /dev/vx/rdisk/archiveDG/archiveVol**.
 - Create mount point on each node of the cluster.


```

mkdir /input_path
mkdir /archive

```
 - Add disk groups to the cluster configuration.


```

/opt/VRTS/bin/cfsdgadm add <disk_group> <cluster_node_1>=sw
<cluster_node_2>=sw

/opt/VRTS/bin/cfsdgadm add inputDG pnxtappsucs460-1=sw
pnqaextapps460-7=sw

/opt/VRTS/bin/cfsdgadm add archiveDG pnxtappsucs460-1=sw
pnqaextapps460-7=sw

```
 - Create service group and add mount point.


```

/opt/VRTS/bin/cfsmntadm add <disk_group> <volume_name> <mount_point>
<service_group_name> <cluster_node_1>=rw <cluster_node_2>=rw

/opt/VRTS/bin/cfsmntadm add ArchiveDG ArchiveVol /archive murSG
pnxtappsucs460-1=rw pnqaextapps460-7=rw

/opt/VRTS/bin/cfsmntadm add inputDG inputVol /input_path murSG
pnxtappsucs460-1=rw pnqaextapps460-7=rw

```

5. Setup MUR application as per the requirements for Scalability.

While installing, the RDPs sharing the shared disk need to have the same user id.

Shared disk needs to have all permission so that it will be accessed by all RDPs and this will prevent the errors while adding gateway.

The sample commands are:

```

User : mur_rdp , group : murgroup

chown -R mur_rdp /input_path

chgrp -R murgroup /archive

```

```
chmod -R 777 /input_path
```

6. Tune the VxFS file system.

VxFS file system can be tuned for better performance. **Vxtunefs** command is used to set the tuning parameters. The default values of these parameters are set when the volume is mounted. MUR application shows improved performance when the following tuning parameters are changed:

read_pref_io: The preferred read request size. The filesystem uses this in conjunction with the *read_nstream* value to determine how much data to read ahead. The default value is 64K. MUR performance is improved when this value is set to 128K.

read_nstream: This is the desired number of parallel read requests of size *read_pref_io* to have outstanding at one time. The file system uses the product of *read_nstream* multiplied by *read_pref_io* to determine its read ahead size. The default value for *read_nstream* is 1. If you know the hardware RAID configuration on the external storage then set *read_nstream* to be the number of columns (disks) in the disk array.

write_pref_io: The preferred write request size. The filesystem uses this in conjunction with the *write_nstream* value to determine how to do flush behind on writes. The default value is 64K. MUR performance is improved when this value is set to 128K.

write_nstream: This is the desired number of parallel write requests of size *write_pref_io* to have outstanding at one time. The file system uses the product of *write_nstream* multiplied by *write_pref_io* to determine when to do flush behind on writes. The default value for *write_nstream* is 1. For disk striping configurations, set *write_pref_io* and *write_nstream* to the same values as *read_pref_io* and *read_nstream*.

Use the following commands to tune Veritas file system.

```
$ /opt/VRTS/bin/vxtunefs -o
read_pref_io=131072,read_unit_io=131072,write_pref_io=131072,write_unit_io=131072
/input_path

$ /opt/VRTS/bin/vxtunefs -o
read_pref_io=131072,read_unit_io=131072,write_pref_io=131072,write_unit_io=131072
/archive
```

To ensure these values sustain a reboot and setup the file in the directory */etc/vx/tunefstab*. This file does not exist by default. Add entries into these files by using the following commands:

```
$ cat /etc/vx/tunefstab

/dev/vx/dsk/inputDG/inputVol
read_pref_io=131072,read_unit_io=131072,write_pref_io=131072,write_unit_io=131072

/dev/vx/dsk/archiveDG/archiveVol
read_pref_io=131072,read_unit_io=131072,write_pref_io=131072,write_unit_io=131072
```

Create these files on both the nodes.

I/O Fencing

I/O fencing is a mechanism to prevent uncoordinated access to the shared storage. It also works in the case of faulty cluster communications causing a split-brain condition.

Understanding Split Brain and the Need for I/O Fencing

To provide high availability, the cluster must be capable of taking corrective action when a node fails.

Problems arise when the mechanism that detects the failure breaks down because symptoms appear identical to those of a failed node. For example, if a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects and the remaining node takes corrective action. However, the failure of private interconnects (instead of the actual nodes) would present identical symptoms and cause each node to determine its peer has departed. This situation typically results in data corruption because both nodes attempt to take control of data storage in an uncoordinated manner.

I/O fencing is used to remove the risk associated with split brain. I/O fencing allows write access for members of the active cluster and blocks access to storage from non-members (even a node that is alive is unable to cause damage).

Configuring Disk-based I/O Fencing Using installsfcfs

Perform the following instructions as per the installer for the package.

Given below is a sample trace file:

```
# ./installsfcfs -fencing

Veritas Storage Foundation Cluster File System 5.1 SP1 Install Program

Copyright (c) 2010 Symantec Corporation. All rights reserved. Symantec, the
Symantec Logo are trademarks or registered trademarks of Symantec Corporation or
its affiliates in the U.S. and other countries. Other names may be trademarks of
their respective owners.

The Licensed Software and Documentation are deemed to be "commercial computer
software" and "commercial computer software documentation" as defined in FAR
Sections 12.212 and DFARS Section 227.7202.

Logs are being written to /var/tmp/installsfcfs-201409111116nHY while
installsfcfs is in progress.

Checking communication on extapps460-5 .....
Done

Checking release compatibility on extapps460-5 .....
Done

Checking VCS installation on extapps460-5 .....
5.1.100.000

Veritas Storage Foundation Cluster File System 5.1 SP1 Configure Program

Cluster information verification:

Cluster Name: iofencingcluster

Cluster ID Number: 9999

Systems: extapps460-5 extapps460-6

Would you like to configure I/O fencing on the cluster? [y,n,q] y

Checking communication on extapps460-5 .....
Done
```

```

Checking release compatibility on extapps460-5 .....
Done

Checking VCS installation on extapps460-5 .....
5.1.100.000

Checking communication on extapps460-6 .....
Done

Checking release compatibility on extapps460-6 .....
Done

Checking VCS installation on extapps460-6 .....
5.1.100.000

Veritas Storage Foundation Cluster File System 5.1 SP1 Configure Program

Fencing configuration

1) Configure CP client based fencing
2) Configure disk based fencing
3) Configure fencing in disabled mode

Select the fencing mechanism to be configured in this Application Cluster: [1-3,q] 2

Do you have SCSI3 PR enabled disks? [y,n,q] (y) y

Since you have selected to configure Disk based fencing, you would be asked to
give either the Disk group to be used as co-ordinator or asked to create disk
group and the mechanism to be used

Select one of the options below for fencing disk group:

1) Create a new disk group
b) Back to previous menu

Enter the choice for a disk group: [1-1,b,q] 1

List of available disks to create a new disk group

1) st2540-0_2
2) st2540-0_3
3) st2540-0_4
b) Back to previous menu

Select odd number of disks and at least three disks to form a disk group. Enter
the disk options, separated by spaces:

[1-3,b,q] 1 2 3

```

Enter the new disk group name: [b] fendg

Created disk group fendg

It is strongly recommended to run the 'VxFen Test Hardware' utility located at '/opt/VRTSvcs/vxfen/bin/vxfentsthdw' before continuing. The utility verifies if the shared storage you intend to use is able to support I/O fencing. Use the disk group you just selected for this verification. Come back here after you have completed the above step to continue with the configuration.

As per the 'vxfentsthdw' run you performed, do you want to continue with this disk group? [y,n,q] (y) y

Using disk group fendg

Enter fencing mechanism name (raw/dmp): [b,q,?] raw

Veritas Storage Foundation Cluster File System 5.1 SP1 Configure Program

I/O fencing configuration verification

Disk Group: fendg

Fencing mechanism: raw

Is this information correct? [y,n,q] (y)

Installer will stop VCS before applying fencing configuration. To make sure VCS shuts down successfully, unfreeze any frozen service group in the cluster.

Are you ready to stop VCS on all nodes at this time? [y,n,q] (n) y

Stopping VCS on extapps460-5
Done

Stopping Fencing on extapps460-5
Done

Stopping VCS on extapps460-6
Done

Stopping Fencing on extapps460-6
Done

Starting Fencing on extapps460-5
Done

Starting Fencing on extapps460-6
Done

Updating main.cf with fencing
Done

Starting VCS on extapps460-5
Done

```

Starting VCS on extapps460-6 .....
Done

I/O Fencing configuration .....
Done

I/O Fencing configuration completed successfully

installsfcfs log files, summary file, and response file are saved at:

/opt/VRTS/install/logs/installsfcfs-201409111116nHY

```

Perform the following prerequisite before selecting the option for this prompt **As per the 'vxfcntlsthdw' run you performed, do you want to continue with this disk group** that appears as part of the above mentioned procedure.

The following are sample logs while performing this procedure.

```
# /opt/VRTS/bin/vxfcntlsthdw
```

```

Veritas vxfcntlsthdw version 5.1.100.000-SP1GA LinuxThe utility vxfcntlsthdw works
on the two nodes of the cluster.The utility verifies that the shared storage one
intends to use isconfigured to support I/O fencing. It issues a series of
vxfenadmcommands to setup SCSI-3 registrations on the disk, verifies
theregistrations on the disk, and removes the registrations from the
disk.***** WARNING!!!!!!!!!! *****THIS UTILITY WILL DESTROY THE DATA ON THE
DISK!!

```

```
Do you still want to continue : [y/n] (default: n) y
```

```

The logfile generated for vxfcntlsthdw is
/var/VRTSvcS/log/vxfen/vxfcntlsthdw.log.22591

```

```
Enter the first node of the cluster:
```

```
extapps460-5
```

```
Enter the second node of the cluster:
```

```
extapps460-6
```

```
Enter the disk name to be checked for SCSI-3 PGR on node extapps460-5 in the
format:
```

```
for dmp: /dev/vx/rdmp/sdx
```

```
for raw: /dev/sdx
```

```
Make sure its the same disk as seen by nodes extapps460-5 and extapps460-6
```

```
/dev/vx/rdmp/st2540-0_2
```

```
Enter the disk name to be checked for SCSI-3 PGR on node extapps460-6 in the
format:
```

```
for dmp: /dev/vx/rdmp/sdx
```

```
for raw: /dev/sdx
```

Make sure its the same disk as seen by nodes extapps460-5 and extapps460-6

```
dev/vx/rmdp/st2540-0_2
```

```
*****
```

```
Testing extapps460-5 /dev/vx/rmdp/st2540-0_2 extapps460-6 /dev/vx/rmdp/st2540-0_2
```

```
Evaluate the disk before testing ..... No Pre-existing keys
```

```
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_2 from node extapps460-5 Passed
```

```
Verify registrations for disk /dev/vx/rmdp/st2540-0_2 on node extapps460-5 Passed
```

```
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_2 from node extapps460-6 Passed
```

```
Verify registrations for disk /dev/vx/rmdp/st2540-0_2 on node extapps460-6 Passed
```

```
Unregister keys on disk /dev/vx/rmdp/st2540-0_2 from node extapps460-5 . Passed
```

```
Verify registrations for disk /dev/vx/rmdp/st2540-0_2 on node extapps460-6 Passed
```

```
Unregister keys on disk /dev/vx/rmdp/st2540-0_2 from node extapps460-6 . Passed
```

```
Check to verify there are no keys from node extapps460-5 ..... Passed
```

```
Check to verify there are no keys from node extapps460-6 ..... Passed
```

```
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_2 from node extapps460-5 Passed
```

```
Verify registrations for disk /dev/vx/rmdp/st2540-0_2 on node extapps460-5 Passed
```

```
Read from disk /dev/vx/rmdp/st2540-0_2 on node extapps460-5 ..... Passed
```

```
Write to disk /dev/vx/rmdp/st2540-0_2 from node extapps460-5 ..... Passed
```

```
Read from disk /dev/vx/rmdp/st2540-0_2 on node extapps460-6 ..... Passed
```

```
Write to disk /dev/vx/rmdp/st2540-0_2 from node extapps460-6 ..... Passed
```

```
Reserve disk /dev/vx/rmdp/st2540-0_2 from node extapps460-5 ..... Passed
```

```
Verify reservation for disk /dev/vx/rmdp/st2540-0_2 on node extapps460-5 Passed
```

```
Read from disk /dev/vx/rmdp/st2540-0_2 on node extapps460-5 ..... Passed
```

```
Read from disk /dev/vx/rmdp/st2540-0_2 on node extapps460-6 ..... Passed
```

```
Write to disk /dev/vx/rmdp/st2540-0_2 from node extapps460-5 ..... Passed
```

```
Expect no writes for disk /dev/vx/rmdp/st2540-0_2 on node extapps460-6 . Passed
```

```
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_2 from node extapps460-6 Passed
```

```
Verify registrations for disk /dev/vx/rmdp/st2540-0_2 on node extapps460-5 Passed
```



```

Verify registrations for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-6 Passed
Write to disk /dev/vx/rdmp/st2540-0_2 from node extapps460-5 ..... Passed
Write to disk /dev/vx/rdmp/st2540-0_2 from node extapps460-6 ..... Passed
Preempt and abort key KeyA using key KeyB on node extapps460-6 ..... Passed
Test to see if I/O on node extapps460-5 terminated ..... Passed
RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_2 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-5 Passed
Preempt key KeyC using key KeyB on node extapps460-6 ..... Passed
Test to see if I/O on node extapps460-5 terminated ..... Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-5 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-6 Passed
Verify reservation for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-5 Passed
Verify reservation for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-6 Passed
Remove key KeyB on node extapps460-6 ..... Passed
Check to verify there are no keys from node extapps460-5 ..... Passed
Check to verify there are no keys from node extapps460-6 ..... Passed
Check to verify there are no reservations on disk /dev/vx/rdmp/st2540-0_2 from
node extapps460-5 Passed

Check to verify there are no reservations on disk /dev/vx/rdmp/st2540-0_2 from
node extapps460-6 Passed

RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_2 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-5 Passed
RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_2 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_2 on node extapps460-5 Passed
Clear PGR on node extapps460-5 ..... Passed
Check to verify there are no keys from node extapps460-5 ..... Passed
ALL tests on the disk /dev/vx/rdmp/st2540-0_2 have PASSED

The disk is now ready to be configured for I/O Fencing on node extapps460-5.

ALL tests on the disk /dev/vx/rdmp/st2540-0_2 have PASSED.

The disk is now ready to be configured for I/O Fencing on node extapps460-6.

```

```

Removing test keys and temporary files, if any...

# /opt/VRTS/bin/vxfentsthdw

Veritas vxfentsthdw version 5.1.100.000-SP1GA Linux

The utility vxfentsthdw works on the two nodes of the cluster.

The utility verifies that the shared storage one intends to use is
configured to support I/O fencing. It issues a series of vxfenadm
commands to setup SCSI-3 registrations on the disk, verifies the
registrations on the disk, and removes the registrations from the disk.

***** WARNING!!!!!!!!!! *****

THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!

Do you still want to continue : [y/n] (default: n) y

The logfile generated for vxfentsthdw is
/var/VRTSvcs/log/vxfen/vxfentsthdw.log.28186

Enter the first node of the cluster:

extapps460-5

Enter the second node of the cluster:

extapps460-6

Enter the disk name to be checked for SCSI-3 PGR on node extapps460-5 in the
format:

for dmp: /dev/vx/rdmp/sdx

for raw: /dev/sdx

Make sure its the same disk as seen by nodes extapps460-5 and extapps460-6

/dev/vx/rdmp/st2540-0_3

Enter the disk name to be checked for SCSI-3 PGR on node extapps460-6 in the
format:

for dmp: /dev/vx/rdmp/sdx

for raw: /dev/sdx

Make sure its the same disk as seen by nodes extapps460-5 and extapps460-6

/dev/vx/rdmp/st2540-0_3

*****

```

```

Testing extapps460-5 /dev/vx/rdmp/st2540-0_3 extapps460-6 /dev/vx/rdmp/st2540-0_3
Evaluate the disk before testing ..... No Pre-existing keys
RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_3 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-5 Passed
RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_3 from node extapps460-6 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-6 Passed
Unregister keys on disk /dev/vx/rdmp/st2540-0_3 from node extapps460-5 . Passed

Verify registrations for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-6 Passed
Unregister keys on disk /dev/vx/rdmp/st2540-0_3 from node extapps460-6 . Passed
Check to verify there are no keys from node extapps460-5 ..... Passed
Check to verify there are no keys from node extapps460-6 ..... Passed
RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_3 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-5 Passed
Read from disk /dev/vx/rdmp/st2540-0_3 on node extapps460-5 ..... Passed
Write to disk /dev/vx/rdmp/st2540-0_3 from node extapps460-5 ..... Passed
Read from disk /dev/vx/rdmp/st2540-0_3 on node extapps460-6 ..... Passed
Write to disk /dev/vx/rdmp/st2540-0_3 from node extapps460-6 ..... Passed
Reserve disk /dev/vx/rdmp/st2540-0_3 from node extapps460-5 ..... Passed
Verify reservation for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-5 Passed
Read from disk /dev/vx/rdmp/st2540-0_3 on node extapps460-5 ..... Passed
Read from disk /dev/vx/rdmp/st2540-0_3 on node extapps460-6 ..... Passed
Write to disk /dev/vx/rdmp/st2540-0_3 from node extapps460-5 ..... Passed
Expect no writes for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-6 . Passed
RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_3 from node extapps460-6 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-5 Passed
Verify registrations for disk /dev/vx/rdmp/st2540-0_3 on node extapps460-6 Passed
Write to disk /dev/vx/rdmp/st2540-0_3 from node extapps460-5 ..... Passed
Write to disk /dev/vx/rdmp/st2540-0_3 from node extapps460-6 ..... Passed

```

```

Preempt and abort key KeyA using key KeyB on node extapps460-6 ..... Passed
Test to see if I/O on node extapps460-5 terminated ..... Passed
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_3 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_3 on node extapps460-5 Passed
Preempt key KeyC using key KeyB on node extapps460-6 ..... Passed
Test to see if I/O on node extapps460-5 terminated ..... Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_3 on node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_3 on node extapps460-6 Passed
Verify reservation for disk /dev/vx/rmdp/st2540-0_3 on node extapps460-5 Passed
Verify reservation for disk /dev/vx/rmdp/st2540-0_3 on node extapps460-6 Passed
Remove key KeyB on node extapps460-6 ..... Passed
Check to verify there are no keys from node extapps460-5 ..... Passed
Check to verify there are no keys from node extapps460-6 ..... Passed
Check to verify there are no reservations on disk /dev/vx/rmdp/st2540-0_3 from
node extapps460-5 Passed

Check to verify there are no reservations on disk /dev/vx/rmdp/st2540-0_3 from
node extapps460-6 Passed

RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_3 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_3 on node extapps460-5 Passed
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_3 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_3 on node extapps460-5 Passed
Clear PGR on node extapps460-5 ..... Passed
Check to verify there are no keys from node extapps460-5 ..... Passed
ALL tests on the disk /dev/vx/rmdp/st2540-0_3 have PASSED.

The disk is now ready to be configured for I/O Fencing on node extapps460-5.

ALL tests on the disk /dev/vx/rmdp/st2540-0_3 have PASSED.

The disk is now ready to be configured for I/O Fencing on node extapps460-6.

Removing test keys and temporary files, if any...

```

```
# /opt/VRTS/bin/vxfentsthdw
```

```
Veritas vxfentsthdw version 5.1.100.000-SP1GA Linux
```

The utility `vxfentsthdw` works on the two nodes of the cluster. The utility verifies that the shared storage one intends to use is configured to support I/O fencing. It issues a series of `vxfenadm` commands to setup SCSI-3 registrations on the disk, verifies the registrations on the disk, and removes the registrations from the disk.

***** WARNING!!!!!!!!!! *****

THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!

Do you still want to continue : [y/n] (default: n) y

The logfile generated for `vxfentsthdw` is
`/var/VRTSvcS/log/vxfen/vxfentsthdw.log.27461`

Enter the first node of the cluster:

`extapps460-5`

Enter the second node of the cluster:

`extapps460-6`

Enter the disk name to be checked for SCSI-3 PGR on node `extapps460-5` in the format:

for dmp: `/dev/vx/rdmp/sdx`

for raw: `/dev/sdx`

Make sure its the same disk as seen by nodes `extapps460-5` and `extapps460-6`

`/dev/vx/rdmp/st2540-0_4`

Enter the disk name to be checked for SCSI-3 PGR on node `extapps460-6` in the format:

for dmp: `/dev/vx/rdmp/sdx`

for raw: `/dev/sdx`

Make sure its the same disk as seen by nodes `extapps460-5` and `extapps460-6`

`/dev/vx/rdmp/st2540-0_4`

Enter the disk name to be checked for SCSI-3 PGR on node `extapps460-6` in the format:

for dmp: `/dev/vx/rdmp/sdx`

for raw: `/dev/sdx`

Make sure its the same disk as seen by nodes `extapps460-5` and `extapps460-6`

`/dev/vx/rdmp/st2540-0_4`

```

*****

Testing extapps460-5 /dev/vx/rdmp/st2540-0_4 extapps460-6 /dev/vx/rdmp/st2540-0_4

Evaluate the disk before testing ..... No Pre-existing keys

RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_4 from node extapps460-5 Passed

Verify registrations for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-5 Passed

RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_4 from node extapps460-6 Passed

Verify registrations for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-6 Passed

Unregister keys on disk /dev/vx/rdmp/st2540-0_4 from node extapps460-5 . Passed

Verify registrations for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-6 Passed

Unregister keys on disk /dev/vx/rdmp/st2540-0_4 from node extapps460-6 . Passed

Check to verify there are no keys from node extapps460-5 ..... Passed

Check to verify there are no keys from node extapps460-6 ..... Passed

RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_4 from node extapps460-5 Passed

Verify registrations for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-5 Passed

Read from disk /dev/vx/rdmp/st2540-0_4 on node extapps460-5 ..... Passed

Write to disk /dev/vx/rdmp/st2540-0_4 from node extapps460-5 ..... Passed

Read from disk /dev/vx/rdmp/st2540-0_4 on node extapps460-6 ..... Passed

Write to disk /dev/vx/rdmp/st2540-0_4 from node extapps460-6 ..... Passed

Reserve disk /dev/vx/rdmp/st2540-0_4 from node extapps460-5 ..... Passed

Verify reservation for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-5 Passed

Read from disk /dev/vx/rdmp/st2540-0_4 on node extapps460-5 ..... Passed

Read from disk /dev/vx/rdmp/st2540-0_4 on node extapps460-6 ..... Passed

Write to disk /dev/vx/rdmp/st2540-0_4 from node extapps460-5 ..... Passed

Expect no writes for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-6 . Passed

RegisterIgnoreKeys on disk /dev/vx/rdmp/st2540-0_4 from node extapps460-6 Passed

Verify registrations for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-5 Passed

Verify registrations for disk /dev/vx/rdmp/st2540-0_4 on node extapps460-6 Passed

Write to disk /dev/vx/rdmp/st2540-0_4 from node extapps460-5 ..... Passed

Write to disk /dev/vx/rdmp/st2540-0_4 from node extapps460-6 ..... Passed

```

```

Preempt and abort key KeyA using key KeyB on node extapps460-6 ..... Passed
Test to see if I/O on node extapps460-5 terminated ..... Passed
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_4 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_4 on node extapps460-5 Passed
Preempt key KeyC using key KeyB on node extapps460-6 ..... Passed
Test to see if I/O on node extapps460-5 terminated ..... Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_4 on node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_4 on node extapps460-6 Passed
Verify reservation for disk /dev/vx/rmdp/st2540-0_4 on node extapps460-5 Passed
Verify reservation for disk /dev/vx/rmdp/st2540-0_4 on node extapps460-6 Passed
Remove key KeyB on node extapps460-6 ..... Passed
Check to verify there are no keys from node extapps460-5 ..... Passed
Check to verify there are no keys from node extapps460-6 ..... Passed
Check to verify there are no reservations on disk /dev/vx/rmdp/st2540-0_4 from
node extapps460-5 Passed

Check to verify there are no reservations on disk /dev/vx/rmdp/st2540-0_4 from
node extapps460-6 Passed

RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_4 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_4 on node extapps460-5 Passed
RegisterIgnoreKeys on disk /dev/vx/rmdp/st2540-0_4 from node extapps460-5 Passed
Verify registrations for disk /dev/vx/rmdp/st2540-0_4 on node extapps460-5 Passed
Clear PGR on node extapps460-5 ..... Passed
Check to verify there are no keys from node extapps460-5 ..... Passed
ALL tests on the disk /dev/vx/rmdp/st2540-0_4 have PASSED.

The disk is now ready to be configured for I/O Fencing on node extapps460-5.

ALL tests on the disk /dev/vx/rmdp/st2540-0_4 have PASSED.

The disk is now ready to be configured for I/O Fencing on node extapps460-6.

Removing test keys and temporary files, if any...

```

Verifying I/O Fencing Configuration

Perform the following to verify I/O fencing configuration.

Given below is an example of one of the node types:

```
[root@extapps460-5 bin]# ./vxfenadm -d

I/O Fencing Cluster Information:

=====

Fencing Protocol Version: 201

Fencing Mode: SCSI3

Fencing SCSI3 Disk Policy: raw

Cluster Members:

* 0 (extapps460-5)

1 (extapps460-6)

RFSM State Information:

node 0 in state 8 (running)

node 1 in state 8 (running)
```

Testing I/O Fencing Configuration

To simulate a split brain scenario, the following sample commands are used on one of the nodes to disable heartbeat link.


```
lltconfig -t eth4 -L 0

lltconfig -t eth5 -L 0

lltconfig -t eth2 -L 0
```

The option **ethX** used in the above commands needs to be replaced with the actual details to suit customer setup.

One of the nodes is ejected from the cluster. The following is a screenshot of the console.



```
LLT INFO U-14-1-18518 sent hbreq (NULL) on link 2 (eth2) node 0. 1 more to go.
LLT INFO U-14-1-18518 sent hbreq (NULL) on link 2 (eth2) node 0. 0 more to go.
LLT INFO U-14-1-18518 link 2 (eth2) node 0 inactive 15 sec (189999)
LLT INFO U-14-1-18518 link 2 (eth2) node 0 expired
GAB INFO U-15-1-28836 Port h gen 7f8237 membership :1
GAB INFO U-15-1-28836 Port u gen 7f8244 membership :1
GAB INFO U-15-1-28836 Port v gen 7f8248 membership :1
GAB INFO U-15-1-28836 Port f gen 7f823a membership :1
GAB INFO U-15-1-28836 Port b gen 7f822d membership :1
ofs: msgcat 15 I/O fencing is ON
GAB INFO U-15-1-28836 Port w gen 7f8242 membership :1
GAB INFO U-15-1-28836 Port a gen 7f8282 membership :1
OXFEN WARNING U-11-1-65 Could not eject node 0 from disk
with serial number 68888f588824613c888886c53e88a8 since
keys of node 1 are not registered with it
OXFEN WARNING U-11-1-65 Could not eject node 0 from disk
with serial number 68888f588824613c888886c53e88a8 since
keys of node 1 are not registered with it
OXFEN WARNING U-11-1-65 Could not eject node 0 from disk
with serial number 68888f588824613c888886c53e88b11 since
keys of node 1 are not registered with it
Kernel panic - not syncing: OXFEN CRITICAL U-11-1-28 Local cluster node
ejected from cluster to prevent potential data corruption.
```

The following lines are found in log file “engine_A.log” located in the directory */var/VRTSvc/log/*.

2014/09/25 16:56:39 VCS WARNING V-16-1-11141 LLT heartbeat link status changed.
Previous status = eth4 UP

eth5 UP eth2 UP; Current status = eth4 DOWN eth5 UP eth2 UP.

2014/09/25 16:57:29 VCS INFO V-16-1-10077 Received new cluster membership

2014/09/25 16:57:29 VCS NOTICE V-16-1-10112 System (extapps460-5) - Membership:
0x3, DDNA: 0x2

2014/09/25 16:57:29 VCS ERROR V-16-1-10111 System extapps460-6 (Node '1') is in
Regular and Jeopardy

Memberships - Membership: 0x3, Jeopardy: 0x2

2014/09/25 16:57:34 VCS WARNING V-16-1-11141 LLT heartbeat link status changed.
Previous status = eth4

DOWN eth5 UP eth2 UP; Current status = eth4 DOWN eth5 DOWN eth2 UP.

2014/09/25 17:00:56 VCS INFO V-16-1-10077 Received new cluster membership

2014/09/25 17:00:56 VCS NOTICE V-16-1-10112 System (extapps460-5) - Membership:
0x1, DDNA: 0x0

2014/09/25 17:00:56 VCS NOTICE V-16-1-10034 RECONFIG received. VCS waiting for
I/O fencing to be completed

2014/09/25 17:00:57 VCS NOTICE V-16-1-10036 I/O fencing completed

2014/09/25 17:00:57 VCS ERROR V-16-1-10079 System extapps460-6 (Node '1') is in
Down State - Membership: 0x1

2014/09/25 17:00:57 VCS ERROR V-16-1-10322 System extapps460-6 (Node '1') changed
state from RUNNING to FAULTED

2014/09/25 17:00:57 VCS NOTICE V-16-1-10446 Group cvm is offline on system
extapps460-6

2014/09/25 17:01:03 VCS WARNING V-16-1-11141 LLT heartbeat link status changed.
Previous status = eth4

DOWN eth5 DOWN eth2 UP; Current status = eth4 DOWN eth5 DOWN eth2 DOWN.



Important: Once the RDP is fenced off to avoid data corruption, the fenced off RDP will not be able to read or write anything on the shared disk and the RDP in this case will not be able to process the files.

Manual intervention is needed to get the fenced off RDP online.

Fix the fenced off RDP with the help of the system administrator or Symantec TAC Support and get it back online for the service.

Configurations in MUR for Clustering Support

MUR configuration for Scalability involves installation of individual RDP and a Master.

There is no special action required during the installation of MUR. However, care has to be taken while assigning the paths of each component.

- MUR application path has to be on internal disks.
- MUR database path has to be on internal disks.
- MUR archival path can be on internal or external path designated at the time of sizing.
- Choose the component as RDP / Master depending on the node of cluster.
- Choose the same user id, user name, and user group on all the RDP nodes so that there is no permission related errors while dealing with input files and exported files between RDP and MUR.

There are special recommendations for configuring gateways and RDP under Scalability deployment.

Configuring MUR for Use in Scalable model

The following section provides the step-by-step guidelines for these configurations.

1. Configuring Region:

If for each GGSN, separate region needs to be configured than "NOC" then add necessary regions per GGSN.

Note that all the RDPs which are going to be used to support single GGSN has to be configured under same region.

2. Adding RDP:

Add all the RDPs required for supporting the throughput of the GGSNs.


Here, choose the region for the RDP if needed. There is no option to configure it later.

The screenshot shows the 'Update RDP' configuration window. On the left is a sidebar menu with options like GATEWAYS, RDP, USERS, GROUPS, DATA, LOGGING, E-MAIL, AUDIT, PURGING, BACKUP, BULKSTATS, KPI ALARMS, and REPORT GRANULARITY. The 'RDP' option is selected. The main form has the following fields:

- Ipaddress:** 10.105.81.41 (RDP Host IPAddress)
- Is SSL Enabled:** ☐ (Check if RDP is on HTTPS)
- Port:** 8080 (RDP json port)
- Name:** RDP2_pnqaextapps460 (RDP name)
- Region:** NOC (Associated Region)
- Location:** pune (RDP location)
- Timezone:** Asia/Calcutta (GMT+0530)

At the bottom are 'Update' and 'Cancel' buttons. A 'Show RDP' button is in the top right corner of the form area. The page number 381323 is visible in the bottom right corner.

The following screen shows all the four RDPs added for a 40 Gbps throughput.



3. Adding Gateways:

The way scalability is supported is by having one primary Gateway and other remaining gateways as virtual alias of this primary gateway called as pseudo gateway.

The primary gateway can be added on either of the serving RDPs and then pseudo gateways should be added on rest of the RDPs.

Number of pseudo gateways required depends on the predefined load balancing strategy which is conducted before deployment as part of sizing.

Primary gateway should be assigned the actual IP address of the GGSN and rest of the pseudo gateways can be assigned non-existing or dummy IP addresses. Apart from bulkstat related configuration, RDP never accesses ASR node. Since bulkstat files would always be processed on Master through primary RDP, other nodes can be reluctant of bulkstat configuration.

The following screens indicate how a primary gateway and pseudo gateways are added to MUR.

CISCO

Gateway Settings

Gateways

Users

Groups

Data

Logging

E-Mail

Alert

Purging

Backup

BulkStats

WLANs

Report Granularity

Home

UPI

CF

HTTP

BulkStats

RPI

Favorites

Background Tasks

Admin

Welcome admin | Change password | Help | About | System | Log out

Update Gateway

Show Gateway List

Name

gw1_pne+appsruce460-

Required: Name of the gateway

IP address

10.105.89.211

Required: Gateway IP Address

Netmask

255.255.255.0

Required: Gateway Netmask

Pseudo Gateway

☐

Is this Pseudo to an existing gateway?

Existing Gateway of which this gateway is a pseudo:

Gateway Type

OSGN

Associated RDP Master

RDPI_pne+appsruce460-1

Required: RDP Master for Gateway

Associated Region

NOC

Required: associated region for Gateway

Location

pune

Timezone

Asia/Calcutta (GMT+0530)

Flow-Edr

☒

Support for Flow-Edr Data Format

Flow-Edr Filename Pattern

flow*[0-1].gz

Required: Flow-Edr Filename Pattern e.g. "flow"

Flow-Edr Filename Path

/input_data/data

NOTE: Please cross-check the incoming flow-edr filename path in case of RDP

Http-Edr

☒

Support for Http-Edr Data Format

Http-Edr Filename Pattern

http*[0-1].gz

Required: Http-Edr Filename Pattern e.g. "http"

Http-Edr Filename Path

/input_data/data

NOTE: Please cross-check the incoming http-edr filename path in case of RDP

CF-Edr

☐

Support for CF-Edr Data Format

CF-Edr Filename Pattern

Required: CF-Edr Filename Pattern e.g. "cfedr"

CF-Edr Filename Path

NOTE: Please cross-check the incoming cf-edr filename path in case of RDP

Edr Filename Format

STR_STR_YYYYMMDDHHmmSS_NUM_NUM

Edr Filename Format, it is applicable for all i.e. Flow-Edr, Http-Edr and CF-Edr

BulkStats

☐

Support for BulkStats

BS Filename Pattern

Required: BulkStats Filename Pattern e.g. "bulkstat"

BS Filename Path

Required: BulkStats Filename Path

Update

Cancel

381319

Mobility Unified Reporting

[Welcome admin](#)
[Change password](#)
[Help](#)
[About](#)
[System](#)
[Log out](#)

[HOME](#)
[DPI](#)
[CF](#)
[HTTP](#)
[BULKSTATS](#)
[KPI](#)
[FAVOURITES](#)
[BACKGROUND TASKS](#)
[ADMIN](#)

GATEWAYS
RDP
USERS
GROUPS
DATA
LOGGING
EMAIL
AUDIT
PURGING
BACKUP
BULKSTATS
KPI ALARMS
REPORT GRANULARITY

Update Gateway

Show Gateway List

Name:
Required. Name of the gateway.

IP address:
Required. Gateway IP Address.

Netmask:
Required. Gateway Netmask.

Pseudo Gateway: ☒
Is this Pseudo to an existing gateway?

Existing Gateway of which this gateway is a pseudo:

Gateway Type:

Associated RDP Master:
RDP Master for Gateway

Associated Region:
Required, associated region for Gateway

Location:

Timezone:

Flow Edit: ☒
Support for Flow-Edit Data Format

Flow Edit Filename Pattern:
Incoming Flow-Edit Filename Pattern e.g. "flow"

Flow Edit Filename Path:
NOTE: Please cross-check the incoming flow-edit filename path in case of RDP

Http Edit: ☒
Support for Http-Edit Data Format

Http Edit Filename Pattern:
Incoming Http-Edit Filename Pattern e.g. "http"

Http Edit Filename Path:
NOTE: Please cross-check the incoming http-edit filename path in case of RDP

CF Edit: ☐
Support for CF-Edit Data Format

CF Edit Filename Pattern:
Incoming CF-Edit Filename Pattern e.g. "cfedit"

CF Edit Filename Path:
NOTE: Please cross-check the incoming cf-edit filename path in case of RDP

Edit Filename Format:
Edit Filename Format, % is applicable for all i.e. Flow-Edit, Http-Edit and CF-Edit

Bulkstats: ☐
Support for BulkStat

BS Filename Pattern:
Incoming BulkStats Filename Pattern e.g. "bulkstat"

BS Filename Path:
Incoming BulkStats Filename Path

381320

CISCO

MOBILITY UNIFIED REPORTING

WELCOME admin | CHANGE PASSWORD | HELP | ABOUT | SYSTEM | LOG OUT

GATEWAYS

RDP

USERS

GROUPS

DATA

LOGGING

E-MAIL

AUDIT

PURGING

BACKUP

BULKSTATS

WIRALARMS

REPORT GRANULARITY

HOME

UPI

CF

HTTP

BULKSTATS

RPI

FAVOURITES

BACKGROUND TASKS

ADMIN

Update Gateway

Show Gateway List

Name: pseudo2_460-7

Required: Name of the gateway

IpAddress: 10.105.81.103

Required: Gateway IP Address

Netmask: 255.255.255.0

Required: Gateway Netmask

Pseudo Gateway: ☒

Is this Pseudo to an existing gateway?

Existing Gateway of which this gateway is a pseudo: gw1_pnextappsucc460-1

Gateway Type: OOSN

Associated RDP Master: RDP3_inqandapps460-7

RDP Master for Gateway

Associated Region: NOC

Required, associated region for Gateway

Location: pune

Timezone: Asia/Calcutta (GMT+0530)

Flow Edr: ☒

Support for Flow-Edr Data Format

Flow-Edr Filename Pattern: %flow%[%4s] gz

Incoming Flow-Edr Filename Pattern e.g. "flow"

Flow-Edr Filename Path: /input_data/data

NOTE: Please cross-check the incoming flow-edr filename path in case of RDP

Http Edr: ☒

Support for Http-Edr Data Format

Http-Edr Filename Pattern: %http%[%4s] gz

Incoming Http-Edr Filename Pattern e.g. "http"

Http-Edr Filename Path: /input_data/data

NOTE: Please cross-check the incoming http-edr filename path in case of RDP

CF Edr: ☐

Support for CF-Edr Data Format

CF-Edr Filename Pattern:

Incoming CF-Edr Filename Pattern e.g. "chdr"

CF-Edr Filename Path:

NOTE: Please cross-check the incoming cf-edr filename path in case of REP

Edr Filename Format: STR_STR_YYYYMMDDHHmmSS_NUM_NUM

Edr Filename Format, it is applicable for all i.e. Flow-Edr, Http-Edr and CF-Edr

Bulkstats: ☐

Support for BulkStat

BS Filename Pattern:

Incoming BulkStat Filename Pattern e.g. "bulkstat"

BS Filename Path:

Incoming BulkStat Filename Path

Update


Cancel

381321

Cisco Mobility Unified Reporting System Installation and Administration Guide

398

OL-27216-09



Update Gateway

Name: pseudo_460-5
Required. Name of the gateway.

IP Address: 10.105.89.209
Required. Gateway IP Address.

Netmask: 255.255.255.0
Required. Gateway Netmask.

Pseudo Gateway: ☒
Is this Pseudo to an existing gateway?

Existing Gateway of which this gateway is a pseudo: gw1_pnqaetapps460-1

Gateway Type: GGSN

Associated RDP Master: RDP4_pnqaetapps460-5
RDP Master for Gateway

Associated Region: NOC
Required, associated region for Gateway

Location: pune

Timezone: Asia/Calcutta (GMT+0530)

Flow Edit: ☒
Support for Flow-Edit Data Format

Flow-Edit Filename Pattern: ^flow-[1-9].gz
Incoming Flow-Edit Filename Pattern e.g. ^flow*

Flow-Edit Filename Path: /input_data/data
NOTE: Please cross-check the incoming flow-edit filename path in case of RDP

Http Edit: ☒
Support for Http-Edit Data Format

Http-Edit Filename Pattern: ^http-[1-9].gz
Incoming Http-Edit Filename Pattern e.g. ^http*

Http-Edit Filename Path: /input_data/data
NOTE: Please cross-check the incoming http-edit filename path in case of RDP

CI Edit: ☐
Support for CI-Edit Data Format

CI-Edit Filename Pattern:
Incoming CI-Edit Filename Pattern e.g. ^ciedit*

CI-Edit Filename Path:
NOTE: Please cross-check the incoming ci-edit filename path in case of RDP

Edit Filename Format: STR_STR_YYYYMMDD+mmSS_NUM_NUM
Edit Filename Format, it is applicable for all i.e. Flow-Edit, Http-Edit and CI-Edit

BulkStats: ☐
Support for BulkStat

BS Filename Pattern:
Incoming BulkStats Filename Pattern e.g. bulkstat*

BS Filename Path:
Incoming BulkStats Filename Path

Buttons: [Update] [Cancel]

381322



Gateways

Name	IP Address	Netmask	Pseudo's Gateway	Associated RDP	Region	Type	Timezone	Action
gw1_pnqaetapps460-1	10.105.89.211	255.255.255.0		RDP1_pnqaetapps460-1	NOC	GGSN	Asia/Calcutta	[Edit] [Delete]
pseudo_460-2	10.105.81.41	255.255.255.0	gw1_pnqaetapps460-1	RDP2_pnqaetapps460-2	NOC	GGSN	Asia/Calcutta	[Edit] [Delete]
pseudo_460-7	10.105.81.103	255.255.255.0	gw1_pnqaetapps460-1	RDP3_pnqaetapps460-7	NOC	GGSN	Asia/Calcutta	[Edit] [Delete]
pseudo_460-5	10.105.89.209	255.255.255.0	gw1_pnqaetapps460-1	RDP4_pnqaetapps460-5	NOC	GGSN	Asia/Calcutta	[Edit] [Delete]

381325

For configuring the pseudo gateway it has to be marked as pseudo and should be specified for which gateway (primary) it is going to act as pseudo.

Also the filename pattern should be regular expression chosen such that together this set would form an equal division of all the files across added RDPs.

Input file path for each of the gateway has to be the shared drive path where ASR is pushing the files.

4. Support Bulkstats:

Bulkstats has to be configured only on primary gateway. All the bulkstat files are processed by Master node only.

5. Configuring ASR to push files:

The destination host as a receiver node for ASR can be only one RDP. So the data from ASR node should be pushed only through either of the IP addresses of the RDP machine. However, the path where files should be pushed has to be the shared drive path which is also configured as input file path on each of the gateways.

6. Viewing reports:

All the reports comprising pseudo gateways and primary gateway would appear under primary gateway.

Post Installation Actions

Perform the following necessary and highly recommended configuration for prevention and recognition of faults in MUR.

- Configuring SNMP alarms in MUR
- Configuring backup in MUR
- Configuring purging of database and files in MUR
- Configuring Archiving in MUR

Installation / Upgrade to Scalable Model

This section discusses the installation and upgrade procedure for MUR in scalable model under various deployment cases.

In this section, the following cases are considered:

- Current standalone deployment, upgrade to scalable model
- Current hierarchical deployment, upgrade to scalable model.

Upgrade from Standalone Deployment to Scalable Model

This section provides instructions on how to upgrade from the standalone deployment of MUR application to scalable model.

 **Important:** Whenever an MUR is upgraded from an older version to 12.2, the logical regions for the MUR gets void and the user will not be able to see any gateways under the **DPI/Bulkstats/CF/KPI** tab. In this scenario, the user should add that gateway under NOC.

It is expected that the user must first upgrade the MUR application to a hierarchical model before deploying MUR in scalable model.

- Step 1** Consider the current standalone node as a master. Then, upgrade the software to latest MUR version.
- Step 2** Install two or more RDPs in scalable mode.
- Step 3** Retain the existing gateway as it is (attached to Master) if you are interested in viewing existing gateway reports. Else, remove the gateway through the Gateway Configuration screen on the GUI.
- Step 4** Configure the same gateway (with different name and IP address) through master for one of the RDPs and mark as pseudo to another RDP.

Rest of the configuration remains the same.

- Step 5** Reconfigure the gateway to push data to one of the RDPs. In the earlier configuration, the gateway pushed the files to the standalone server (now it is master MUR).

If user is interested in viewing data for the earlier configured gateway, the gateway should remain in the configuration. After upgrading to scalable model, the following procedure should be performed:

- Invoke MUR GUI. On the home page, even if there were gateways configured earlier, the following message will be displayed “no gateways added”.
- Click **GATEWAY** from **Admin** menu, and then edit each of the gateways to have the region as NOC.

All the reports will now be available to users.

 **Important:** Note that there might be some files that would remain on the old master instance and would get processed there. So, during the upgrade process, some of the data might be reported on old master instance and the old metadata configuration changes would not be reflected on the RDPs.

Upgrade from Hierarchical Deployment to Scalable Model

This section provides instructions on how to upgrade from the hierarchical deployment of MUR application to scalable model.



Important: Whenever an MUR is upgraded from an older version to 12.2, the logical regions for the MUR gets void and the user will not be able to see any gateways under the **DPI/Bulkstats/CF/KPI** tab. In this scenario, the user should add that gateway under NOC.

It is assumed that there is at least one RDP attached to the master. Follow the steps as described here.

- Step 1** Stop the MUR on both master and RDP-1.
- Step 2** Upgrade MUR on both of them independently to the latest version.
- Step 3** Create appropriate regions if required. Then, update RDP-1 and gateway configuration appropriately from master MUR.
- Step 4** Setup a new node as RDP-2 and configure it in cluster mode, attached to a shared storage.
While configuring new node [RDP-2] make sure that you specify the id of the node-1 (older RDP) during cluster installation so that cluster will get installed on node-1 also.
Make sure that node-1 and node-2 are physically interconnected to have cluster functionality. Also, make sure the share disk is created and mounted as described in the [Basic Scalability Model](#) section.
- Step 5** Install the latest MUR on RDP-2. Make sure that the parameters for user id, and user name are the same as that of MUR on RDP-1.
- Step 6** Reconfigure the gateway to push traffic through this new node RDP-2 on the shared disk.
- Step 7** Add this gateway as pseudo gateway on RDP-2 through master. This would be a pseudo to the gateway on RDP-1. Initially, configure the gateway to pick up all the files without setting any filters.
Files will keep on getting processed on this RDP-2.
- Step 8** On the old RDP-1, check if all the files have been parsed. If yes, check if there are any files pending in the `<starbi home>/server/sql_export_data` directory. When there are no files, and you have waited for 20 minutes after the files have been parsed, take a backup of all the metadata so that RDP can be restored later, and stop this RDP.
- Step 9** Now uninstall the MUR on this server (RDP-1). Bring the node RDP-1 in cluster, and re-install RDP on it with the same parameters with which it was installed (installation path, user id, postgres id, postgres port, postgres password, apache port, RPC port, etc.).
- Step 10** Perform RDP recovery procedure and then reconfigure the gateway from master on the RDPs to take the files from appropriate path with required filters.
- Step 11** If there are more than one RDP, repeat the above procedure for other RDPs. Then, edit the earlier gateway configuration to pick up only odd (or as desired) numbered files on RDP-2.
- Step 12** Add another gateway, but as a pseudo to earlier one, on RDP-2. Configure it to pick up only even numbered files (or as desired) so that the gateway traffic is evenly distributed across the RDPs.
- Step 13** Repeat the above procedure for other RDPs.



Important: The same RDP can be restored if they are reinstalled with the same parameters.



Important: Note that there might be some files that would remain on the old RDP instance and would get processed there. So, during the upgrade process, some of the data might be reported on old RDP instance and the old metadata configuration changes would not be reflected on the RDPs.

Chapter 9

Troubleshooting MUR System

This chapter provides information on how to resolve situations you might encounter with using MUR software. This chapter provides problem definitions, their likely cause(s), and solutions.

This chapter describes the following topics:

- [MUR Preventive and Control Measures](#)
- [Install and Upgrade Pre-requisites](#)
- [Issues Pertaining to Installation Upgrade](#)
- [Issues Pertaining to MUR Processes](#)
- [Issues Related to MUR GUI](#)
- [Issues Related to Master-RDP communication](#)
- [Issues Related to BS-KPI-Alarms](#)
- [Issues Related to Backup and Restore](#)
- [Miscellaneous](#)

MUR Preventive and Control Measures

This section highlights the preventive measures which will help to control MUR breakdown and also manages to restore the MUR functionality back in case of failures.

- SNMP

MUR has a self-monitoring capability designed to monitoring the following conditions at periodic intervals and raise an SNMP alert (to a configured Management entity/entities), upon violation of pre-configured thresholds.

- CPU usage
- Disk usage (thresholds can be configured)
- RAM usage (swap)
- Core processors restarting or going down (except Postgres)
- Monitoring of unprocessed files (EDRs and bulkstats)
- Upon encountering syntactical errors in EDR files (e.g. missing headers in incoming EDR file or issues with uncompressing the files received from ASR5K)
- Inordinate delay experienced by certain scripts e.g. parsing and aggregation scripts

This feature runs independently on the RDP and Master servers. At present, this feature has a limitation in the sense that it cannot raise alerts in the event of Postgres outage itself.

- Email notifications

MUR also has the capability to send emails to pre-configured email group. This feature is used when monitoring the KPI thresholds, absence of the incoming EDR files, etc.

In addition, on enabling SNMP alerts, MUR server also monitors disk space overruns (i.e. disk full scenario) and in the event of this situation it suspends the scheduler operation (i.e. all scripts executions). MUR sends an email notification of this suspension. Upon disk space recovery the scheduler operation is recommended.

- Database purging

MUR purges reporting data as well as archived EDR files as per the data retention configuration. This way it will restrict MUR database or archive partition not to become 100% full.

- Database backup and recovery

MUR purges reporting data as well as archived EDR files as per the data retention configuration. This way it will restrict MUR database or archive partition not to become 100% full.

For additional details on these features, refer to *Mobility Unified Reporting System Administration and Management* chapter of this guide.

MUR Install and Upgrade Prerequisites

These are the checks to be done before proceeding with installation/upgrade.

- Check if shared memory settings are as per the recommendation:
kernel.shmmax=10737418240
kernel.shmall=4294967296
- Check if fstab entries are proper so that reboot of server should not cause any damage to the installation.
- Check if disk partitioning and sizes are as per the recommendation.
- Check if time zone is properly set on the system. List of valid time zones is as below:

<http://www.vmware.com/support/developer/vc-sdk/visdk400pubs/ReferenceGuide/timezone.html>

Refer to the MOP to be followed after upgrade under "Upgrading MUR" section in this guide.

Refer to the "Release Note" section before proceeding with Install / Upgrade.

Issues Pertaining to MUR Installation / Upgrade

Problem:	Install / Upgrade failure
Action(s):	Please perform pre-requisite checks mentioned above.

Problem:	Offline Upgrade Fails in Master and RDP.
Possible Cause(s):	The aggregation failures are caused due to lack of the column 'ossign' in tables 'sbi_edr_daily_aggr_user_agent_%' or 'sbi_edr_http_summary_'.
Action(s):	Perform the following steps: 1. Stop MUR by executing the following command from the <code><MUR_Install_Path>/starbi/bin/</code> directory: <code>serv stop</code> 2. Start the Postgres server by executing the following command from the <code><MUR_Install_Path>/starbi/bin/</code> directory: <code>serv postgres start</code> 3. Connect to the Postgres server by executing the following command from the <code><MUR_Install_Path>/starbi/postgres/bin</code> directory: <code>psql -p <port> -U <postgres-user> starbidb</code> 4. Execute the following query: <pre>"delete from sbi_edr_http_script_state where status='f' and (tablename like 'sbi_edr_daily_aggr_user_agent_%' or tablename like 'sbi_edr_http_summary_');"</pre> 5. Stop the Postgres server. 6. Start the upgrade procedure.

Issues Pertaining to MUR Processes

Problem:	Parser server is not running after upgrade.
Possible Cause(s):	When upgrading from 12.0 to any newer version, file parsing configurations are NOT synced automatically at all RDPs. As a result, EDR file parsing does not happen at RDPs.
Action(s):	Please see if post-upgrade MoP is followed properly or not for RDP upgrades. - After the upgrade, manually attach appropriate RDPs to all relevant gateways. - Create appropriate RDP regions through the System menu. - Attach all RDPs to their respective regions through Edit RDP page viewed by clicking RDP from Admin tab. - Attach appropriate gateways to their corresponding RDPs and regions through Edit Gateway page viewed by clicking GATEWAYS from Admin tab. - Navigate to System > File-Parsing Configs menu and then manually save the file parsing configurations for all the gateways which are attached to RDPs.

Problem:	No reports are generated.
Possible Cause(s):	This may be due to file processing issue. File processing issues may be: - Insufficient file/directory permissions. - Incorrect input file path / filename format in gateway configuration.
Action(s):	- Use appropriate user ID while pushing the files in input directory. - Correct the gateway configurations.

Problem:	Reports are delayed by few days.
Possible Cause(s):	This may be a result of aggregation backlog which may be caused due to: - Postgres errors are observed for some tables resulting in aggregation backlog. - EDR file size and number of EDRs coming to MUR might be not as per the recommendation.
Action(s):	MUR tuning or GGSN configuration change should be performed as per the recommended rotation time (300 seconds) and rotation volume (40MB).

Problem:	Scheduler found to stop working (from scheduler logs)
Possible Cause(s):	Scheduler is found to throw Exception: <Fault 0: 'Java heap space'>" Scheduler stops working and results in stopping MUR processing.
Action(s):	Restart the MUR application, this brings the system back in condition and starts MUR reporting.

Problem:	The EDR files are generated and moved out from the input directory. However, there are no reports getting generated.
----------	--

Possible Cause(s):	- The files may not be available in the archive directory i.e. <code><MUR_install_dir>/starbi/archive</code>. - The incoming files might not be correct.
Action(s):	- Check if the files are available in the archive directory. - Check if they are marked invalid. If yes, check if there are any headers present in the files. If not, you need to configure ECS appropriately. - If the headers are present, check if all the required headers are present in the files.

Problem:	Parser is not handling data files properly.
Possible Cause(s):	The file might be corrupted.
Action(s):	- Files are marked as 'UNPROCESSED.<file>' and moved to archive directory if one of the following conditions are met: - file is '.gz' and corrupted with CRC error. - file is empty. - file does not have a header. - Files are marked as 'CORRUPTED.<file>' and moved to archive directory if the file is '.gz' and corrupted (other than CRC error) like 'invalid compressed data--format violated'.

Problem:	<code>./serv status</code> shows Postgres processes as NOT RUNNING
Possible Cause(s):	The shared memory configuration in the <code>/etc/system</code> directory might not be correct.
Action(s):	- Check if "shmmax" has been appropriately configured in the <code>/etc/system</code> directory (for Solaris users) or <code>/etc/sysctl.conf</code> directory (for Linux users). - Check the available disk space (especially swap or /tmp) using <code>df -hk</code> command. - Try stopping other MUR instances on the machine. Each MUR instance will consume 2.5 GB of system's shared memory. Use the <code>prstat --a</code> command to check the used and free memory. For hardware details, refer to the <i>Mobility Unified Reporting System Overview</i> chapter of this guide.

Problem:	<code>./serv status</code> shows cache server is not running.
Possible Cause(s):	The port used by cache server implicitly is Postgres port + 2. This is not a configurable port. Make sure that this port is not occupied by any other process.
Action(s):	- Free the required port for cache server. - Reinstall MUR with Postgres port such that postgres port + 2 will be free and can be used by cache server.

Problem:	Scheduler remains in "Paused" state
----------	-------------------------------------

Possible Cause(s):	Disk normal alarm is not sent due to exec_rpc_call failure.
Action(s):	Set the environment variables and execute the following command to clear the paused state forcefully: server/scripts/traps/execute_script.sh 4

Problem:	Archive/ Postgres disk getting full 100%.
Possible Cause(s):	Check if partitioning is done as per recommendation. Refer to the <i>MUR System Requirements</i> section in the <i>Mobility Unified Reporting System Overview</i> chapter of this guide. Purging might not be enabled. Purging is enabled but the dimensions are not configured as per the disk sizes. Purging is enabled but not working. Check if data rate is beyond supported capacity.
Action(s):	• Correct the hardware configurations. • Enable purging. Please contact Cisco service personnel for supported data rate and disk size recommendation.

Problem:	If the EDR files are not getting parsed on the RDP and the files still remain in the input directory.
Possible Cause(s):	The EDR configuration might be incorrect.
Action(s):	• Click ADMIN tab from the MUR GUI. • Click Edit for the gateway for which the EDR files are not getting parsed. • Check the values for the “Flow-Edr Filename Path” and “Http-Edr Filename Path” parameters and compare them with the actual path of files on RDP.

Issues Related to MUR GUI

Problem:	MUR reporting client cannot be started.
Possible Cause(s):	The web browser cache might be full. Check if firewall is disabled.
Action(s):	The browser cache must be cleared. - In the case of Firefox, follow these steps: - On the Tools menu, click Clear Private Data. - Select Cache check box. - Click Clear Private Data Now. - In the case of Internet Explorer, follow these steps: - On the Tools menu, click Internet Options. - Click Delete. - Select Temporary Internet files check box. - Click Delete.  Important: The Firefox version supported for MUR is 3.0.10 and later. For Internet Explorer it is 7.0 and later. Firewall can be disabled using the following command: service iptables stop

Problem:	Unable to login to MUR client.
Possible Cause(s):	When a favorite is created with huge KPI list defined in it, the number of apache requests getting created is equal to no. of KPIs configured. As a result, the MUR is unable to execute further processes due to maximum request limit.
Action(s):	Limit the maximum number of KPI's per favorite in order to resolve this issue.

Problem:	Unable to add / edit / delete gateways.
Possible Cause(s):	The gateway configuration may be incorrect.
Action(s):	- Check if correct IP is provided while adding the gateway. - Check if gateway host is reachable from MUR. For more information, refer to the <i>Cisco Mobility Unified Reporting System Online Help</i> documentation.

Problem:	Unable to add / edit / delete RDPs.
Possible Cause(s):	The RDP configuration may be incorrect.

Action(s):	• Check if correct IP and port are provided while adding RDP. • Check if RDP is actually running on remote machine. • Check if RDP host is reachable from MUR. For more information, refer to the <i>Cisco Mobility Unified Reporting System Online Help</i> documentation.
------------	--

Problem:	Duplicate reports are generated and/or the reports are incorrect.
Possible Cause(s):	MUR might have parsed half-cooked files.
Action(s):	The chassis tags the EDR files with a prefix 'prog.' while transferring to MUR. After the transfer is complete, the chassis removes the 'prog.' tag. The 'prog.' prefix indicates that the file is half cooked. • Check if the EDR files with the prefix 'prog.' are ignored. • Check if EDR file formats are configured properly.

Problem:	The archived files are not getting purged even after configured purging interval.
Possible Cause(s):	MUR might have parsed half-cooked files.
Action(s):	• Check the ownership of files in the archive directory. They must be owned by MUR group user. • The entity pushing the files to MUR, for example, L-ESS should be added to MUR user group. For details, refer to the <i>Managing Mobility Unified Reporting System Installation</i> chapter of this guide.

Problem:	GUI usability issues like some of the menus may not be appearing/functioning as intended after an upgrade of MUR software application.
Possible Cause(s):	It must be because of the "cached media files" in the browser.
Action(s):	Please clear the browser cache and then invoke the MUR GUI.

Problem:	Web browser hangs while fetching huge bulkstats data.
Possible Cause(s):	The browser's java script engine throws error when it exceeds certain time limit of around 30-35 seconds, (time limit is different for different browsers) to render/generate the report on the GUI due to huge set of data.
Action(s):	• Check if the Background functionality for bulkstat is enabled through the MUR GUI. If not, please enable this feature through the System menu. Check if the configurable time limit is set to 30 seconds through System > Bulkstat Background Task Configurations menu. • Though the user can modify and set any higher value, it is recommended to keep this as the maximum value. Also, the "Data Check " when enabled, will intelligently check whether the data set is huge or not. • If data set is huge then it will automatically transfer the bulkstats report in the background. • The task will be moved to Background considering whichever condition is hit first, either the time limit or the data check. (assuming data check is enabled).

■ Issues Related to MUR GUI

Problem:	KPI reports not available after MUR reinstallation - 12.2.214. The following error message "There is some problem while getting the response from the server. Please try after sometime" is observed while running BS or KPI reports, when specific dates are selected.
Possible Cause(s):	Corrupted code in one of the source file (utils.py)
Action(s):	Edit the file "utils.py" in the <i>/MUR/starbi/client/reports</i> directory to resolve this issue.

Problem:	Unable to see BS reports immediately after the upgrade of MUR software application.
Possible Cause(s):	Browser cookies might be causing the reports display even though they are generated.
Action(s):	This issue is solved by deleting the browser cookies.

Problem:	Unable to see the reports in MUR GUI.
Possible Cause(s):	The EDR attributes might not be configured properly.
Action(s):	In the ASR5K CLI command console, check if these two attributes “sn-start-time” and “sn-end-time” in EDR Format are configured in seconds.

Problem:	Created favorite for HTTP report is not redirecting correctly after clicking the Favorite button
Possible Cause(s):	This could be because the Summary and TOP N are two tabs in HTTP, which do not have distinct URLs.
Action(s):	Please select the appropriate fields from HTTP tab with details from Favorites to view the correct report.

Problem:	Schedule and Email feature may not work as intended for DPI and HTTP after an upgrade of MUR software application
Possible Cause(s):	This could be because the browser cookies were not deleted.
Action(s):	Please clear the cookies and then recreate the favorites for these reports.

Issues Related to Master-RDP Communication

Problem:	Master-RDP communication issue.
Action	• Check if json port on master GUI is matching with installed RDP's apache port. • Check if firewall is not enabled on RDP. • Check if all RDP services are running properly. • Check from logs if RDP installation is successful. • Check if Master Details and RDP details are properly configured.

Problem:	While adding RDP, if the MUR GUI throws the following error message “Could not communicate with RDP”
Action(s):	• Check if the Apache server of RDP is running. • If it is not running, start the server using serv start command. • If it is running, check if the Firewall is running on RDP. To check the Firewall status, use the following command: service iptables status • If the output indicates “No firewalls running”, check the Firewall settings for the associated ports and also check if the ports configured for RDP are in use. • From the command output, if you find the Firewall to be running, stop it with the service iptables stop command and then retry. • If you receive this message “RDP is already configured”, please contact Cisco Advanced Service team for additional support and guidance.

Problem:	If sftpying of EDR/UDR files failed
Possible Cause(s):	• SSH keys and SFTP server might not be configured appropriately. • SFTP might not be running on MUR server.

Action(s):	SSH keys and SFTP server need to be configured on the chassis and also SFTP should be running on MUR server. • Check if the following variables in the <i>sshd_config</i> file present in the <i>/etc/ssh</i> directory are set appropriately. • PermitRootLogin = yes • PasswordAuthentication = yes • PAMAuthenticationViaKBDInt = no (Applicable ONLY for SOLARIS) • UsePAM = no (Applicable ONLY for RHEL) • Comment the line “PAMAuthenticationViaKBDInt yes” as “#PAMAuthenticationViaKBDInt yes” • Update the SFTP parameters as necessary if the variables are not set properly. • After updating restart SSH daemon using the following commands: For Solaris 9: <code>/etc/init.d/ssh restart</code> For Solaris 10: <code>svcadm restart svc:/network/ssh:default</code> For RHEL: <code>service sshd restart</code> If the problem still persists, remove the EDR generation configuration from the gateway and reconfigure them.
------------	--

Problem:	Though the EDR files are parsed at the RDP, the reports are not available for a gateway attached to the RDP.
----------	--

Action(s):	• Check if RDP is actually running on remote machine. • Check if there are any pending files in the RDP's <i>/starbi/server/sql_export_data/</i> directory. • If there are any pending files, check if the following log is shown in the <i>starbi_server_devel</i> file located in the <i>/starbi/logs/server/</i> directory. <i>Unable to open host keys file" [rdptomaster_file_mover.py ...]"</i> • If the logs are found, perform the following procedure: • Log on to the RDP machine. Switch the user as RDP admin. If you have set RDP admin as myrdp during RDP installation, then execute the following command su - myrdp. • Create an SFTP session to master (sftp masteradmin@masterhost) The output will look something similar to the following: <pre>Connecting to <masterhost>... The authenticity of host <masterhost> (<master ip address>)' can't be established. RSA key fingerprint is <some key>. Are you sure you want to continue connecting (yes/no)?</pre> • Enter yes and quit the SFTP session. • Check the files again in the <i>starbi/server/sql_export_data</i> directory to identify if they are present now. Then, check the reports after a while. • If the reports are still not seen, check the master configuration in System menu. Check if the master login and password are specified. Ensure that master admin password has been set using password command on the master. If the solution provided above does not help to resolve your problem, please contact Cisco Advanced Service Team for additional support and guidance.
------------	---

Issues Related to BS-KPI Alarms

Problem:	The bulkstats or KPI reports are not generated.
Possible Cause(s):	The bulkstats file might not be parsed.
Action(s):	- Check if the bulkstats schemas are properly configured on the gateway through the BULKSTATS menu in the ADMIN tab. - Check if the prerequisites described in the <i>Configuring Bulkstats Schemas</i> section of the <i>Configuring Chassis for MUR</i> chapter are met. - On the ADMIN menu, check the bulkstats audit under AUDIT. The audit should indicate whether the bulkstats files are being parsed or not. For more information, refer to the <i>Cisco Mobility Unified Reporting System Online Help</i> documentation. - Check if FTP is enabled on the MUR server. - Check if the bulkstats are FTPed to correct location on the MUR server from the gateway. The path should be as follows: <code>\$STARBI_HOME/server/data/\$gwname/bs</code> Where \$STARBI_HOME = MUR installation directory, \$gwname = Gateway name - Check if the Bulkstats/KPIs UI show only one day data - Check if the counters of same schemas are added in the formula while configuring KPIs. For example, if you are adding KPI in SGSN schema, then you should add counters of SGSN only in the formula. - Check if the bulkstat files are always pushed from gateway to master MUR and not to RDP.

Problem:	If user is not able to configure bulkstats schema through Add Schema configuration screen that appears by selecting ADMIN > Bulkstats menu.
Possible Cause(s):	The initial prerequisites might not be met.
Action(s):	- Check if the prerequisites described in the <i>Configuring Bulkstats Schemas Using GUI</i> section of the <i>Configuring Chassis for MUR</i> chapter are met. - Check if the SSH Username in the Add Bulkstats schema configuration screen is specified correctly. This user name is used to connect to gateway via SSH for schema configuration. For information on how to configure the schemas, refer to the <i>Configuring Bulkstats Schemas Using GUI</i> section. Also, see the <i>Cisco Mobility Unified Reporting System Online Help</i> documentation.

Problem:	The bulkstats files are not pushed from the chassis to the MUR server even after successfully configuring schemas.
Possible Cause(s):	The related configurations might be incorrect.

Action(s):	- Check if Username specified in the Add Bulkstats schema configuration screen is present in MUR group on the MUR server. To create the username if it does not exist on MUR server, use the following command: useradd <i><new user name></i> If the user is already present then use the following command to add the user in the MUR group. usermod -G <i><MUR Group></i> <i><user name></i> - Check if Destination specified in the Add Bulkstats schema configuration screen is correct or not.
------------	---

Problem:	While configuring bulkstat schema if configuration screen hangs for a long time.
Possible Cause(s):	The session might be timed out.
Action(s):	- Close the browser and try to configure the schema again. - Restart apache server by executing the following command from the <i><MUR_install_dir>/starbi/bin</i> directory and try again to configure the schema. ./serv start apache

Problem:	If user is not able to edit the schema configuration through the Add Bulkstat schema configuration screen.
Possible Cause(s):	The schemas may not be configured properly.
Action(s):	- Follow the steps mentioned for above 3 cases. - If you are still not able to configure then delete the schema configuration for that particular schema from the GUI and try to configure again.

Issues Related to Backup and Restore

Problem:	Favorites list is not recovered from a database backup in MUR running version 12.2.186
Possible Cause(s):	MUR is not taking backup of the tables which are used in Favorites and Offline tasks. The models for the 'Favorites' functionality were not included in the backup feature.
Action(s):	The models for the 'Favorites' functionality should be included in the backup feature.

Miscellaneous Issues

Problem:	getSupportDetails script does not work on Linux and Solaris.
Possible Cause(s):	• The script is trying to include Veritas cluster logs irrespective of whether or not the setup is present on the machine. • This script is unable to load the appropriate perl modules for Solaris.
Action(s):	This is a known issue with some versions of MUR. 1. Check if the environment variables are set using the following commands: <code>cd <MUR installation path></code> <code>source bin/setenv.sh</code> 2. Copy the latest script, replace MUR installation path correctly and execute after setting environment variables.