Tivoli. software

IBM

# IBM TotalStorage Productivity Center Advanced Topics

Learn to effectively use TotalStorage Productivity Center

Create customized reports

Maintain the database repository

Mary Lovelace
Jason Bamford
Doug Dunham
Daniel Frueh
Adelson Lovatto
Amit Vallabh

Redbooks

ibm.com/redbooks

International Technical Support Organization

# IBM TotalStorage Productivity Center Advanced Topics

June 2007

SG24-7348-00

**Note:** Before using this information and the product it supports, read the information in “Notices” on page vii.

**First Edition (June 2007)**

This edition applies to Version 3, Release 2 of IBM TotalStorage Productivity Center (product number 5608-VC0).

# Contents

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

# Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo) ®
z/OS®
AIX®
DB2 Universal Database™
DB2®
DFSMS™
DS4000™
DS6000™
DS8000™
Enterprise Storage Server®
IBM®
Redbooks®
Tivoli Enterprise™
Tivoli Enterprise Console®
Tivoli®
TotalStorage®
WebSphere®

The following terms are trademarks of other companies:

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Network Appliance, NetApp, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

Java, JVM, Solaris, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Excel, Microsoft, MS-DOS, Visual Basic, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

# Preface

You have installed and performed the basic customization of IBM® TotalStorage® Productivity Center. You have successfully completed performance data collection and have generated reports. But how do you best use the TotalStorage Productivity Center to manage your storage infrastructure?

This IBM Redbooks® publication shows how to best set up TotalStorage Productivity Center based on the storage environment infrastructure, and then manage that storage infrastructure with TotalStorage Productivity Center. It includes experiences from client accounts and our own internal experiences. This book includes the following topics:

- TotalStorage Productivity Center installation considerations:
  - Number of servers
  - Database placement
  - Firewall considerations
  - Agent deployment
- CIMOM management

  How many are required and how to customize them
- Performance monitoring:
  - Setting up thresholds and alerts
  - Gathering data
  - Which reports to use
- Custom report creation with TPCTOOL
- TotalStorage Productivity Center maintenance:
  - Data retention
  - Database backup
  - Debugging tools

This Redbooks publication is intended for use by storage administrators, who are responsible for the performance and growth of the IT storage infrastructure.

## The team that wrote this IBM Redbooks publication

This IBM Redbooks publication was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center (see Figure 1 on page x).

*Figure 1 Amit, Daniel, Jason, Mary, Lovatto, Doug*

**Mary Lovelace** is a Consulting IT specialist at the International Technical Support Organization. She has more than 20 years of experience with IBM in large systems, storage and storage networking product education, system engineering and consultancy, and systems support. She has written many IBM Redbooks publications on TotalStorage Productivity Center and z/OS® storage products.

**Jason Bamford, MBCS CITP**, is a Certified IT Specialist in the IBM Software Business, United Kingdom. He has 22 years of client experience in finance, commercial, and public sector accounts, consulting and deploying mid-range systems in AIX®, Windows®, and other UNIX® variants. An IBM employee for nine years, Jason specializes in IBM software storage products and is a subject matter expert in the UK for Tivoli® Storage Manager and TotalStorage Productivity Center.

**Doug Dunham** is from Palo Alto and Sunnyvale, California. He served in the U.S. Army and a year in Vietnam. His Army career gave him a start in digital electronics, which provided the basis for his employment with two startup electronic companies in the early days of Silicon Valley, and varied careers within IBM from manufacturing test on 3340 and 3350 storage devices to read channel Product Engineering on 3380s, to software development in DFSMS™ and storage resource management, to his position today as a senior member of the Tivoli Storage SWAT team. Doug assists clients with problems related to storage environments and uses IBM solutions to help those clients understand and control their storage environments.

**Daniel Frueh** is an Advisory IT Specialist in IBM Switzerland Strategic Outsourcing. He has four years of experience in the Open Storage field. He holds a degree in Computer Science from the University of Rapperswil. His areas of expertise include Tivoli Storage Manager, SVC, DS8k, DS6k, SAN, and NAS.

**Adelson Lovatto** is an IT Architect, with 19 years of experience in Information Technology, working since 1993 for IBM Brazil, where he currently focuses on designing solutions for Information Lifecycle Management (ILM). He is certified in these areas: a certified IT Architect by IBM, Storage Networking Industry Association (SNIA), The Open Group Architecture

Framework (TOGAF), and also for some IT products, including TPC. As a member of the Brazil Technology Leadership Council (TLC), an affiliation of the IBM Academy of Technology (AoT), and as an IBM Ambassador through the program IBM Academic Initiatives, he has written papers and delivered presentations to help in the dissemination of IBM technologies in the market and at universities.

**Amit Vallabh** is an IT Storage Specialist for IBM Global Technology Services, South Africa. He has six years of experience implementing, designing, and supporting IT Solutions in storage and Open Systems. He is a Certified IBM Tivoli Storage Manager Specialist. His expertise includes SAN Storage Area Network, IBM Tivoli Storage Manager, IBM TotalStorage Productivity Center, and Information Life Cycle Management (Data Classification). Prior to returning to IBM in November 2005, he was a senior Tivoli Storage Manager administrator at a large insurance company and supported Disaster Recovery, Business Continuity, and the backup environment. He holds a three year National Diploma in Computer Systems Engineering from Technikon Witwatersrand (Johannesburg), South Africa.

Thanks to the following people for their contributions to this project:

Sangam Racherla
International Technical Support Organization, San Jose Center

Diana Duan
Stephan Jaquet
Brian Smith
Larry Mills
Lloyd Dean
IBM San Jose

Yijie Zhang
Greg Van Hise
IBM Tucson

Wendy L Henson
IBM Raleigh

Frank Chodacki
IBM Austin

Adam Crosskey
Storage Systems Engineer
Oxfordshire County Council, Oxfordshire, United Kingdom

Matthew James
Deputy Service Support Manager
Oxfordshire County Council, Oxfordshire, United Kingdom

# Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbooks publication dealing with specific products or solutions, while getting hands-on experience with leading edge technologies. You will have the opportunity to team with IBM technical professionals, IBM Business Partners, and clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

**ibm.com**/redbooks/residencies.html

# Comments welcome

Your comments are important to us!

We want our IBM Redbooks publications to be as helpful as possible. Send us your comments about this or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review IBM Redbooks publications form found at:

  **ibm.com**/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

1

# Component overview

In this chapter, we describe the components of IBM TotalStorage Productivity Center and the functions that they provide, including the functions provided by the combination of the components. We then describe the types of agents and the information that is provided by each type.

## 1.1 TotalStorage Productivity Center overview

IBM TotalStorage Productivity Center (TPC) has evolved from a suite of products that provides end-to-end storage infrastructure management from the host application to the target storage device in a heterogeneous platform environment. TPC provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance, and assets in the IT environment. TPC also provides disk and tape subsystem configuration and management, performance management, SAN fabric management and configuration, and usage reporting and monitoring from the perspective of the database application as well as from the perspective of the filesystem (see Figure 1-1).

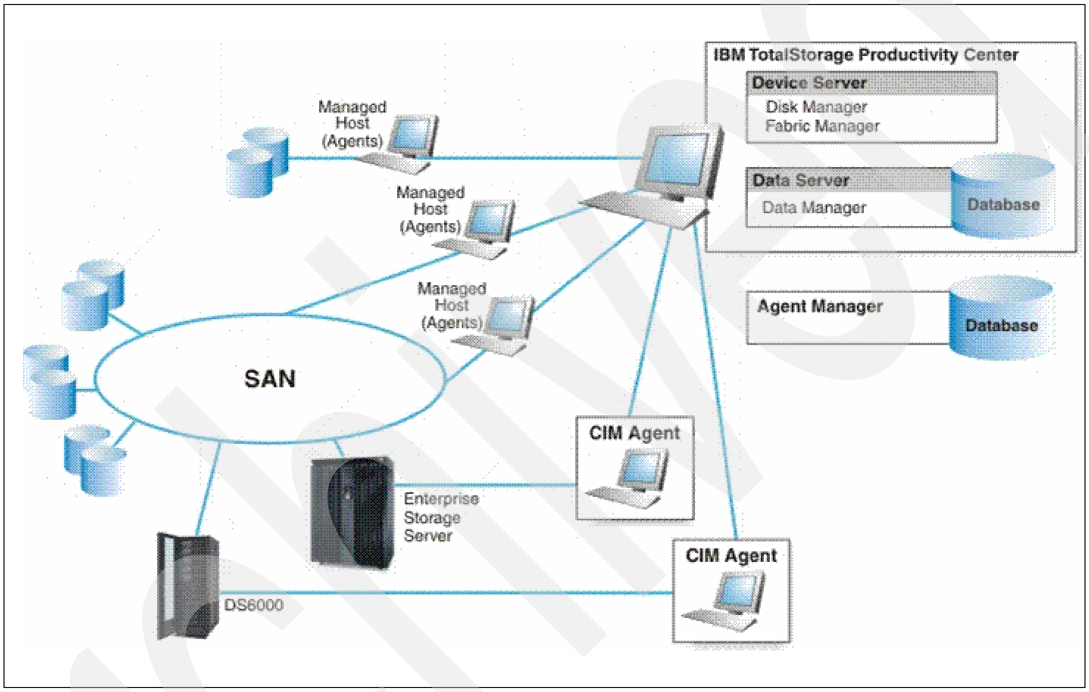

Figure 1-1 TotalStorage Productivity Center component interaction

### 1.1.1 TPC components

In this section, we review the key components of TPC that provide storage area networks (SAN) infrastructure management.

#### Data manager

The *Data manager* is the component that is responsible for understanding the behavior of the user data that originated from the filesystems and databases. The Data manager includes enterprise-wide reporting and monitoring, policy-based management, and automated capacity provisioning for Direct-Attached Storage (DAS), Network-Attached Storage (NAS), and SAN environments.

The purpose of this component is to discover the usage patterns and capacity utilization of the storage from an application's point of view. The Data manager is also responsible for gathering application usage demographics, reporting on storage from the application's perspective, and performing and automating actions that provision application storage either by itself or invoking other components to do so. The Data manager is additionally responsible for collecting and analyzing performance information from the filesystems and the databases that are used by applications.

More details about this component are provided in 1.2, “TotalStorage Productivity Center for Data” on page 4.

### Disk manager

The *Disk manager* is the component that is responsible for discovering and monitoring storage subsystems and performing and automating actions that provision disk resources. The Disk manager is additionally responsible for collecting and analyzing performance information from storage subsystems and helping in the management of SANs and heterogeneous storage devices from a single console.

More details about this component are provided in 1.3, “TotalStorage Productivity Center for Disk” on page 6.

### Fabric manager

The *Fabric manager* is the component that is responsible for discovering, monitoring, and controlling the SAN topology, primarily concentrating on the switched SAN fabric and on the hosts and storage systems that are zoned for access. The Fabric manager is a comprehensive management solution for multi-vendor SANs. The Fabric manager is responsible also for collecting and analyzing performance information from SAN fabrics.

More details about this component are provided in 1.4, “TotalStorage Productivity Center for Fabric” on page 8.

### TotalStorage Productivity Center database

The TotalStorage Productivity Center database is a single DB2 database instance that serves as a repository for all TPC components. Detailed information about the TPC database is provided in Chapter 3, “IBM Total Productivity Center database considerations” on page 61.

### Device server

The *Device server* discovers storage subsystems and SAN fabrics, and then it gathers information about storage subsystems and SAN fabrics and analyzes their performance. The Device server controls the communication with agents and the data collection from agents that scan SAN fabrics. It is also responsible for the creation and monitoring of replication relationships between storage devices.

### Data server

The *Data server* hosts the control points for product scheduling functions, configuration, event information, reporting, and graphical user interface support. It coordinates the communication with agents and the data collection from agents that scan filesystems and databases to gather storage demographics and populate the TPC database with results. Automated actions can be defined to drive functions, such as filesystem extension, data deletion, TSM backup or archiving, or event reporting when defined thresholds are encountered. The Data server is the primary contact point for all user interface functions. It also includes functions that schedule data collection and discovery for the Device server.

### Agent Manager

The *Agent Manager* gathers host, application, and SAN fabric information and sends it to the Data server and the Device server. The TPC agents are based on Tivoli Common Agent Services (CAS). The Common Agent Services provides a way to deploy agent code across multiple user machines or application servers throughout the enterprise. The deployed agent code collects data from and performs operations on managed resources on behalf of TPC.

### Agents

TPC needs a combination of *agents* to gather data about the devices that will be monitored and managed. Different combinations of these agents are required to effectively enable the functions of the Data Manager, Fabric Manager, Disk Manager, and Tape Manager. The agents are required to enable the function of the Topology Viewer that reflects the entities in the management scope of TPC.

More information about these agents is in 1.5, “Agents” on page 9.

### Other components

The *Graphic User Interface (GUI)* provides an integrated user interface (UI) for all TPC interactions. TPC provides two methods for running the user interface: as a downloadable Java applet or as an application. The benefit of running the user interface as a Java applet is that it is not necessary to install the GUI component on every workstation. Users can simply access the TPC applet from any Web-enabled workstation and the appropriate applets will be automatically downloaded for them on an as-needed basis.

The *Command Line Interface (CLI)* provides scriptable access to major TPC functions. Generic commands are available regardless of which TPC components have been licensed and configured. The only requirements are an active Device server and a valid login. However, in many cases, these commands do provide richer function depending on the components that have been enabled. For example, disk commands are available when disk components have been installed and fabric commands are available when fabric components have been installed.

Provisioning function is provided by the Tivoli Provisioning Manager and by the Tivoli Intelligent Orchestrator. These Tivoli products interface with TPC and can provide automated provisioning of storage through workflows.

## 1.2 TotalStorage Productivity Center for Data

The Data Manager component is also known as TotalStorage Productivity Center for Data (TPC for Data). The Data Manager can be installed separately from the other components of TPC. For historic reasons, TPC for Data is also the way that many professionals used to name the component Data Manager.

When you install only the TPC for Data component, the Navigation Tree on the main TPC GUI window in the left pane shows the functions Administrative Services and IBM TotalStorage Productivity Center. Administrative Services and IBM TotalStorage Productivity Center are common for all TPC components, as well as Data Manager, Data Manager for Databases, and Data Manager for Chargeback (see Figure 1-2 on page 5).

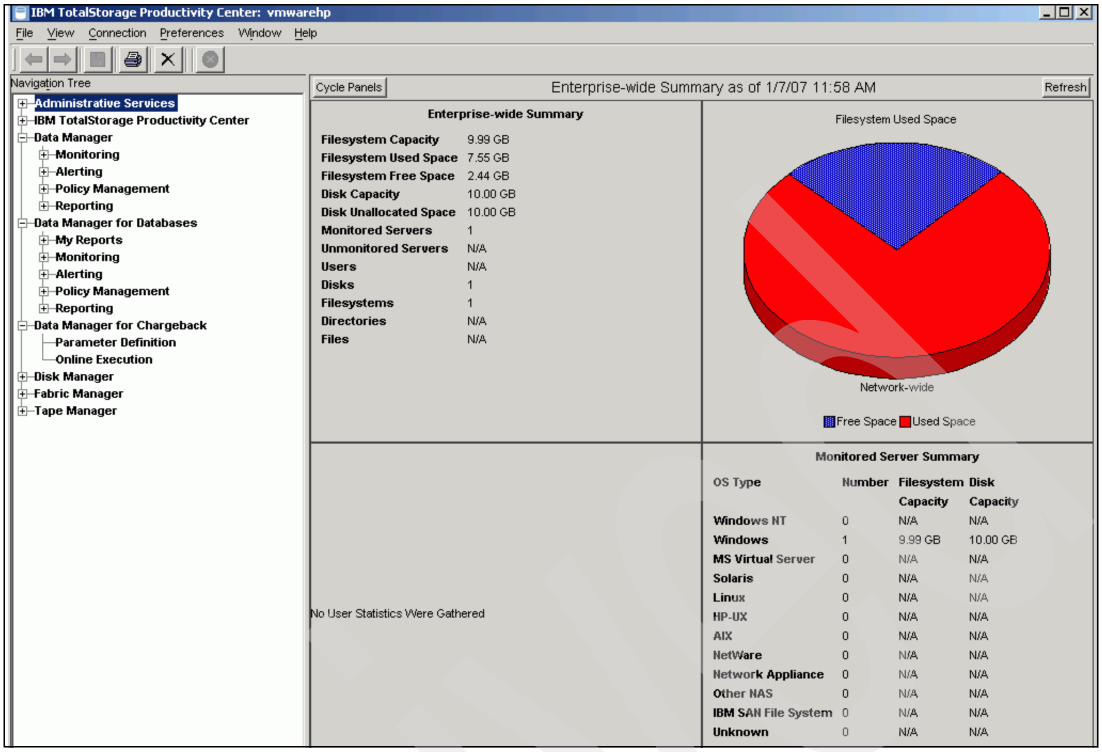
IBM TotalStorage Productivity Center: vmwarehp
File View Connection Preferences Window Help
Navigation Tree
Cycle Panels
Refresh
Filesystem Used Space
Network-wide
Free Space
Used Space

- Administrative Services
- IBM TotalStorage Productivity Center
- Data Manager
  - Monitoring
  - Alerting
  - Policy Management
  - Reporting
- Data Manager for Databases
  - My Reports
  - Monitoring
  - Alerting
  - Policy Management
  - Reporting
- Data Manager for Chargeback
  - Parameter Definition
  - Online Execution
- Disk Manager
- Fabric Manager
- Tape Manager

Enterprise-wide Summary as of 1/7/07 11:58 AM

Enterprise-wide Summary

| | |
|---|---|
| Filesystem Capacity | 9.99 GB |
| Filesystem Used Space | 7.55 GB |
| Filesystem Free Space | 2.44 GB |
| Disk Capacity | 10.00 GB |
| Disk Unallocated Space | 10.00 GB |
| Monitored Servers | 1 |
| Unmonitored Servers | N/A |
| Users | N/A |
| Disks | 1 |
| Filesystems | 1 |
| Directories | N/A |
| Files | N/A |

No User Statistics Were Gathered

Monitored Server Summary

| OS Type | Number | Filesystem Capacity | Disk Capacity |
|---|---|---|---|
| Windows NT | 0 | N/A | N/A |
| Windows | 1 | 9.99 GB | 10.00 GB |
| MS Virtual Server | 0 | N/A | N/A |
| Solaris | 0 | N/A | N/A |
| Linux | 0 | N/A | N/A |
| HP-UX | 0 | N/A | N/A |
| AIX | 0 | N/A | N/A |
| NetWare | 0 | N/A | N/A |
| Network Appliance | 0 | N/A | N/A |
| Other NAS | 0 | N/A | N/A |
| IBM SAN File System | 0 | N/A | N/A |
| Unknown | 0 | N/A | N/A |

*Figure 1-2 Navigation tree for Data component*

TPC for Data can help improve storage utilization, plan for future capacity, and ensure availability by providing storage on demand for filesystems. TPC for Data performs the following functions:

- Discover and monitor disks, partitions, shared directories, and servers
- Monitor and report on capacity and utilization across platforms, helping to identify trends and prevent problems
- Provides a wide variety of standardized reports about filesystems and storage infrastructure to track usage and availability
- Provide file analysis across platforms, helping to identify and reclaim space used by nonessential files
- Provide policy-based management and automated capacity provisioning for filesystems when user-defined thresholds are reached

When only TPC for Data is installed and not the entire suite of products, only the following data management functions are available:

- Host-centric:
  - Discovery
  - Monitoring
  - File System extension
- Application-centric:
  - Monitor DB2, Oracle, SQL Server, and Sybase
  - Discovery
  - Monitoring
  - Chargeback

The following functions are not available:

- For Storage subsystems (disk management):
  - Discovery
  - Monitoring
  - Configuration (for example, creating volumes)
  - Performance management
- For fabrics (fabric management):
  - Discovery
  - Monitoring
  - Configuration (for example, zoning)
  - Performance management
- For tape (tape management):
  - Discovery
  - Monitoring

**Note:** When installing only the TPC for Data, you do not need to install CIMOM agents, because only data management functions are available.

## 1.3 TotalStorage Productivity Center for Disk

The Disk Manager component is known also as TotalStorage Productivity Center for Disk (TPC for Disk). When installing only the TPC for Disk component, the Navigation Tree on the main TPC GUI window in the left pane shows the functions Administrative Services and IBM TotalStorage Productivity Center. Administrative Services and IBM TotalStorage Productivity Center are common for all TPC components, as well as Disk Manager and Tape Manager (see Figure 1-3 on page 7).

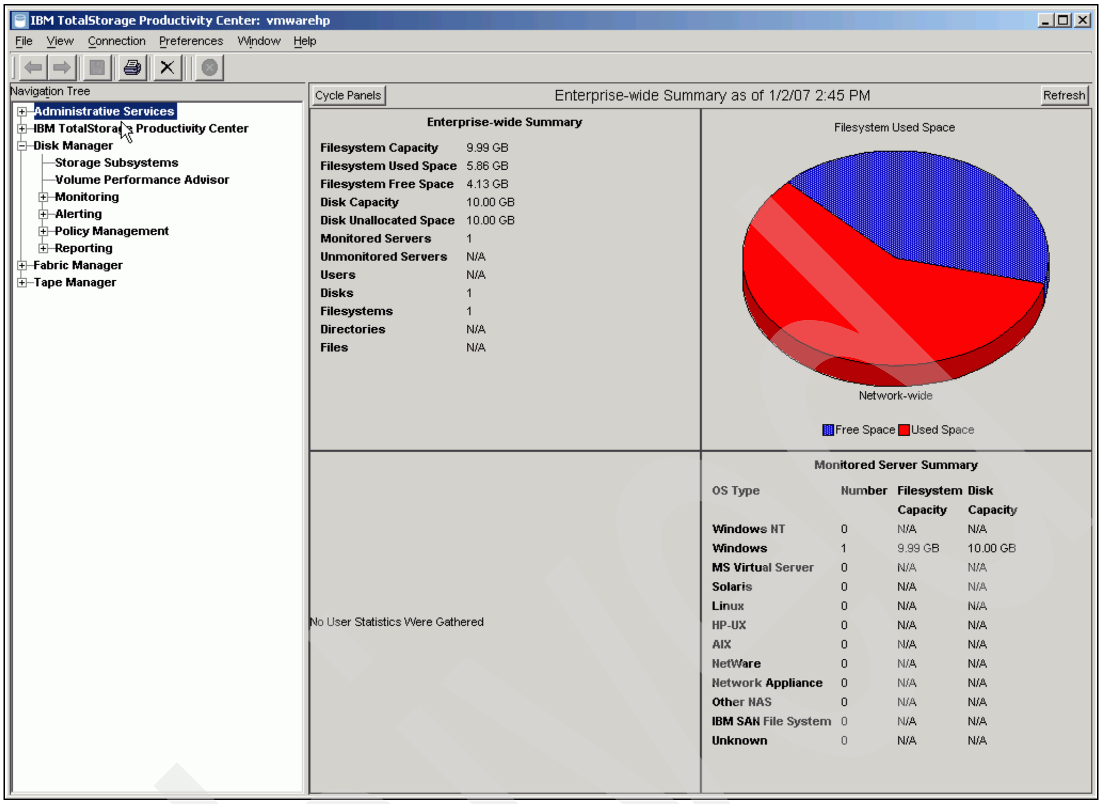

Enterprise-wide Summary as of 1/2/07 2:45 PM

**Enterprise-wide Summary**

| | |
|---|---|
| Filesystem Capacity | 9.99 GB |
| Filesystem Used Space | 5.86 GB |
| Filesystem Free Space | 4.13 GB |
| Disk Capacity | 10.00 GB |
| Disk Unallocated Space | 10.00 GB |
| Monitored Servers | 1 |
| Unmonitored Servers | N/A |
| Users | N/A |
| Disks | 1 |
| Filesystems | 1 |
| Directories | N/A |
| Files | N/A |

No User Statistics Were Gathered

**Monitored Server Summary**

| OS Type | Number | Filesystem Capacity | Disk Capacity |
|---|---|---|---|
| Windows NT | 0 | N/A | N/A |
| Windows | 1 | 9.99 GB | 10.00 GB |
| MS Virtual Server | 0 | N/A | N/A |
| Solaris | 0 | N/A | N/A |
| Linux | 0 | N/A | N/A |
| HP-UX | 0 | N/A | N/A |
| AIX | 0 | N/A | N/A |
| NetWare | 0 | N/A | N/A |
| Network Appliance | 0 | N/A | N/A |
| Other NAS | 0 | N/A | N/A |
| IBM SAN File System | 0 | N/A | N/A |
| Unknown | 0 | N/A | N/A |

*Figure 1-3 Navigation Tree for Disk component*

TPC for Disk helps in the management of SANs and heterogeneous storage from a single console, because the components support the SMI-S standards. Device discovery is performed by using the Service Location Protocol (SLP), as specified by SMI-S. Configuration of the discovered devices is possible in conjunction with CIM agents associated with those devices, using the standard mechanisms defined in SMI-S. TPC for Disk also gathers events and can launch an element manager specific to each device. It performs the following functions:

- Collects and stores performance data and provides alerts
- Provides graphical performance reports
- Helps to optimize storage allocation
- Provides volume contention analysis

When only TPC for Disk is installed and not the entire suite of products, the following functions are available:

- For Storage subsystems (disk management):
  - Discovery
  - Monitoring
  - Configuration (for example, creating volumes)
  - Performance management
- For tape (tape management):
  - Discovery
  - Monitoring

The following functions are not available:

- For fabric management:
  - Discovery
  - Monitoring
  - Configuration (for example, zoning)
  - Performance management
- For data management:
  - Host-centric:
    - Discovery
    - Monitoring
    - Filesystem extension
  - Application-centric:
    - Monitor DB2, Oracle, SQL Server, and Sybase
    - Discovery
    - Monitoring
    - Chargeback

## 1.4 TotalStorage Productivity Center for Fabric

The Fabric Manager component is known also as TotalStorage Productivity Center for Fabric (TPC for Fabric). When installing only the TotalStorage Productivity Center for Fabric component, the Navigation Tree found on the main TPC GUI window in the left pane shows the functions Administrative Services and IBM TotalStorage Productivity Center that are common for all TPC components, as well as the Fabric Manager as shown in Figure 1-4.

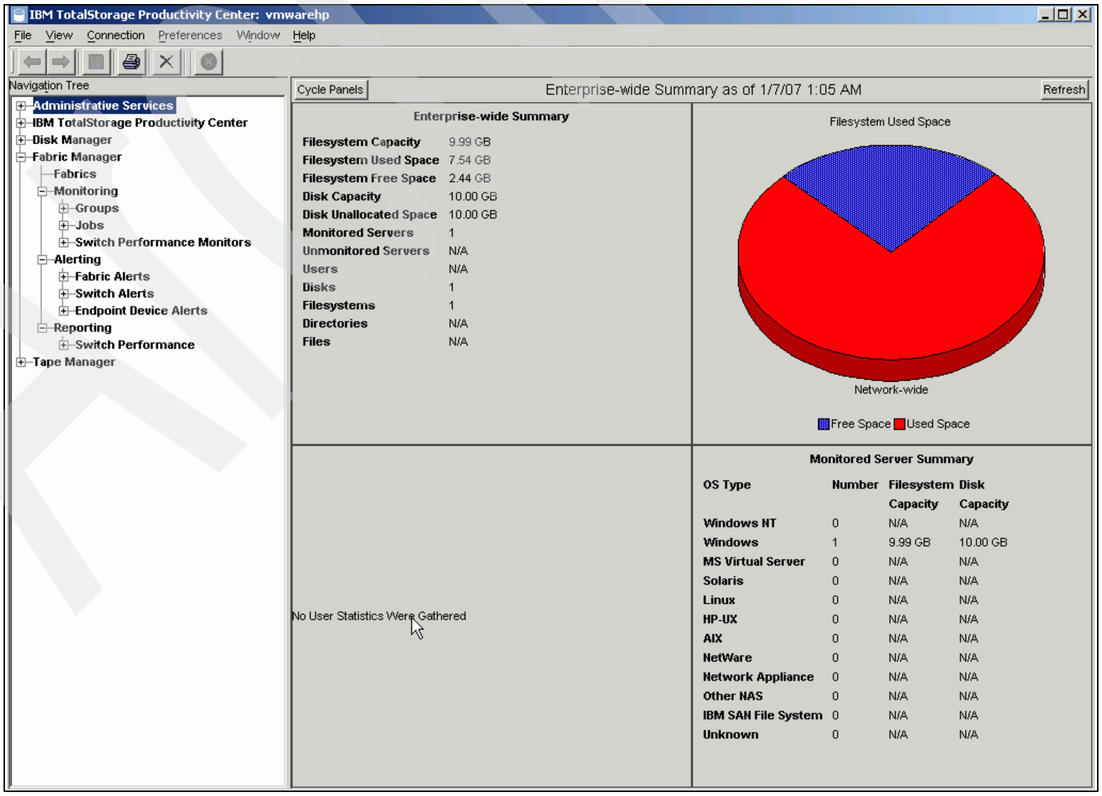

Enterprise-wide Summary as of 1/7/07 1:05 AM

- Administrative Services
- IBM TotalStorage Productivity Center
- Disk Manager
- Fabric Manager
  - Fabrics
  - Monitoring
    - Groups
    - Jobs
    - Switch Performance Monitors
  - Alerting
    - Fabric Alerts
    - Switch Alerts
    - Endpoint Device Alerts
  - Reporting
    - Switch Performance
- Tape Manager

Enterprise-wide Summary

| | |
|---|---|
| Filesystem Capacity | 9.99 GB |
| Filesystem Used Space | 7.54 GB |
| Filesystem Free Space | 2.44 GB |
| Disk Capacity | 10.00 GB |
| Disk Unallocated Space | 10.00 GB |
| Monitored Servers | 1 |
| Unmonitored Servers | N/A |
| Users | N/A |
| Disks | 1 |
| Filesystems | 1 |
| Directories | N/A |
| Files | N/A |

No User Statistics Were Gathered

Monitored Server Summary

| OS Type | Number | Filesystem Capacity | Disk Capacity |
|---|---|---|---|
| Windows NT | 0 | N/A | N/A |
| Windows | 1 | 9.99 GB | 10.00 GB |
| MS Virtual Server | 0 | N/A | N/A |
| Solaris | 0 | N/A | N/A |
| Linux | 0 | N/A | N/A |
| HP-UX | 0 | N/A | N/A |
| AIX | 0 | N/A | N/A |
| NetWare | 0 | N/A | N/A |
| Network Appliance | 0 | N/A | N/A |
| Other NAS | 0 | N/A | N/A |
| IBM SAN File System | 0 | N/A | N/A |
| Unknown | 0 | N/A | N/A |

*Figure 1-4 Navigation Tree for Fabric Manager*

TPC for Fabric helps in the management of SAN fabrics that connect the host systems and applications to the storage devices. TPC for Fabric is a comprehensive management solution for multi-vendor SANs and includes automatic resource and topology discovery, monitoring and alerts, and zone control. TPC for Fabric helps:

- Simplify the task of SAN management and configuration
- Ensure SAN availability
- Improve SAN return on investment

If you install only TPC for Fabric, the following functions are available:

- For fabric management:
  - Discovery
  - Monitoring
  - Configuration (for example, zoning)
  - Performance management

The following functions are not available:

- For storage subsystems (disk management):
  - Discovery
  - Monitoring
  - Configuration (for example, creating volumes)
  - Performance management
- For tape (tape management):
  - Discovery
  - Monitoring
- For data management:
  - Host-centric:
    - Discovery
    - Monitoring
    - Filesystem extension
  - Application-centric:
    - Monitor DB2, Oracle, SQL Server, and Sybase
    - Discovery
    - Monitoring
    - Chargeback

## 1.5 Agents

TPC uses four types of agents to gather data about the storage subsystems and servers that will be monitored and managed:

- CIMOM agent
- SNMP agent
- Data agent
- Fabric agent

Figure 1-5 on page 10 shows agent interaction in a TPC environment.

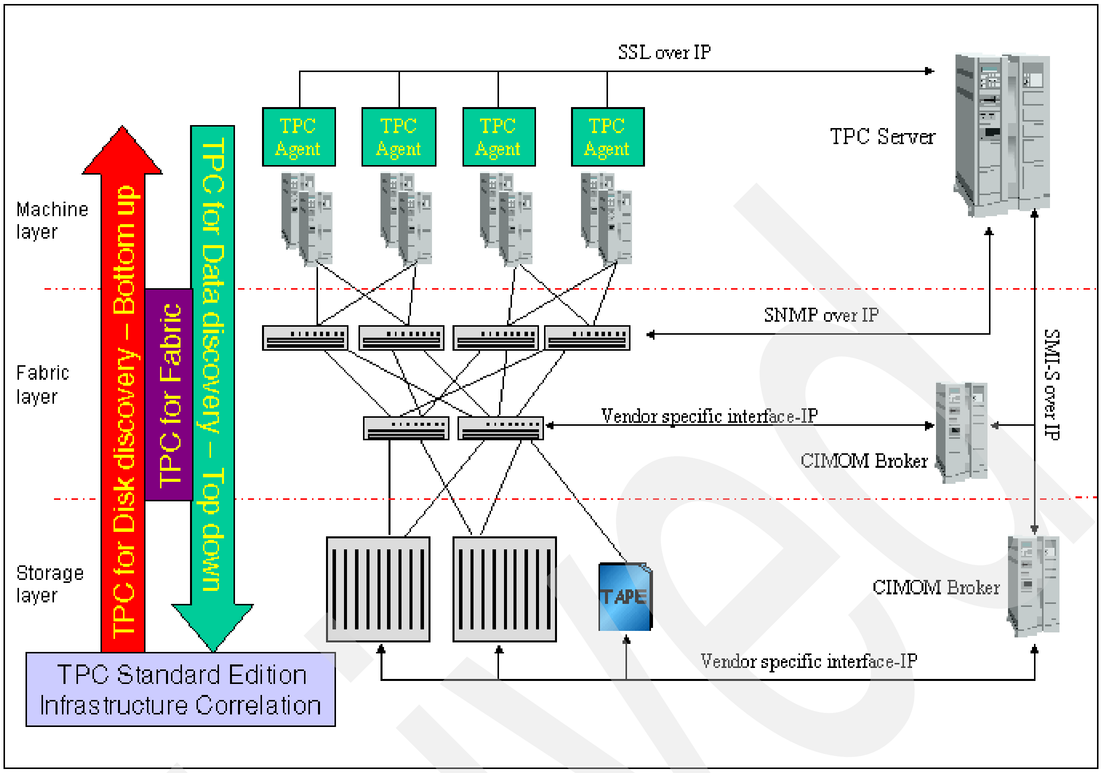

*Figure 1-5 Agent interaction*

### 1.5.1 CIMOM agent

The *Common Information Model Object Manager (CIMOM)* provides the means by which a device can be managed by common building blocks rather than proprietary software. If a device is CIM-compliant, TPC can manage it, because TPC is also CIM-compliant. Actually, the *CIM agent* is an interpreter between TPC and the device. There is an SNIA SMI-S interface using an XML transport for data and command interchange from TPC to the CIMOM. From the CIMOM to the device layer, there are proprietary interfaces provided by the device vendor to convert those commands and answers from the SNIA XML language to a language that the device can understand:

- For storage, the CIM agents are needed for storage asset information, provisioning, alerting, and performance monitoring.
- For fabric switches, the CIM agents are only used for performance monitoring.
- For tape libraries, the CIM agents are used for asset and inventory information.

The CIM agents can be referred to by a variety of names, including CIMOM agent and SMI-S Provider. The CIM agent can be a separate agent installation or can be imbedded in the device itself, which is the case with Cisco fabric switches. In this case, there is no proxy agent to install and TPC is configured to point to the managed device.

After you install and configure the CIM agent, you can configure TPC to communicate with it.

### 1.5.2 SNMP agent

Actually, the correct name is Out-of-Band Fabric (OOBF) agent, and SNMP is in fact the protocol that is used by it to collect topology information from fabric switches through queries using the IP network. However, to simplify the classification of the types of agents that TPC uses, *SNMP agent* is used to identify it because SNMP is the standard.

The Out-of-Band Fabric agents are used by the Fabric Manager to collect topology information from the fabric switches through the IP network using SNMP queries to the switches.

In order to manage the SAN fabric effectively, it is fundamental to have an OOBF agent pointing to each switch in the SAN fabrics that you are monitoring. The use of OOBF agents is required to collect VSAN information for Cisco switches and zoning information for Brocade switches where the admin user ID and password are needed.

### 1.5.3 Data agent

The *Data agents* collect information from the machine or host on which they are installed. The agents collect asset information, file and filesystem attributes, and any other information needed from the computer system. Data agents can also gather information on database managers installed on the server, Novell NDS tree information, and NAS device information. You create ping, probe, and scan jobs to run against the servers that have Data agents installed.

The Data agents are installed on all of the computer systems that you want TPC to manage.

### 1.5.4 Fabric agent

*Fabric agents* use scanners to collect information. The Fabric agents are installed on computer systems that have fiber optic connectivity (through HBAs) into the SAN fabrics that you want to manage and monitor.

The scanners are written in operating system code and communicate through the HBA to collect fabric topology information, port state information, and zoning information. The Fabric agents can also identify other SAN-attached devices (if they are in the same zone). Using operating system calls, the agents collect information about the machine on which they are installed.

## 1.6 Licensing structure

TPC is packaged in two marketing package distributions:

- TotalStorage Productivity Center Standard Edition
- TotalStorage Productivity Center Limited Edition

There are package distributions for individual components:

- TotalStorage Productivity Center for Data (discussed in “TotalStorage Productivity Center for Data” on page 4)
- TotalStorage Productivity Center for Disk (discussed in “TotalStorage Productivity Center for Disk” on page 6)
- TotalStorage Productivity Center for Fabric (discussed in “TotalStorage Productivity Center for Fabric” on page 8)

### 1.6.1 IBM TotalStorage Productivity Center Standard Edition

The TotalStorage Productivity Center Standard Edition can be considered the *full version* of TotalStorage Productivity Center. It is a packaged offering consisting of IBM TotalStorage Productivity Center for Data, Disk, and Fabric. The TotalStorage Productivity Center Standard Edition provides the following functions:

- **Data Management:**
  - Host-centric:
    - Discovery
    - Monitoring
    - File System extension
  - Application-centric:
    - Monitor DB2, Oracle, SQL Server, and Sybase
    - Discovery
    - Monitoring
    - Chargeback
- **Disk Management:**
  - For Storage subsystems:
    - Discovery
    - Monitoring
    - Configuration (for example, creating volumes)
    - Performance management
- **Fabric Management:**
  - For fabrics:
    - Discovery
    - Monitoring
    - Configuration (for example, zoning)
    - Performance management
- **Tape Management:**
  - Discovery
  - Monitoring

### 1.6.2 IBM TotalStorage Productivity Center Limited Edition

The TotalStorage Productivity Center Limited Edition is a special edition of IBM TotalStorage Productivity Center for clients of DS4000, DS6000, DS8000, SAN Volume Controller storage systems, and selected IBM tape libraries.

You can obtain technical support for either The TotalStorage Productivity Center Limited Edition or the TotalStorage Productivity Center Standard Edition on the appropriate IBM TotalStorage Productivity Center component page, which you get to from:

http://www-03.ibm.com/servers/storage/support/software/tpc

2

# Data management techniques

In this chapter, we look at examples of how you can use TotalStorage Productivity Center for Data (TPC for Data) to actively manage and control data in an environment. This discussion includes how to create scripts that TPC can call to perform custom-designed tasks.

We look at how TPC for Data, which is a component of TPC, can integrate with other tools, such as Tivoli Storage Manager, to create backup and archive policies. In this chapter, we describe:

- How TPC for Data can spot files that are missed by your backup software on Windows machines
- How to improve the way that TPC for Data spots duplicate files
- How to define scripted actions (automatic housekeeping and data pruning)
- How TPC can spot and alert you about sudden, unusual data growth in targeted areas
- How to create automatic data archiving tasks

## 2.1 How to spot files that are not backed up on Windows

It is not uncommon for an organization that manages a large number of servers to make an occasional mistake when making changes to their storage infrastructure.

Experience shows one common mishap is when one group of administrators adds a new storage volume to a server but forgets to inform the group that manages the backups. This condition can go unnoticed for months, or sometimes years. The backup administrators do not see any errors, because the existing jobs continue to complete without errors and simply miss the new volume altogether.

Usually the first time that anyone notices this mistake is when someone requests a restore, only to find that there is no data available for the new volume. Incidents, such as this one, are often very costly to an organization, and in some cases, these incidents can lead to legal action if the organization is required to retain copies of data for compliance reasons.

TPC for Data can help you to safeguard against this type of incident by watching for files that are not backed up regularly and then alerting you about this situation.

**Note:** This function is primarily available on Windows, because the Windows OS uses the concept of an “archive attribute” on a file. When you modify a file on a Windows machine, the archive flag is set against the file by the operating system. This function has existed since MS-DOS®. This function provides a way to indicate to the backup software that a file has changed and requires backup. The backup software then resets the archive attribute when it backs up the file, and so the cycle continues.

The ability to spot files works with any vendor’s backup software that resets the archive flag after copying a file. We show this function working with Tivoli Storage Manager.

In addition to the ability to spot these files, TPC has direct integration with Tivoli Storage Manager that you can use to trigger a backup of potentially unprotected data. TPC can work in a similar manner with other vendor’s backup products. The ability to call user-defined scripts allows TPC to pass a list of files to another program that backs up the files.

### 2.1.1 Set up a custom constraint and define data files and types to ignore

To make a “modified but not backed up” alert or report meaningful, we need to filter out all of the files that are intentionally excluded from a backup process. These files can be temporary files, or they can be directories or files with a given extension. If you have specified for your backup software to exclude this type of data, then you need to specify for TPC to ignore this type of data as well. Otherwise, meaningless alerts will trigger every day, and then when a real problem occurs, you will likely miss or ignore the alert.

You can achieve this by setting up a custom-defined constraint within TPC for Data that defines the exact criteria for the alerts that we want.

**Tip:** You will need the details of data that has been specifically EXCLUDED from backup cycles. Your backup administrator has these details. Often, organizations have a standard set of EXCLUDES that apply across many machines. If there are several distinct sets of excludes that are applied across your organization, you need to set up a separate constraint in TPC for each exclude and apply it to the appropriate machines.

## Define a constraint in TPC

Define a new constraint in TPC Data Manager. TPC Data Manager is built into TPC. A *constraint* is a file, filesystem, or something that we do not want in our environment. An example of a constraint is "We do not want files that have not been backed up, but ignore those files that we know are excluded." If the constraint is breached, then TPC has found data that is not backed up and will then do something about it based on what you have defined.

> **Note:** This section assumes a working knowledge of TPC for Data and is not a highly detailed step-by-step guide. We explain all key information; however, you need to be familiar with the panels.

The steps are:

1. Select **Data Manager** → **Policy Management** → **Constraints**. Right-click **Constraints** to create a new constraint as seen in Figure 2-1.

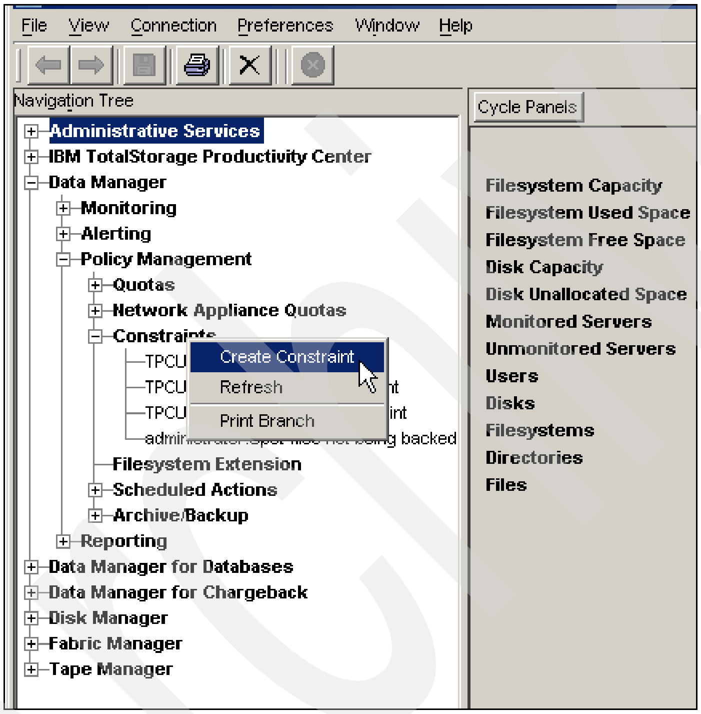

*Figure 2-1 Create constraint*

2. Next, select the filesystem in your environment to which this constraint will apply. This might be all filesystems or a subset if you plan to create more constraints with varying file exclude values.

> **Important:** By selecting all filesystems as in Figure 2-2, TPC automatically starts checking newly added filesystems that it finds on a host with an installed TPC agent. It does not matter if UNIX hosts are included, because they get filtered out by the condition logic.

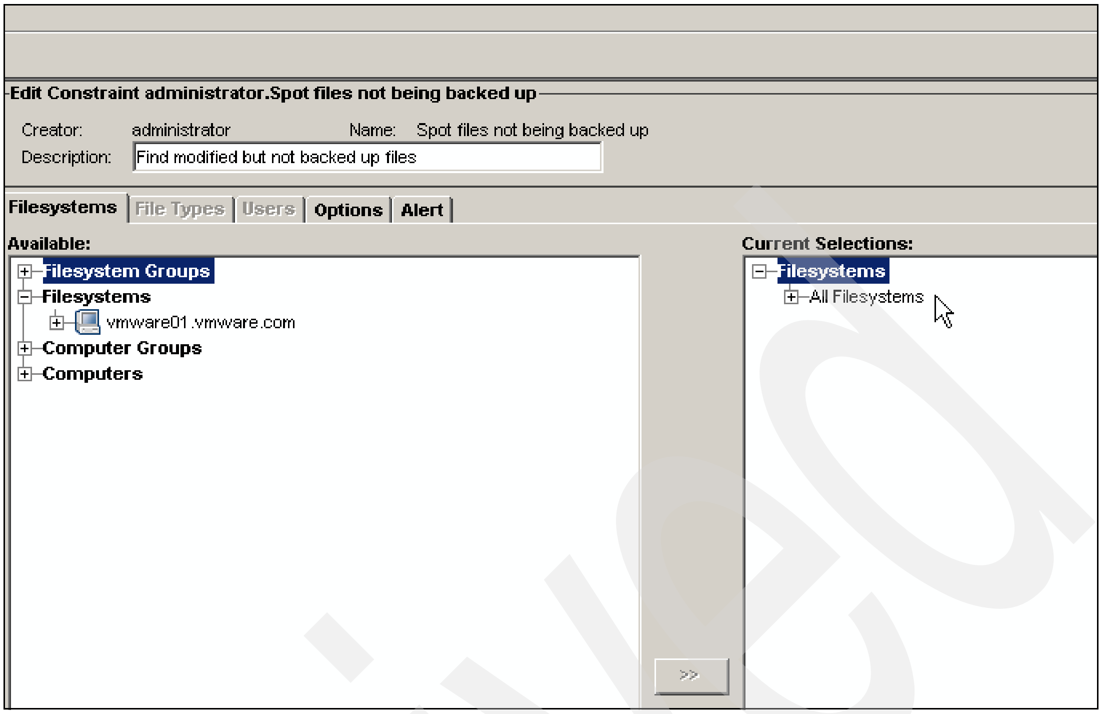

Edit Constraint administrator.Spot files not being backed up

Creator: administrator Name: Spot files not being backed up

Description: Find modified but not backed up files

*Figure 2-2 Select all filesystems*

3. After you make your selection, move to the **Options** tab to continue the constraint definition as seen in Figure 2-3.

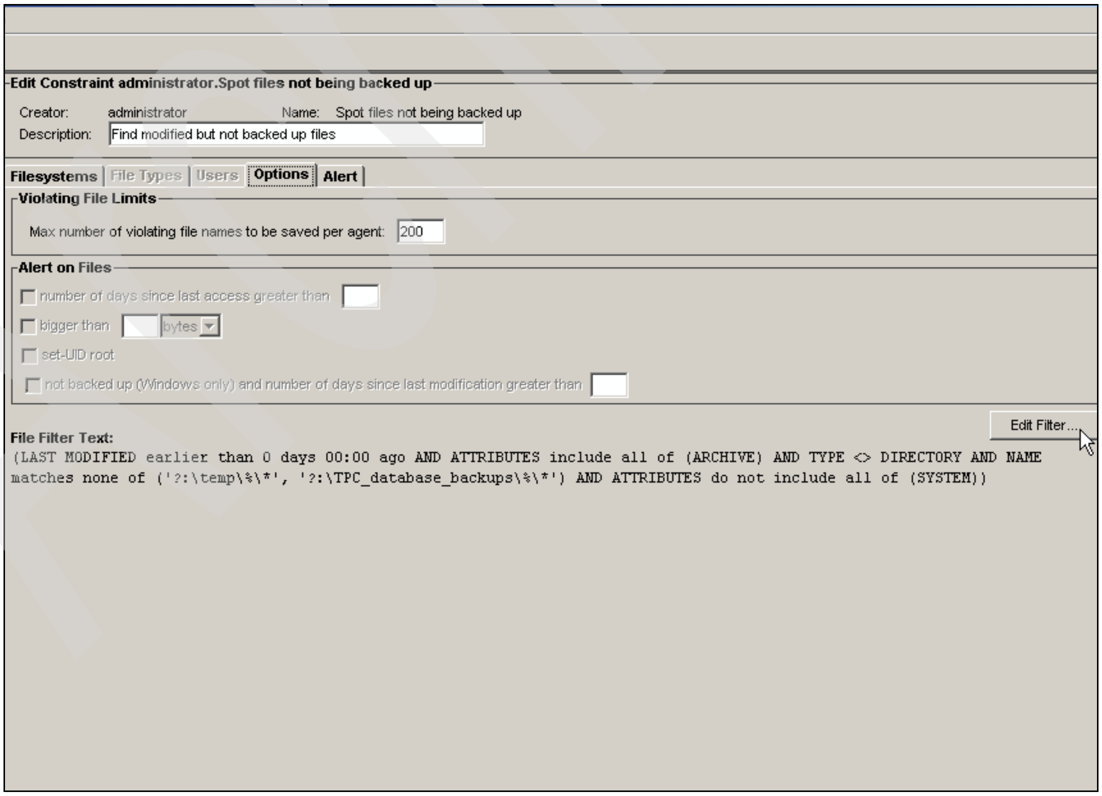

Edit Constraint administrator.Spot files not being backed up

Creator: administrator Name: Spot files not being backed up

Description: Find modified but not backed up files

Max number of violating file names to be saved per agent: 200

□ number of days since last access greater than

□ bigger than bytes

□ set-UID root

□ not backed up (Windows only) and number of days since last modification greater than

File Filter Text:

```
(LAST MODIFIED earlier than 0 days 00:00 ago AND ATTRIBUTES include all of (ARCHIVE) AND TYPE <> DIRECTORY AND NAME
matches none of ('?:\temp\%\*', '?:\TPC_database_backups\%\*') AND ATTRIBUTES do not include all of (SYSTEM))
```

*Figure 2-3 Define constraint filter*

4. You see in Figure 2-3 on page 16 that the constraint is already defined.

   Configure the following:

   a. Give your constraint a name. In our case, the name is "Find modified but not backed up files."

   b. Set the "Max number of violating file names to be saved per agent" as shown in our scenario.

> **Tip:** When TPC finds files that meet this constraint on a machine, TPC records the violating file names in the TPC repository database. This amount of names saved in TPC to your specified maximum number of violating file names limits the number of file names that TPC saves. This limit also keeps the size of the TPC repository database under control, while at the same time, alerting you that these violating files exist.
>
> For example, TPC might come across a newly added filesystem that has not been backed up but has 200,000 new files in it. It is not necessary to know all of the file names in this filesystem, because the alerting limit of 200 in our scenario signals the administrators to see what is happening.
>
> With this in mind, set this max number of violating file names value to a figure that you think is appropriate in your environment.

   c. Create a filter to spot files that are not backed up. To create this filter, click **Edit Filter** as in Figure 2-3 on page 16.

      The criteria that we used in our definition of "modified but not backed up files" in this example are:

      - File last modified more than seven days ago
      - File attribute = ARCHIVE
      - Type is not a directory
      - File names are not in "temp" or "TPC_database_backup" directories in any filesystem. These file name patterns need to match the file name patterns of the files excluded in your backup software.
      - Attribute is not SYSTEM

   Figure 2-4 on page 18 shows how these criteria are fully expressed in TPC.

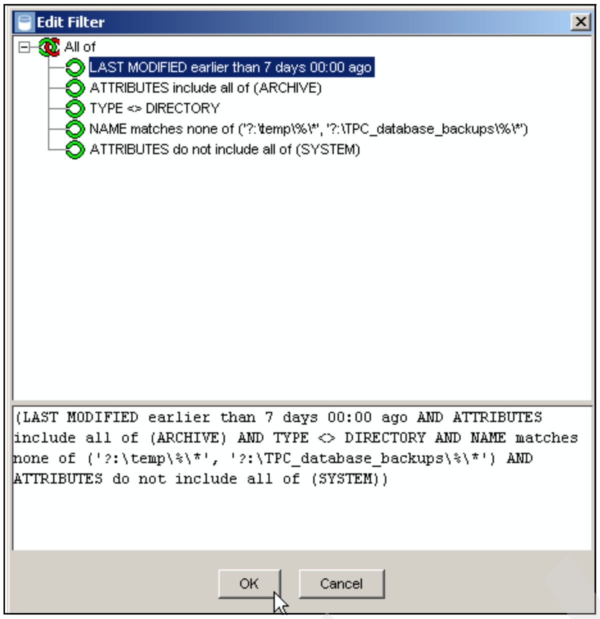

```
All of
    LAST MODIFIED earlier than 7 days 00:00 ago
    ATTRIBUTES include all of (ARCHIVE)
    TYPE <> DIRECTORY
    NAME matches none of ('?:\temp\%\*', '?:\TPC_database_backups\%\*')
    ATTRIBUTES do not include all of (SYSTEM)
```

```
(LAST MODIFIED earlier than 7 days 00:00 ago AND ATTRIBUTES
include all of (ARCHIVE) AND TYPE <> DIRECTORY AND NAME matches
none of ('?:\temp\%\*', '?:\TPC_database_backups\%\*') AND
ATTRIBUTES do not include all of (SYSTEM))
```

*Figure 2-4 Constraint file filter definition*

5. Start adding conditions by right-clicking on an empty filter panel as in Figure 2-5 and start by creating a new group.

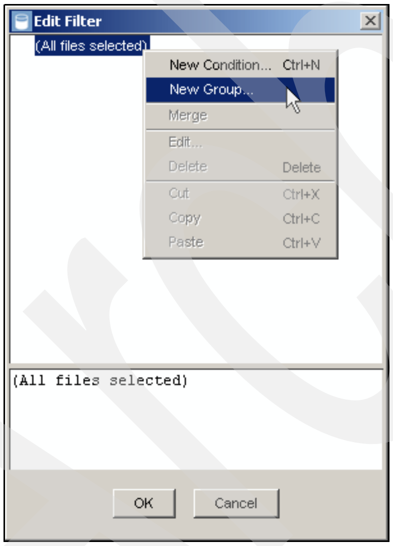

```
(All files selected)
```

*Figure 2-5 Add new filter group*

6. Select **All of** in the list box for grouping-type. In other words, a logical AND for all conditions as seen in Figure 2-6.

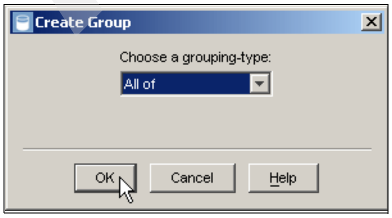

Choose a grouping-type: All of

*Figure 2-6 Create an All of group*

7. Now, start to create the various filter conditions by right-clicking and selecting **New Condition** as shown in Figure 2-7.

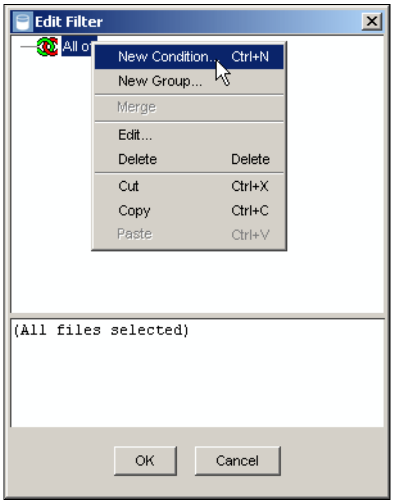

(All files selected)

*Figure 2-7 Create new condition*

8. When you add the files and directories that you know are excluded from the backup within your backup software, use standard TPC pattern matching wild cards.

   The example in Figure 2-8 shows:

   – ?:\temp\%*

     This excludes all files in \temp directories and below in any filesystem.

   – ?:\TPC_database_backups\%*

     This excludes all files in \TPC_database_backups and below in any filesystem.

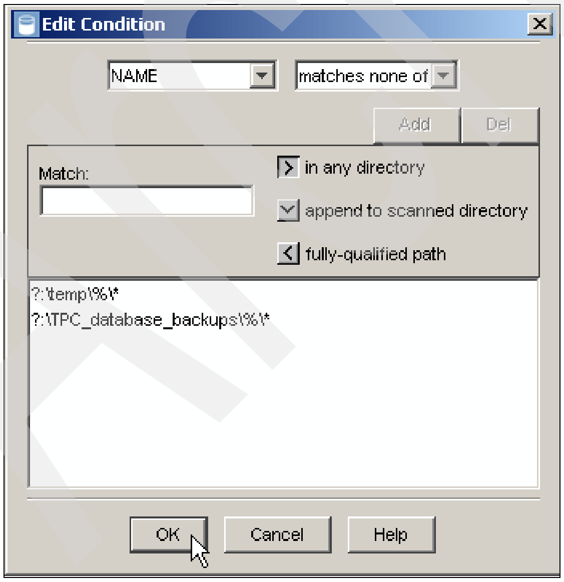

NAME matches none of

?:\temp\%*
?:\TPC_database_backups\%*

*Figure 2-8 Adding files to exclude*

9. Move to the Alert panel (shown in Figure 2-9 on page 20) and set the **Violating Files Consume More Than** condition to 1 Kilobytes. You can also enter an e-mail address to alert a person that files are not getting backed up.

This means that if TPC finds more than 1 Kilobyte of files, which meet the conditions that you specified in the file filter, the constraint is triggered, and an e-mail is sent to the addresses shown in Figure 2-9.

> **Note:** For TPC to send e-mail alerts, you must already have configured an SMTP server to TPC in the configuration panels located in the TPC menus under **Administrative Services → Configuration → Alert Disposition.**

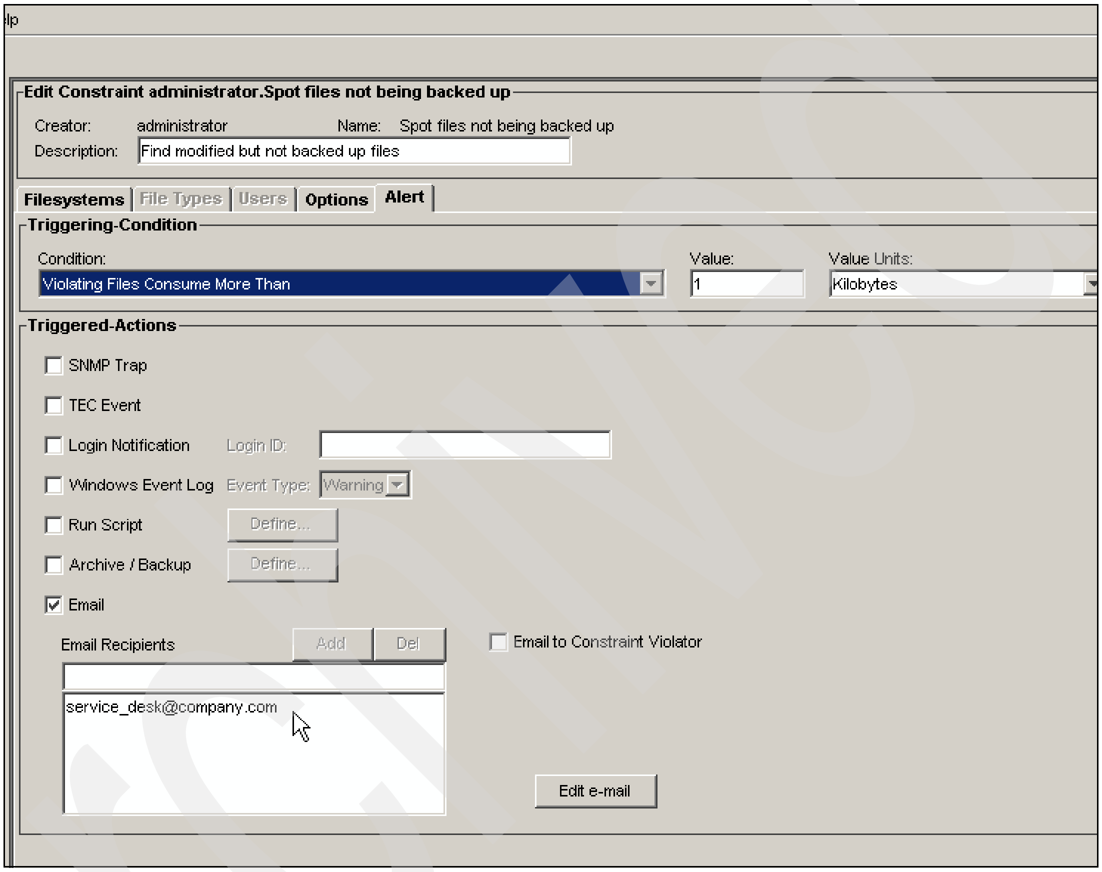

Edit Constraint administrator.Spot files not being backed up

Creator: administrator    Name: Spot files not being backed up

Description: Find modified but not backed up files

Triggering-Condition

Condition: Violating Files Consume More Than    Value: 1    Value Units: Kilobytes

Triggered-Actions

- ☐ SNMP Trap
- ☐ TEC Event
- ☐ Login Notification    Login ID:
- ☐ Windows Event Log    Event Type: Warning
- ☐ Run Script    Define...
- ☐ Archive / Backup    Define...
- ☑ Email

Email Recipients    Add    Del    ☐ Email to Constraint Violator

service_desk@company.com

*Figure 2-9 Define alerting and violating condition*

10. Now, save the constraint by clicking Ctrl+S. A Save As panel appears. Enter a name for this constraint as seen in Figure 2-10. Click **OK**.

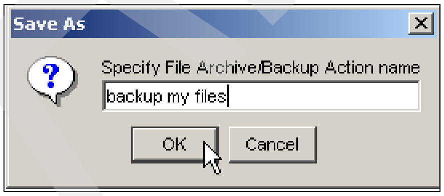

Save As

Specify File Archive/Backup Action name

backup my files

*Figure 2-10 Save constraint*

## Set up backup software

Now that you have defined the constraint, TPC looks for files that meet the constraint conditions every time that a scan is performed against the applicable host or filesystem.

Assuming that your backup software has reset the archive flag correctly on your files and that you have set the file and directory exclusion correctly, a constraint will be triggered only if there are files that the backup process misses.

> **Note:** By default, Tivoli Storage Manager does not reset the archive flag on a file when it backs it up. You need to configure an option file setting on the Tivoli Storage Manager Windows clients and then run one normal backup process (incremental or full). This allows Tivoli Storage Manager to set the archive flag on files to be consistent with what Tivoli Storage Manager has previously stored.
>
> If you use another vendor's backup software, check how it handles the archive flag. Configuration might be necessary.

## Configure Tivoli Storage Manager to reset the archive flag

To activate archive flag handling on Tivoli Storage Manager Windows clients, set the following option setting in the client option file dsm.opt. This option can also be set in a client option set. If you do not know how to do this, consult your Tivoli Storage Manager administrator.

```
dsm.opt setting
```

**RESETARCHIVEATTRIBUTE    YES**

Figure 2-11 shows the Tivoli Storage Manager client GUI preferences panel. Check **Reset archive attribute**, then click **Apply,** and close/restart the client for it to take effect.

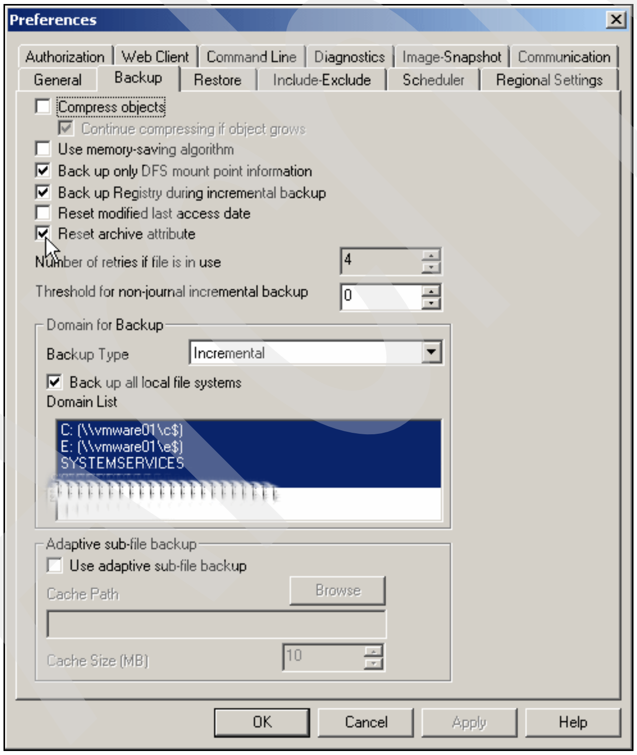

Preferences

Authorization | Web Client | Command Line | Diagnostics | Image-Snapshot | Communication
General | Backup | Restore | Include-Exclude | Scheduler | Regional Settings

☐ Compress objects
☑ Continue compressing if object grows
☐ Use memory-saving algorithm
☑ Back up only DFS mount point information
☑ Back up Registry during incremental backup
☐ Reset modified last access date
☑ Reset archive attribute

Number of retries if file is in use: 4
Threshold for non-journal incremental backup: 0

Domain for Backup
Backup Type: Incremental
☑ Back up all local file systems
Domain List
C: (\\vmware01\c$)
E: (\\vmware01\e$)
SYSTEMSERVICES

Adaptive sub-file backup
☐ Use adaptive sub-file backup
Cache Path
Cache Size (MB): 10

*Figure 2-11 Set the archive flag on the Tivoli Storage Manager preference panel*

After you configure Tivoli Storage Manager, Tivoli Storage Manager must go through one normal backup cycle for it to set the archive flag.

**Considerations:**

- There are other applications, such as Ntbackup.exe, which also manipulate or examine the Windows archive attribute. You should understand the ramifications of using this option in conjunction with these products.
- This option does not apply to archive, image, or server-free operations.
- This option does not apply to backing up Windows 2000 Server or Windows XP system objects, nor to backing up Windows Server® 2003 System State or System Services.
- This option only applies to files; it does not apply to directories.
- The use of this option during incremental backup applies only to those operations that require the examination of the stored client attributes on the server. Therefore, this option does not apply to Journal Based Backup or incremental by date processing. Also, this option does not apply to files that are excluded from backup and therefore not examined during an incremental backup operation.
- The client does not guarantee the accuracy of the current Windows archive attribute setting. For example, this option is selected, and a file examined by a reporting product indicates that the Windows archive attribute is OFF. This does not necessarily equate to having a valid, active backup version of a file on the server. This situation might occur due to several factors, including the following:
  - The third-party product is manipulating the Windows archive attribute.
  - The file space was deleted from the server.
  - The backup tape was lost or destroyed.
- This option can be set on the server by using the server client-options set feature.
- This option does not affect restore processing.

### Running a constraint report

The constraint condition will be checked each time that any scan runs against the machines and a filesystem is specified in the constraint.

If you have specified an e-mail address in the constraint alerting panel, you will receive one e-mail for each filesystem containing files that are not backed up.

To get a report of the files that meet the defined constraint, run a constraint violation report as in Figure 2-12 on page 23. Click **Generate Report**.

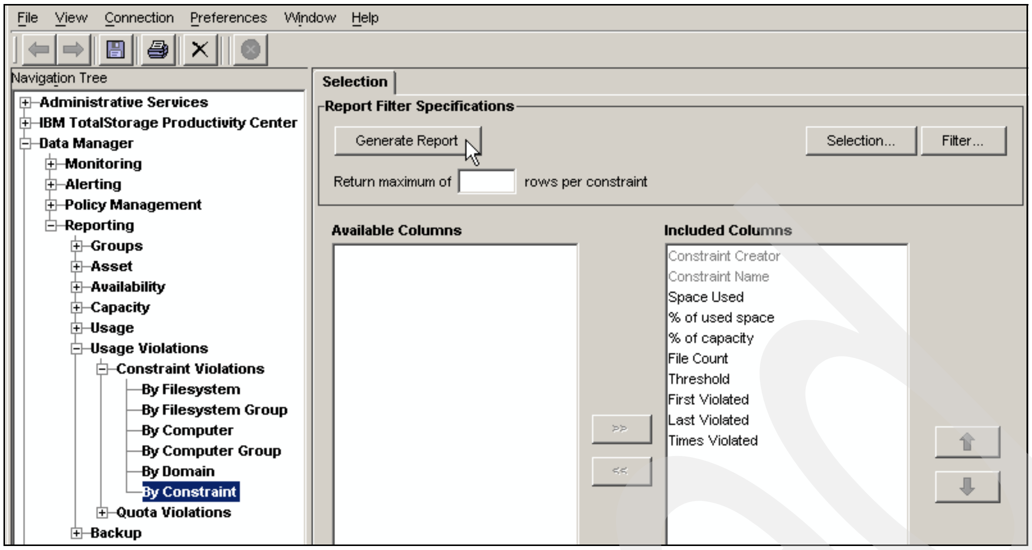

*Figure 2-12 Run a constraint report*

To show the violating files that are not being backed up, right-click the constraint name and select **Show Violating Files** as in Figure 2-13.

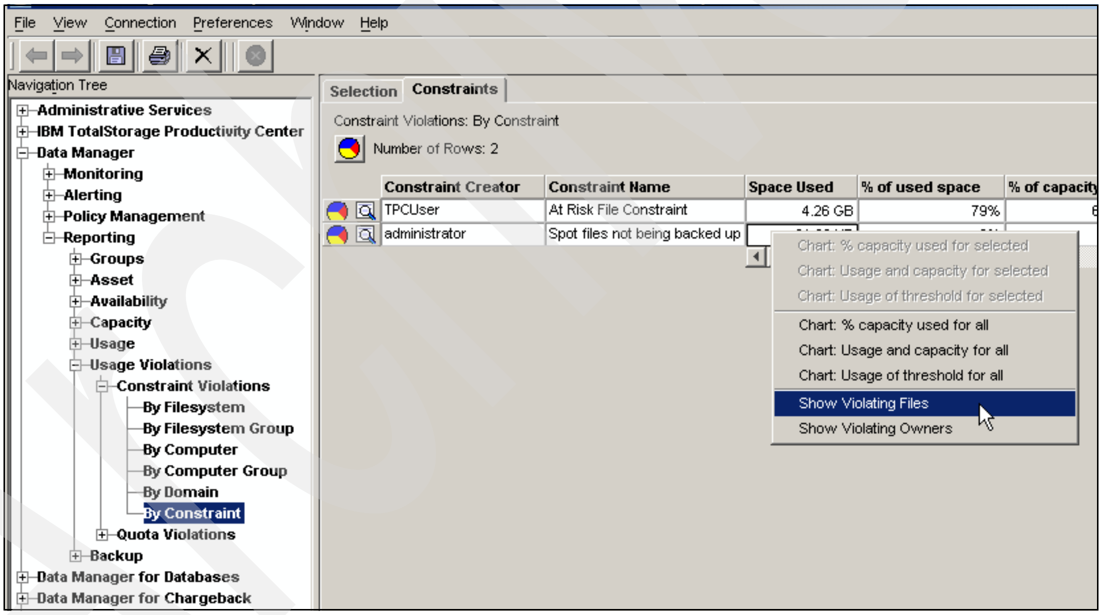
File View Connection Preferences Window Help
Navigation Tree
Administrative Services
IBM TotalStorage Productivity Center
Data Manager
Monitoring
Alerting
Policy Management
Reporting
Groups
Asset
Availability
Capacity
Usage
Usage Violations
Constraint Violations
By Filesystem
By Filesystem Group
By Computer
By Computer Group
By Domain
By Constraint
Quota Violations
Backup
Data Manager for Databases
Data Manager for Chargeback
Selection
Constraints
Chart: % capacity used for selected
Chart: Usage and capacity for selected
Chart: Usage of threshold for selected
Chart: % capacity used for all
Chart: Usage and capacity for all
Chart: Usage of threshold for all
Show Violating Files
Show Violating Owners

Constraint Violations: By Constraint

Number of Rows: 2

| Constraint Creator | Constraint Name | Space Used | % of used space | % of capacity |
|---|---|---|---|---|
| TPCUser | At Risk File Constraint | 4.26 GB | 79% | 6 |
| administrator | Spot files not being backed up | | | |

*Figure 2-13 Show violating files*

Figure 2-14 on page 24 shows a report of the files that violate this constraint and therefore are not getting backed up by your backup software.

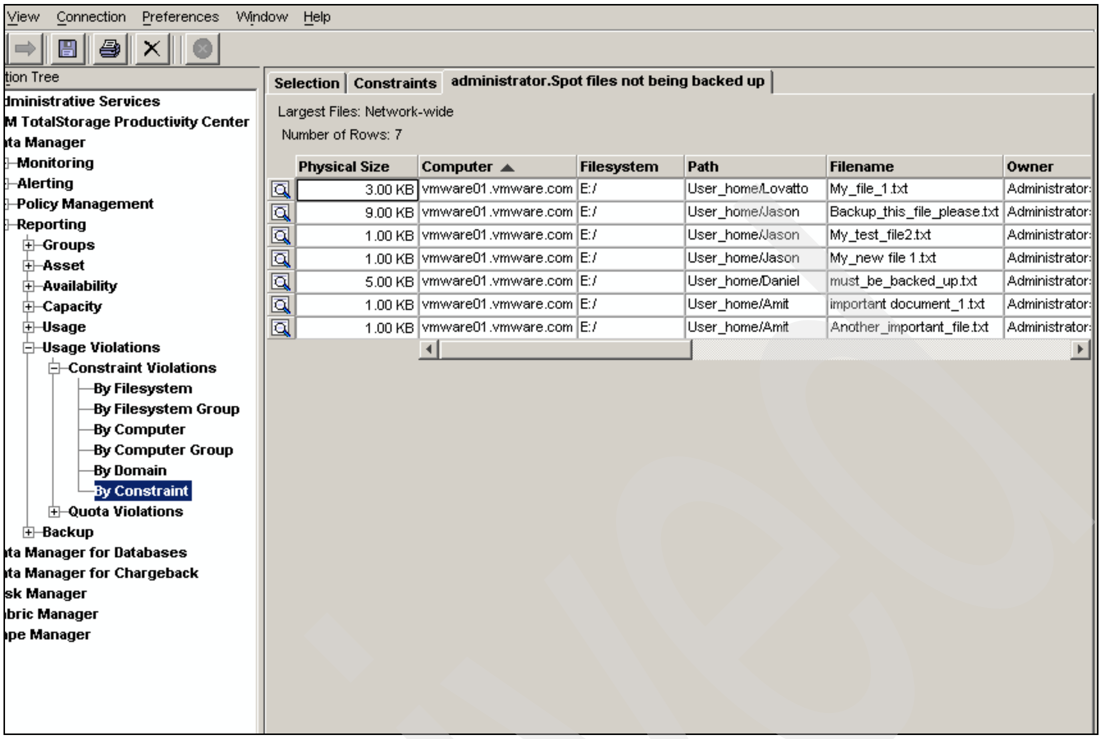

Largest Files: Network-wide

Number of Rows: 7

| Physical Size | Computer | Filesystem | Path | Filename | Owner |
|---|---|---|---|---|---|
| 3.00 KB | vmware01.vmware.com | E:/ | User_home/Lovatto | My_file_1.txt | Administrator: |
| 9.00 KB | vmware01.vmware.com | E:/ | User_home/Jason | Backup_this_file_please.txt | Administrator: |
| 1.00 KB | vmware01.vmware.com | E:/ | User_home/Jason | My_test_file2.txt | Administrator: |
| 1.00 KB | vmware01.vmware.com | E:/ | User_home/Jason | My_new file 1.txt | Administrator: |
| 5.00 KB | vmware01.vmware.com | E:/ | User_home/Daniel | must_be_backed_up.txt | Administrator: |
| 1.00 KB | vmware01.vmware.com | E:/ | User_home/Amit | important document_1.txt | Administrator: |
| 1.00 KB | vmware01.vmware.com | E:/ | User_home/Amit | Another_important_file.txt | Administrator: |

*Figure 2-14 Report of files not being backed up*

Use this information to take remedial steps to ensure that these files are backed up in the future or excluded (by editing the constraint filter) from the report if you are satisfied that these files should be excluded.

## 2.2 Enhance duplicate file searching

This section looks at how TPC searches for and reports on duplicate files in an environment and what configuration changes you can make to customize its focus when necessary.

### 2.2.1 How does TPC identify a duplicate file

There have been a number of misconceptions about how TPC identifies duplicate files in the past. By default, TPC does not find every instance of a duplicate file across your environment.

TPC performs duplicate file processing on the file names that are stored in its repository only. By default, these file names are collected by the following profiles:

- TPCUser.Largest Files
- TPCUser.Largest Orphans
- TPCUser.Most at Risk
- TPCUser.Most Obsolete

These profiles collect 20 file names by default per client for the largest files or the largest orphan files. Therefore, when TPC looks for duplicate files, it only looks at a small number of file names per machine.

Any duplicates that TPC finds are by definition in the top 20 largest, largest orphans, or most at risk files. In many, if not most situations, this limited level of duplicate spotting is too small to be of any great use. This level will not find the true extent of a duplicate file problem in this default configuration.

Conversely, if TPC attempts to perform duplicate file matching for all of the files in your environment and there are ten million (not an unrealistic figure these days) files through which TPC looks, clearly, the TPC repository becomes huge, even if only only one percent of your files are duplicate files.

What you need is a balanced approach. Configure TPC to look for more file names but target its focus at specific filesystems and file types where you suspect there is a problem or where the problem most likely exists.

For example, collecting all of the file names on the C: drives for all of your Windows machines only tells you what you already know in that the same file names exist on all of them. Storing all these file names in TPC for this reason represents poor reporting value for the effort involved in collecting, storing, and managing the file name data.

There is greater value in perhaps collecting file names of all the Microsoft® Office file types (*.doc, *.xls, *.ppt) and media files (*.mp3 or *.avi) in users' file and print directories. Reporting at this level can spot spreadsheets that were created by one person but then e-mailed to many people and then detached multiple times. Also, TPC can spot video files or jokes from external sources that employees have shared with others through e-mail and detaching.

**Note:** TPC matches duplicates by file name and file size only. TPC does not open or examine the content of any files or perform any kind of checksum processing.

### Create a profile to target file names for duplicate processing

With an understanding of how TPC looks for duplicate files, you can create targeted profiles that can collect large amounts of file names but are targeted at specific servers (perhaps file and print) and also focus on specific file types.

From the left panel in Figure 2-15 on page 26, follow the path **Data Manager** → **Monitoring** and right-click **Profiles** to create a new profile. In the right panel, enter a profile **Description** and check the boxes shown, and then enter a number of files.

This example specifies 1,000 files for each type. This means that for each server targeted by this profile, TPC will gather up to 4,000 file names overall for the largest, the most obsolete, the most at risk, and the largest orphan files.

Adjust these numbers as appropriate to your environment. You might decide to look for the 4,000 largest files and not bother with the other types.

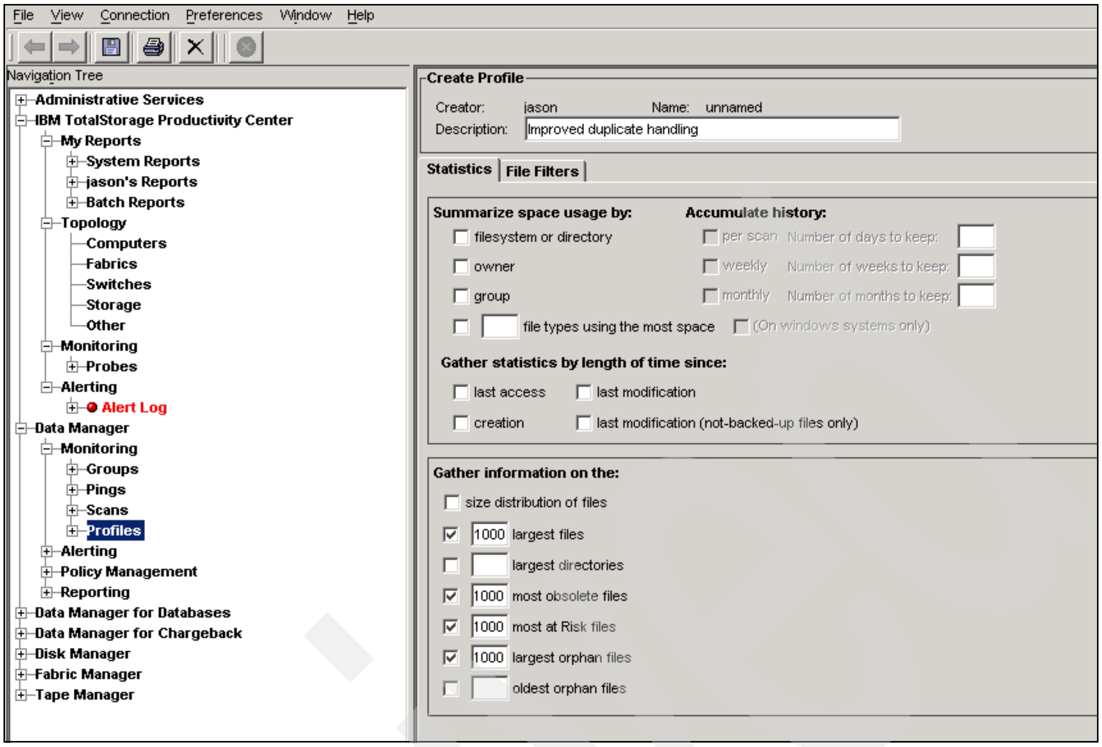

Navigation Tree
- Administrative Services
- IBM TotalStorage Productivity Center
  - My Reports
    - System Reports
    - jason's Reports
    - Batch Reports
  - Topology
    - Computers
    - Fabrics
    - Switches
    - Storage
    - Other
  - Monitoring
    - Probes
  - Alerting
    - ● Alert Log
- Data Manager
  - Monitoring
    - Groups
    - Pings
    - Scans
    - Profiles
  - Alerting
  - Policy Management
  - Reporting
- Data Manager for Databases
- Data Manager for Chargeback
- Disk Manager
- Fabric Manager
- Tape Manager

Create Profile

Creator: jason Name: unnamed

Description: Improved duplicate handling

Statistics | File Filters

Summarize space usage by:
- □ filesystem or directory
- □ owner
- □ group
- □ [ ] file types using the most space

Accumulate history:
- □ per scan Number of days to keep: [ ]
- □ weekly Number of weeks to keep: [ ]
- □ monthly Number of months to keep: [ ]
- □ (On windows systems only)

Gather statistics by length of time since:
- □ last access □ last modification
- □ creation □ last modification (not-backed-up files only)

Gather information on the:
- □ size distribution of files
- ☑ 1000 largest files
- □ [ ] largest directories
- ☑ 1000 most obsolete files
- ☑ 1000 most at Risk files
- ☑ 1000 largest orphan files
- □ [ ] oldest orphan files

*Figure 2-15 Create a profile to collect file names for duplicate matching*

> **Tip:** To understand the impact that this type of profile has on the size of the TPC database, refer to Chapter 3, “IBM Total Productivity Center database considerations” on page 61.

Now, move to the **File Filters** tab and create a filter to target this profile at the types of files in which you are interested.

If you are interested in all file types, leave the file filters blank.

Click **New Condition** to the right of the panel in Figure 2-16 on page 27 to build the file filter.

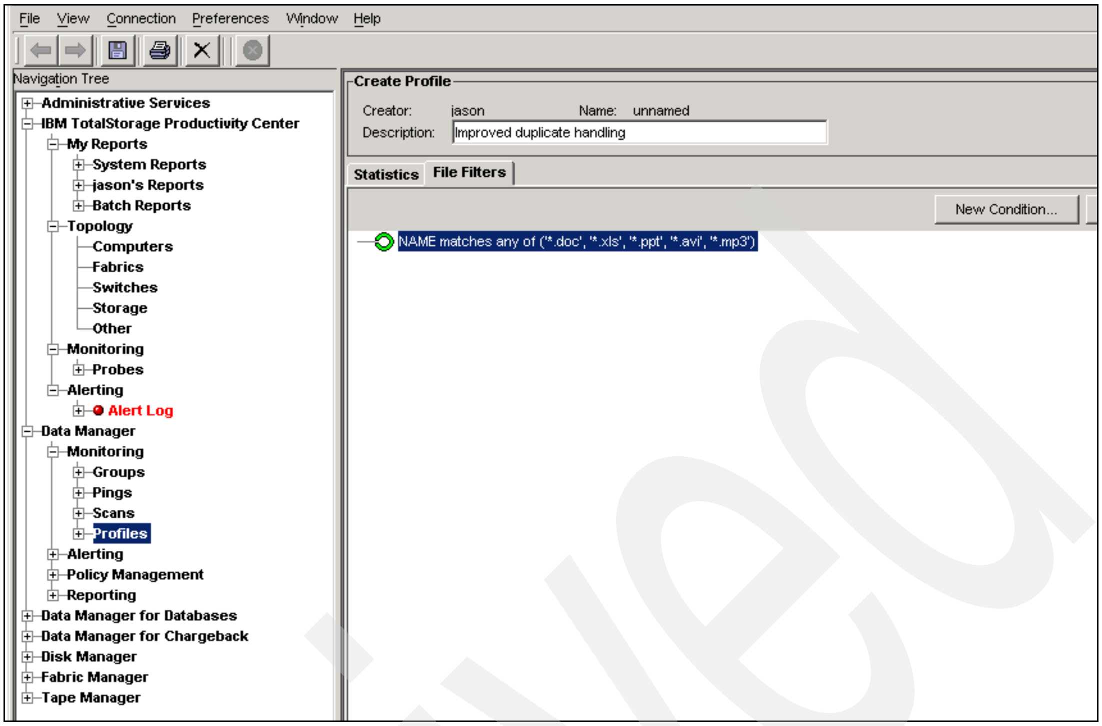

Creator: jason Name: unnamed

Description: Improved duplicate handling

NAME matches any of ('*.doc', '*.xls', '*.ppt', '*.avi', '*.mp3')

*Figure 2-16 Create a file filter*

Figure 2-17 shows the Create Condition dialog box and you can see that we have added a number of file extension names that are of interest to us. You can add more conditions that target file owners, file age, or other criteria as necessary.

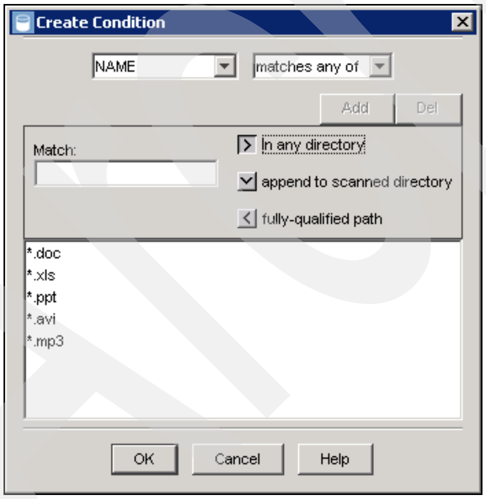

NAME matches any of

☒ In any directory

☑ append to scanned directory

☐ fully-qualified path

```
*.doc
*.xls
*.ppt
*.avi
*.mp3
```

*Figure 2-17 Add file types to filter*

When you have completed your profile configuration, click **OK**. Then click Ctrl+S to save the profile and give it a unique name as in Figure 2-18 on page 28. Click **OK**.

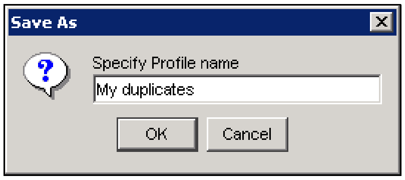

Specify Profile name: My duplicates

Figure 2-18   Save profile

## 2.2.2 Create a targeted scan job

After you define the new profile, add it to a scan job so that it can start to collect data. It is likely that the existing scan jobs in your environment are general and wide ranging. In other words, if you add this new profile to your existing scan jobs, the new profile is applied to all servers and all filesystems in the environment. This gathers a lot of data and enlarges the amount of TPC data unnecessarily.

To target this new profile at the servers and filesystems that are of most interest for duplicate file reporting, create a new scan job with a limited scope of servers and a limited scope of filesystems.

Right-click **Scans** to create a new scan job. Then, select the computers and filesystems at which you want to target this profile for duplicate file spotting. Figure 2-19 shows that we selected only computer *tivoli30*.

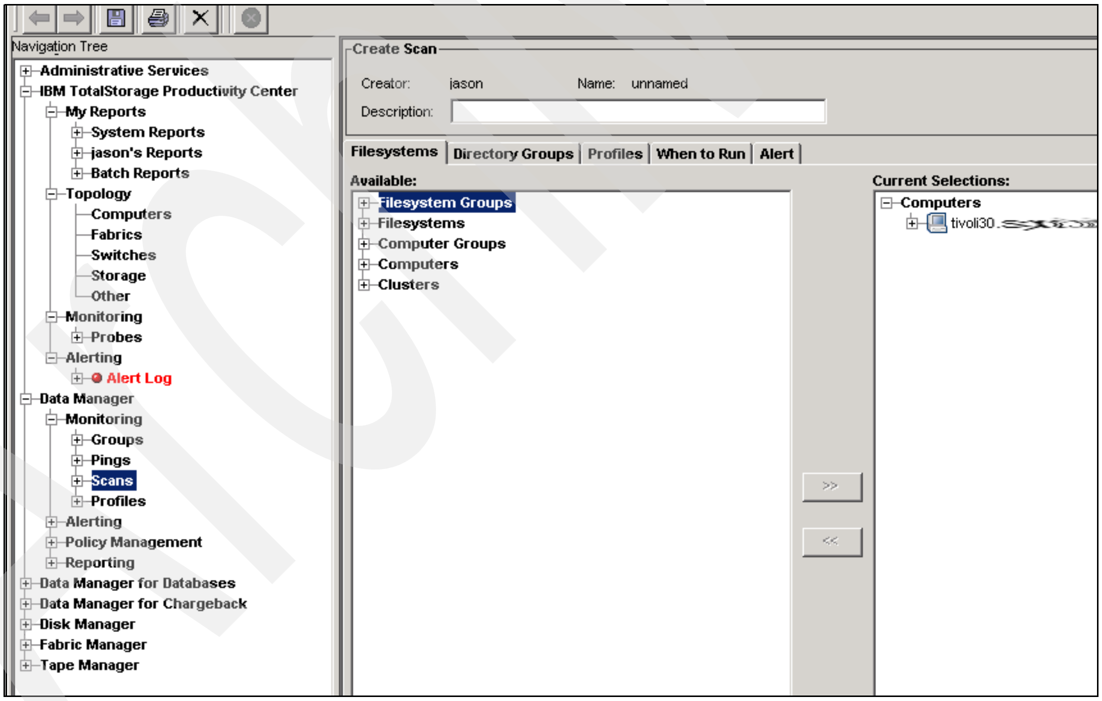

Creator: jason    Name: unnamed

Description:

Current Selections:
Computers
tivoli30.

Figure 2-19   Create a new scan job

Move to the Profiles panel by selecting the **Profiles** tab as in Figure 2-20 on page 29 and select the new profile for inclusion in the scan job.

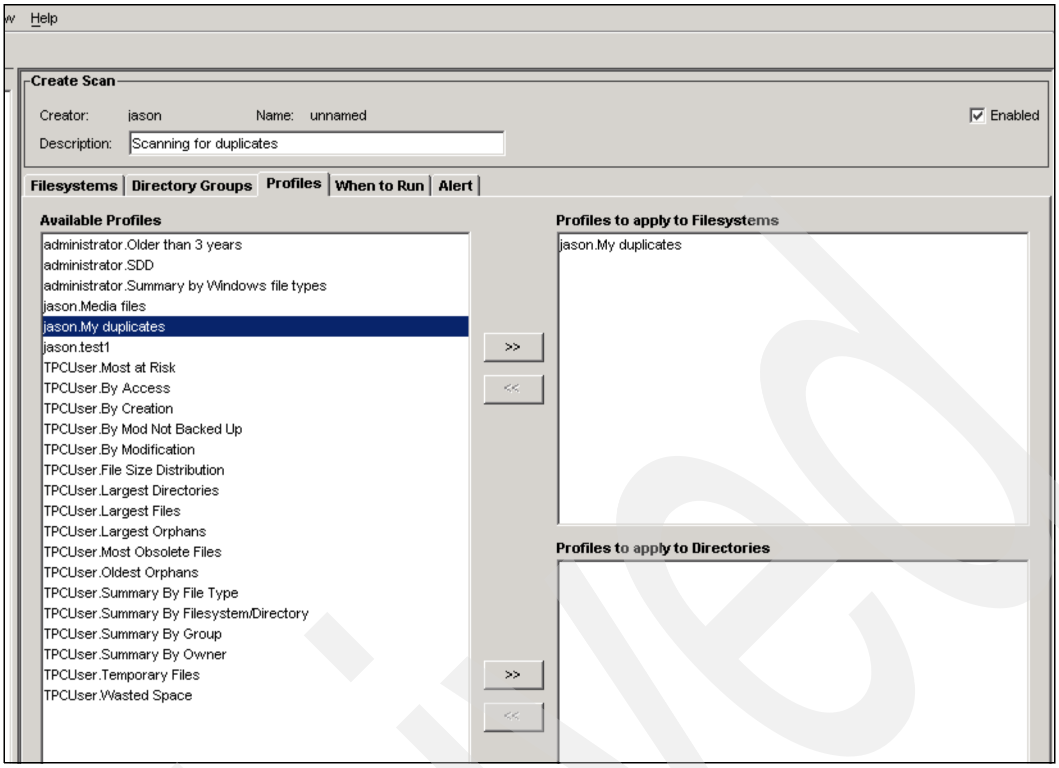

**Create Scan**

Creator: jason Name: unnamed ☑ Enabled

Description: Scanning for duplicates

**Available Profiles**

- administrator.Older than 3 years
- administrator.SDD
- administrator.Summary by Windows file types
- jason.Media files
- jason.My duplicates
- jason.test1
- TPCUser.Most at Risk
- TPCUser.By Access
- TPCUser.By Creation
- TPCUser.By Mod Not Backed Up
- TPCUser.By Modification
- TPCUser.File Size Distribution
- TPCUser.Largest Directories
- TPCUser.Largest Files
- TPCUser.Largest Orphans
- TPCUser.Most Obsolete Files
- TPCUser.Oldest Orphans
- TPCUser.Summary By File Type
- TPCUser.Summary By Filesystem/Directory
- TPCUser.Summary By Group
- TPCUser.Summary By Owner
- TPCUser.Temporary Files
- TPCUser.Wasted Space

**Profiles to apply to Filesystems**

- jason.My duplicates

**Profiles to apply to Directories**

*Figure 2-20 Add new profile to scan job*

Select the **When to Run** tab. On the panel, define how often you want to run this scan or a start time. In this case, we want to run the job immediately, so we click **Run Now** as opposed to scheduling a time to run this job. Note that **Enabled** in the upper right corner is checked.

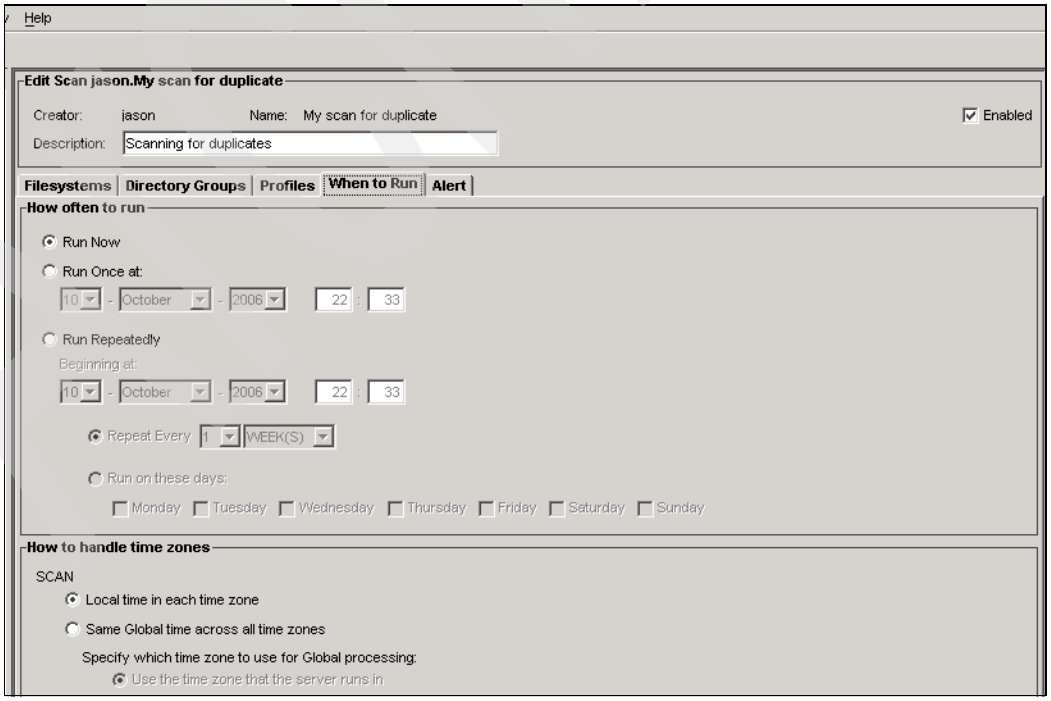

**Edit Scan jason.My scan for duplicate**

Creator: jason Name: My scan for duplicate ☑ Enabled

Description: Scanning for duplicates

**How often to run**

● Run Now

○ Run Once at:

10 - October - 2006 22 : 33

○ Run Repeatedly

Beginning at:

10 - October - 2006 22 : 33

● Repeat Every 1 WEEK(S)

○ Run on these days:

□ Monday □ Tuesday □ Wednesday □ Thursday □ Friday □ Saturday □ Sunday

**How to handle time zones**

SCAN

● Local time in each time zone

○ Same Global time across all time zones

Specify which time zone to use for Global processing:

● Use the time zone that the server runs in

*Figure 2-21 Define when to scan*

Save the new scan job by clicking Ctrl+S and give it a unique name as in Figure 2-22. Click **OK**.

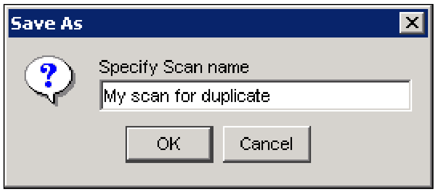

My scan for duplicate

Figure 2-22 Save new scan job

If the scan job is set to **Run Now,** a panel as in Figure 2-23 appears.

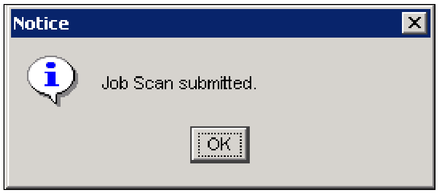

Job Scan submitted.

Figure 2-23 Job Scan submitted

The time that is necessary to complete this scan depends on many factors in your environment.

### 2.2.3 Running duplicate reports

Duplicate file reports differ from other reporting types in that you do not specify a profile for the report to use in order to select data. Instead, duplicate matching looks through all stored file names in the TPC repository.

The example profile that we have previously created in "Create a profile to target file names for duplicate processing" on page 25 will have collected many new file names for the specified pattern match, in this case, *.doc, *.xls, *.ppt, and *.avi. To create a duplicate file report that is targeted at them, you need to build a filter into the report.

To generate a targeted duplicate file report, start by launching the report panel as in Figure 2-24 on page 31.

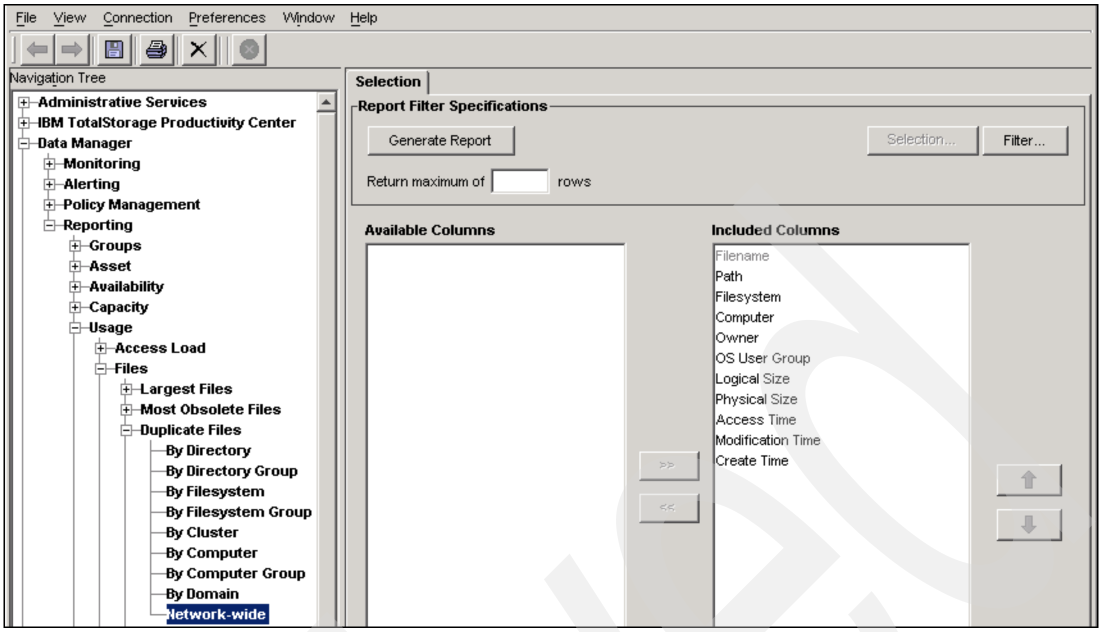

Included Columns: Filename, Path, Filesystem, Computer, Owner, OS User Group, Logical Size, Physical Size, Access Time, Modification Time, Create Time

*Figure 2-24 Running a a duplicate file report*

Click **Filter** in the upper right corner of the panel to build a custom filter. Figure 2-25 shows the filter prepared for *.doc, *.xls, *.ppt, *.avi, and *.mp3 file types. Click **OK** when you finish. When you generate the report, it will only contain duplicate files of the types that are specified in the Value 1 field.

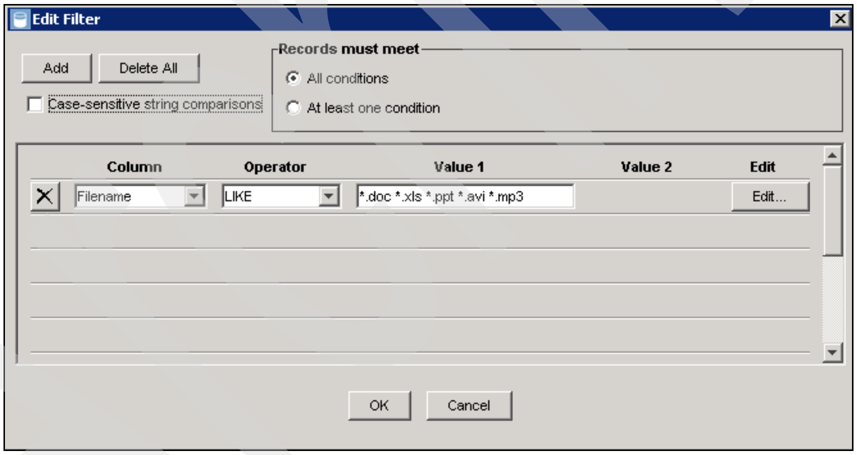

☐ Case-sensitive string comparisons

Records must meet
● All conditions
○ At least one condition

| Column | Operator | Value 1 | Value 2 | Edit |
|---|---|---|---|---|
| Filename | LIKE | *.doc *.xls *.ppt *.avi *.mp3 | | Edit... |

*Figure 2-25 Custom filter*

Figure 2-26 on page 32 shows the filtered report. In this example, you can see a number of duplicate copies of what seem to be music albums.

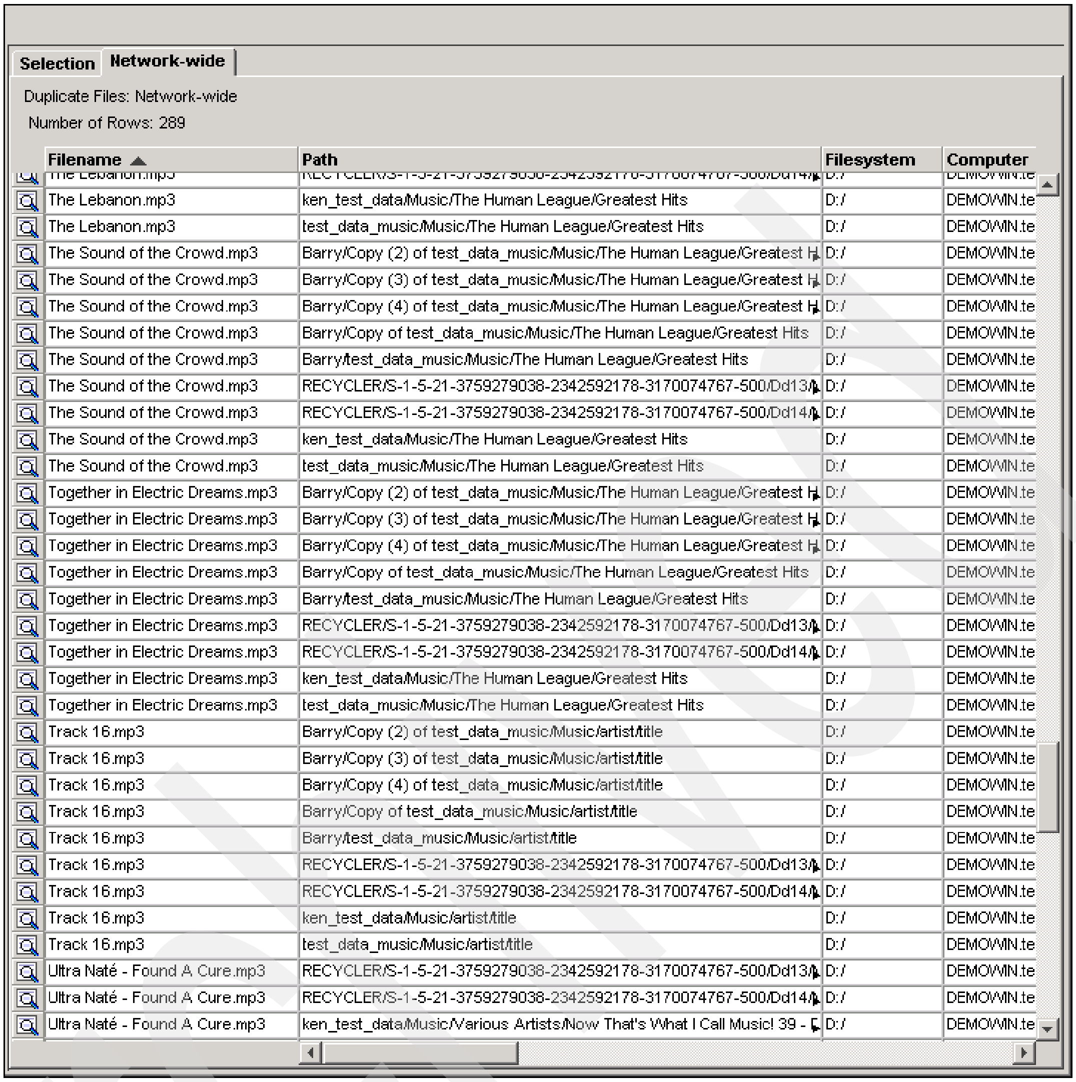

Duplicate Files: Network-wide
Number of Rows: 289

| Filename | Path | Filesystem | Computer |
|---|---|---|---|
| The Lebanon.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd14/ | D:/ | DEMOWIN.te |
| The Lebanon.mp3 | ken_test_data/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| The Lebanon.mp3 | test_data_music/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | Barry/Copy (2) of test_data_music/Music/The Human League/Greatest H | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | Barry/Copy (3) of test_data_music/Music/The Human League/Greatest H | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | Barry/Copy (4) of test_data_music/Music/The Human League/Greatest H | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | Barry/Copy of test_data_music/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | Barry/test_data_music/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd13/ | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd14/ | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | ken_test_data/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| The Sound of the Crowd.mp3 | test_data_music/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | Barry/Copy (2) of test_data_music/Music/The Human League/Greatest H | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | Barry/Copy (3) of test_data_music/Music/The Human League/Greatest H | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | Barry/Copy (4) of test_data_music/Music/The Human League/Greatest H | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | Barry/Copy of test_data_music/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | Barry/test_data_music/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd13/ | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd14/ | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | ken_test_data/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| Together in Electric Dreams.mp3 | test_data_music/Music/The Human League/Greatest Hits | D:/ | DEMOWIN.te |
| Track 16.mp3 | Barry/Copy (2) of test_data_music/Music/artist/title | D:/ | DEMOWIN.te |
| Track 16.mp3 | Barry/Copy (3) of test_data_music/Music/artist/title | D:/ | DEMOWIN.te |
| Track 16.mp3 | Barry/Copy (4) of test_data_music/Music/artist/title | D:/ | DEMOWIN.te |
| Track 16.mp3 | Barry/Copy of test_data_music/Music/artist/title | D:/ | DEMOWIN.te |
| Track 16.mp3 | Barry/test_data_music/Music/artist/title | D:/ | DEMOWIN.te |
| Track 16.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd13/ | D:/ | DEMOWIN.te |
| Track 16.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd14/ | D:/ | DEMOWIN.te |
| Track 16.mp3 | ken_test_data/Music/artist/title | D:/ | DEMOWIN.te |
| Track 16.mp3 | test_data_music/Music/artist/title | D:/ | DEMOWIN.te |
| Ultra Naté - Found A Cure.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd13/ | D:/ | DEMOWIN.te |
| Ultra Naté - Found A Cure.mp3 | RECYCLER/S-1-5-21-3759279038-2342592178-3170074767-500/Dd14/ | D:/ | DEMOWIN.te |
| Ultra Naté - Found A Cure.mp3 | ken_test_data/Music/Various Artists/Now That's What I Call Music! 39 - | D:/ | DEMOWIN.te |

*Figure 2-26 Filtered duplicate file name report*

## 2.3 Defining scripted actions

TPC can call external programs and scripts that are based on the results of quotas and constraints. This capability allows you to create custom functions that take action when predefined conditions arise.

Here are several examples, but you are not limited to these examples:

- **Alert to a third-party monitoring tool:** TPC can send outbound alerts in a number of ways using SNMP, Tivoli Enterprise™ Console (TEC), Windows event log, and e-mail.

  You might choose to send alerts to pagers or other third-party tools. You can use scripts to pass information to these tools in a custom-written format.

- **Delete and clean up unwanted files:** Set a constraint that looks for and deletes files that are known to contain viruses or prohibited content. You can set a constraint that looks for files that contain copyrighted material that is not permitted in your organization. You can use this function to prevent users from storing media files, such as MP3 and video, in their home directories. You can use this capability as a proactive space control mechanism to stop users from storing personal files that consume too much storage.

You can set TPC up to identify the violating files and invoke a script to delete the files.

- **Archive old files:** Set up a constraint to identify old files that have not been accessed for an extended period of time and to archive these files into a third-party product.

  While there is direct integration between TPC and Tivoli Storage Manager, you might use another vendor's backup or archive product. By using scripting functions, you can conceivably integrate TPC with any other vendor tools that support your planned activities.
- **Avoid full filesystems:** You can set a space trigger at 95% filesystem usage. Breaching this threshold can invoke a script that performs client-defined cleanup actions, such as cleaning up old "tmp" files or other data that can be deleted.

## 2.3.1 Scripting example of deleting unwanted *.tmp files

This example demonstrates how you can use scripting to delete files that unnecessarily consume storage and lead to "`filesystem full`" conditions if they are left unchecked.

By using these scripts, TPC can be proactive with these tasks and avoid potentially expensive unplanned downtime. In certain environments today, storage administrators might get an alert about a filesystem that is almost full. Storage administrators have two options: they can either increase the size of the filesystem or reduce the amount of data stored within the filesystem.

The filesystem might be full for a genuine business reason, perhaps temporary work files of application logs have been left to grow unchecked for too long. Administrators can often find themselves trying to resolve an emergency situation, which is inconvenient, stressful, and often unnecessary if you define certain simple automated processes to keep known culprits under control.

**Example description:** This example constraint deletes *.tmp files from the directory named temp in the Document and Settings areas of Windows users. If an application ends in an uncontrolled manner, *.tmp files can build up over time, because these files never get deleted.

The constraint looks for *.tmp files that have not been accessed for three days. This typically indicates that the application that created them is unlikely to ever clean them up.

### Create script to delete files

By default, user-written TPC scripts reside in:

- Windows: C:\Program Files\IBM\TPC\ca\subagents\data\scripts
- UNIX : <*install_dir*>/ep/subagents/TPC/Data/scripts

Each machine on which you want this script to execute needs a copy of it in its script directory.

Create a script as in Figure 2-27 on page 34 and place it on the machines on which you want it to run in the path that we just specified by operating system.

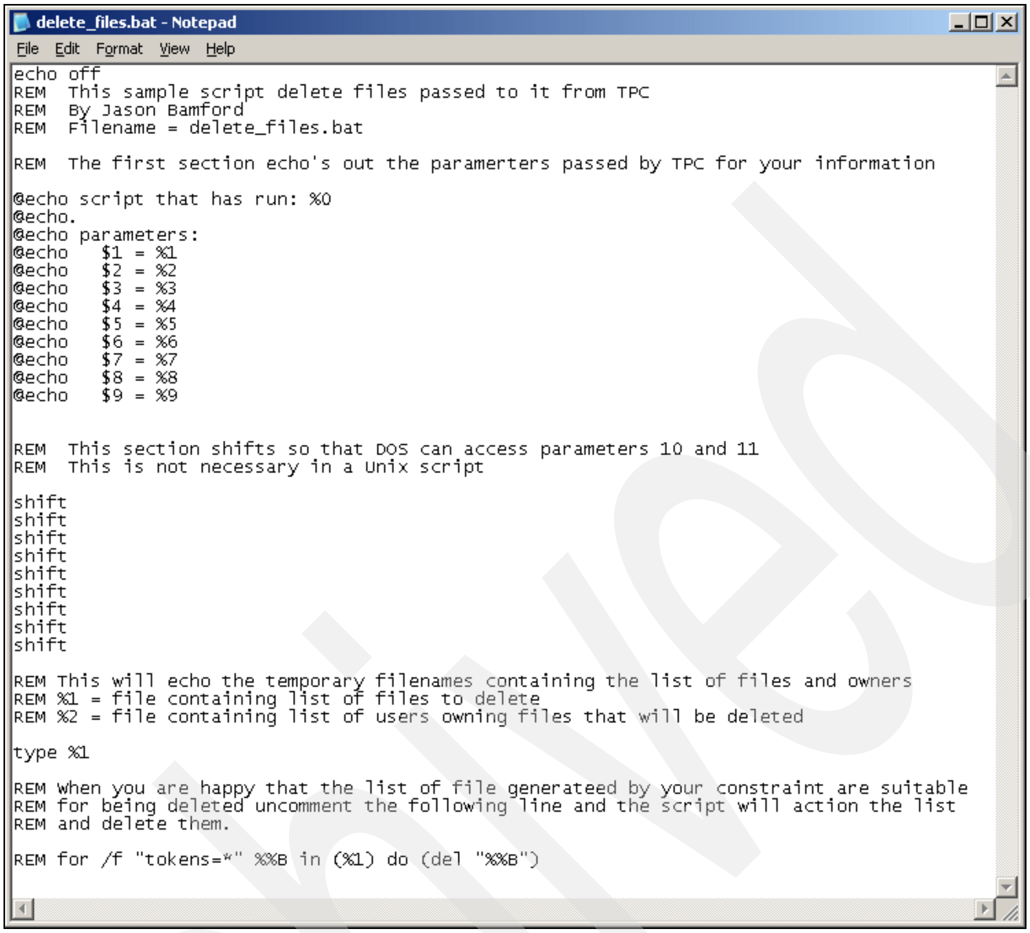

```
echo off
REM  This sample script delete files passed to it from TPC
REM  By Jason Bamford
REM  Filename = delete_files.bat

REM  The first section echo's out the paramerters passed by TPC for your information

@echo script that has run: %0
@echo.
@echo parameters:
@echo   $1 = %1
@echo   $2 = %2
@echo   $3 = %3
@echo   $4 = %4
@echo   $5 = %5
@echo   $6 = %6
@echo   $7 = %7
@echo   $8 = %8
@echo   $9 = %9


REM  This section shifts so that DOS can access parameters 10 and 11
REM  This is not necessary in a Unix script

shift
shift
shift
shift
shift
shift
shift
shift
shift

REM This will echo the temporary filenames containing the list of files and owners
REM %1 = file containing list of files to delete
REM %2 = file containing list of users owning files that will be deleted

type %1

REM When you are happy that the list of file generateed by your constraint are suitable
REM for being deleted uncomment the following line and the script will action the list
REM and delete them.

REM for /f "tokens=*" %%B in (%1) do (del "%%B")
```

*Figure 2-27 delete_files.bat sample script*

By default, this script only lists the files that it will delete, it does not delete them. After your constraint runs a few times and you are satisfied that the files selected for deletion are the correct files, remove the `REM` from the last line of the script. The next time that a scan runs, the constraint is triggered, and this script executes and deletes the files.

## Create constraint to identify old tmp files

Create a constraint that specifies the type of files that you want TPC to clean out on a regular basis. To create the constraint:

1. Follow the path **Administrative Services** → **Data Manager** → **Policy Management** → **Constraints**.
2. Start by naming your constraint and selecting the filesystems against which you want your constraint to run. The example in Figure 2-28 on page 35 shows all filesystems have been selected. If you have a multi-platform environment, select only the Windows machines, because this example is specific to Windows only.

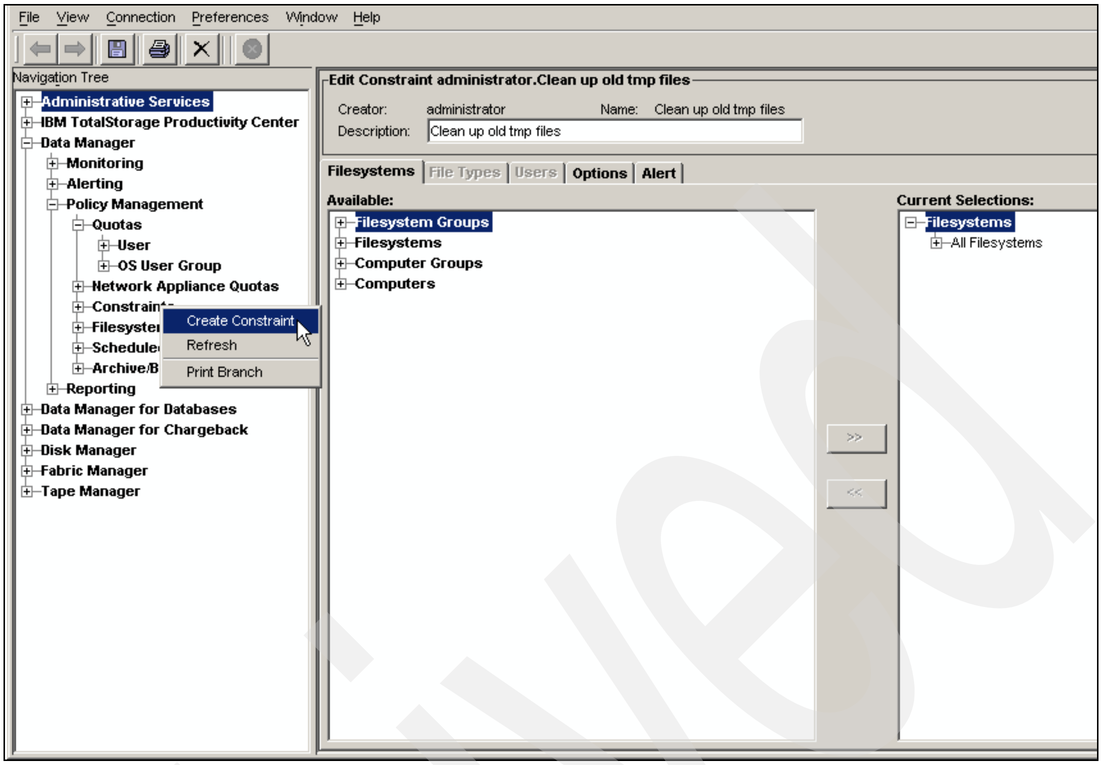

Edit Constraint administrator.Clean up old tmp files

Creator: administrator Name: Clean up old tmp files

Description: Clean up old tmp files

Figure 2-28 Create constraint for tmp files

3. Next, create a filter to identify the files that you want TPC to delete. Click **Edit Filter** on the upper right of the panel in Figure 2-29. An Edit Filter dialog box appears as in Figure 2-30 on page 36.

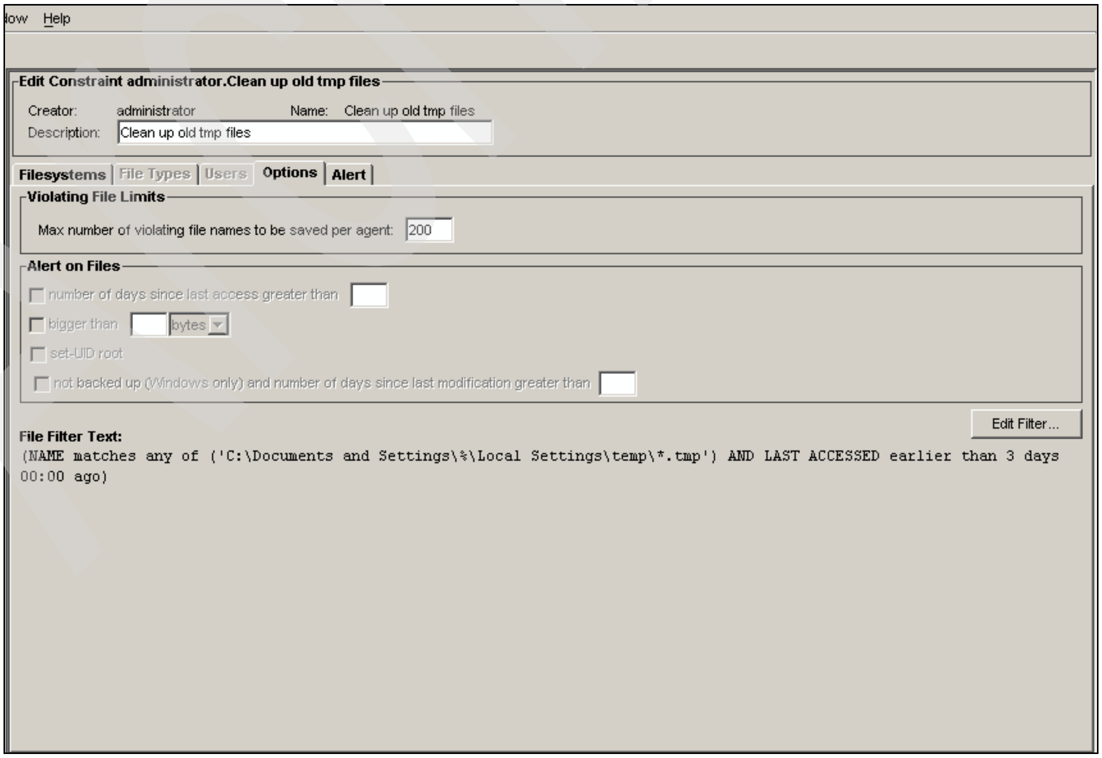

Edit Constraint administrator.Clean up old tmp files

Creator: administrator Name: Clean up old tmp files

Description: Clean up old tmp files

Violating File Limits

Max number of violating file names to be saved per agent: 200

Alert on Files

☐ number of days since last access greater than

☐ bigger than bytes

☐ set-UID root

☐ not backed up (Windows only) and number of days since last modification greater than

File Filter Text:

```
(NAME matches any of ('C:\Documents and Settings\%\Local Settings\temp\*.tmp') AND LAST ACCESSED earlier than 3 days
00:00 ago)
```

Figure 2-29 Create a file filter

4. This example filter has two conditions:
   - File names are *.tmp and exist in C:\Documents and Settings\<*any dir*>\Local Settings\temp.
   - File has not been accessed for at least three days.

> **Note:** Using the percent symbol (%) wildcard in a path statement means ANY. Therefore, the example in Figure 2-30 will search all of the Local Settings\temp in all users profile directories.

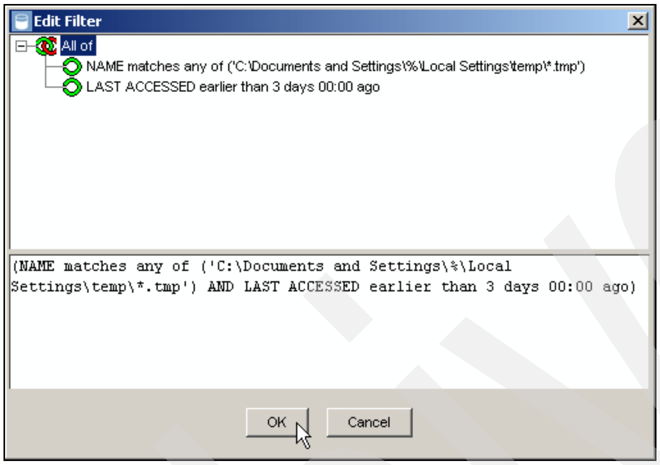

All of
- NAME matches any of ('C:\Documents and Settings\%\Local Settings\temp*.tmp')
- LAST ACCESSED earlier than 3 days 00:00 ago

```
(NAME matches any of ('C:\Documents and Settings\%\Local
Settings\temp\*.tmp') AND LAST ACCESSED earlier than 3 days 00:00 ago)
```

*Figure 2-30 Edit Filter*

5. Next, move to the Alert panel by selecting the **Alert** tab as in Figure 2-31 on page 37. Set the **Condition:** to **Violation Files Consume More Than 1 Kilobytes**. Then, check **Run Script** and click **Define**. If you want TPC to e-mail an administrator when the constraint is triggered, check **Email** and enter a list of the addresses to which to send alerts.

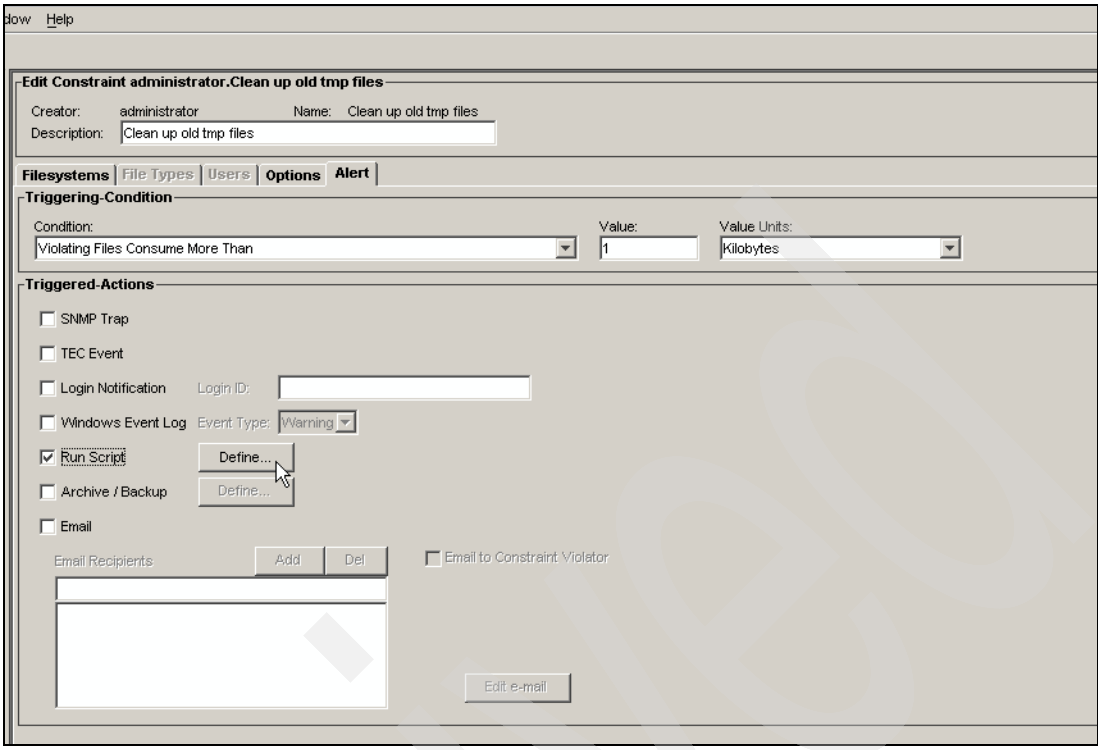

Edit Constraint administrator.Clean up old tmp files

Creator: administrator Name: Clean up old tmp files

Description: Clean up old tmp files

Triggering-Condition

Condition: Violating Files Consume More Than

Value: 1

Value Units: Kilobytes

Triggered-Actions

☐ SNMP Trap

☐ TEC Event

☐ Login Notification Login ID:

☐ Windows Event Log Event Type: Warning

☑ Run Script

☐ Archive / Backup

☐ Email

Email Recipients

☐ Email to Constraint Violator

Figure 2-31 Set alert preferences

6. In the Specify Script dialog box in Figure 2-32, enter the name of the script, which is `delete_files.bat` in this example. Then, ensure that **Where to Run** is set to **(triggering computer)**.

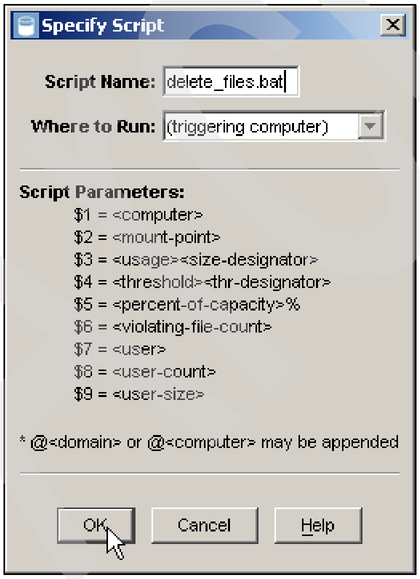

Script Name: delete_files.bat

Where to Run: (triggering computer)

Script Parameters:
$1 = <computer>
$2 = <mount-point>
$3 = <usage><size-designator>
$4 = <threshold><thr-designator>
$5 = <percent-of-capacity>%
$6 = <violating-file-count>
$7 = <user>
$8 = <user-count>
$9 = <user-size>

* @<domain> or @<computer> may be appended

Figure 2-32 Define script name

7. When you have defined the script name, click **OK**. Then click Ctrl+S to save the constraint and give it a unique name as in Figure 2-33 on page 38. Click **OK**.

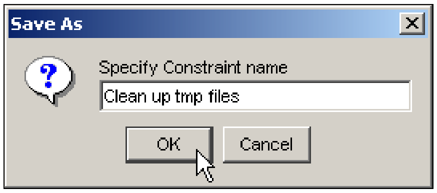

Specify Constraint name: Clean up tmp files

*Figure 2-33 Specify Constraint name and save constraint*

## Running the constraint

TPC processes the newly created constraint at the next scheduled scan of the machines to which the constraint applies. When the scan job finishes, each agent sends a new package of analysis data to the TPC server. The TPC server analyzes this data and if it identifies any files that meet the filter conditions, the TPC server calls the delete_files.bat script on the triggering machine and passes a list of files to the triggering machine for processing.

If you want to perform an immediate scan to test the constraint, you can instruct a scan job to run now as in Figure 2-34.

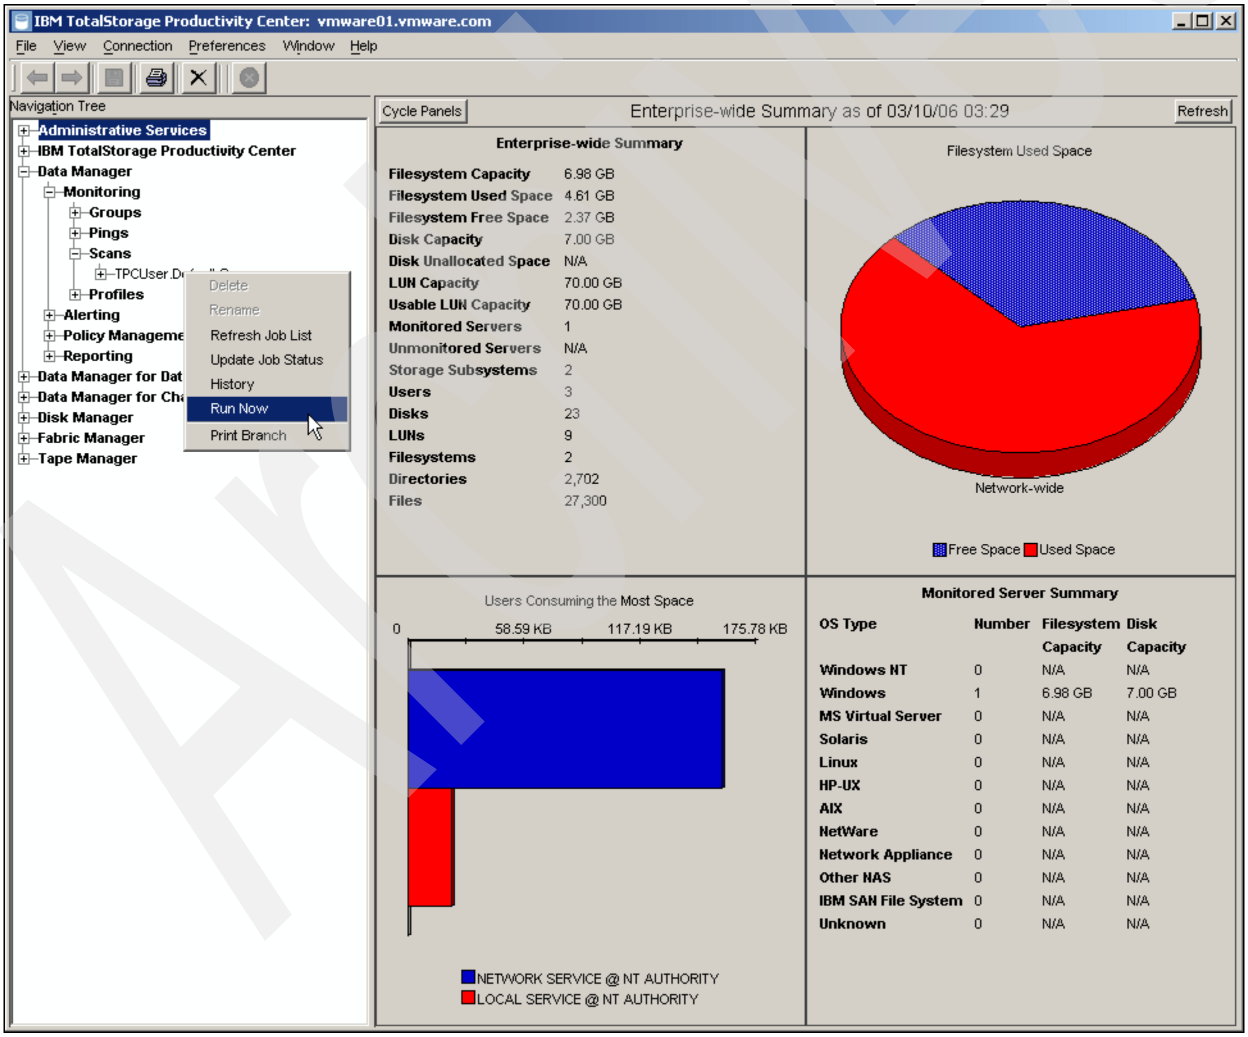

Enterprise-wide Summary as of 03/10/06 03:29

Enterprise-wide Summary

| | |
|---|---|
| Filesystem Capacity | 6.98 GB |
| Filesystem Used Space | 4.61 GB |
| Filesystem Free Space | 2.37 GB |
| Disk Capacity | 7.00 GB |
| Disk Unallocated Space | N/A |
| LUN Capacity | 70.00 GB |
| Usable LUN Capacity | 70.00 GB |
| Monitored Servers | 1 |
| Unmonitored Servers | N/A |
| Storage Subsystems | 2 |
| Users | 3 |
| Disks | 23 |
| LUNs | 9 |
| Filesystems | 2 |
| Directories | 2,702 |
| Files | 27,300 |

Monitored Server Summary

| OS Type | Number | Filesystem Capacity | Disk Capacity |
|---|---|---|---|
| Windows NT | 0 | N/A | N/A |
| Windows | 1 | 6.98 GB | 7.00 GB |
| MS Virtual Server | 0 | N/A | N/A |
| Solaris | 0 | N/A | N/A |
| Linux | 0 | N/A | N/A |
| HP-UX | 0 | N/A | N/A |
| AIX | 0 | N/A | N/A |
| NetWare | 0 | N/A | N/A |
| Network Appliance | 0 | N/A | N/A |
| Other NAS | 0 | N/A | N/A |
| IBM SAN File System | 0 | N/A | N/A |
| Unknown | 0 | N/A | N/A |

*Figure 2-34 Run scan now*

## Viewing the results of the constraint

When this example constraint finds *.tmp files that meet the conditions for deleting, TPC runs the specified script and creates an alert in the Alert Log. You can view the results of the actions taken in Figure 2-35.

Click the magnifying glass to the left of the alert to drill down into the detailed log information.

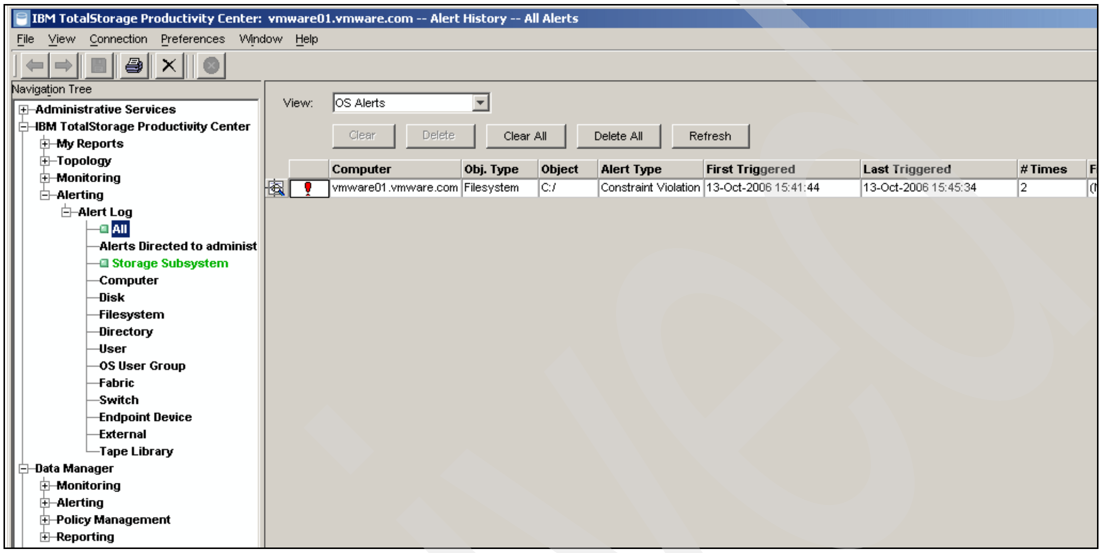

View: OS Alerts

| Computer | Obj. Type | Object | Alert Type | First Triggered | Last Triggered | # Times | F |
|---|---|---|---|---|---|---|---|
| vmware01.vmware.com | Filesystem | C:/ | Constraint Violation | 13-Oct-2006 15:41:44 | 13-Oct-2006 15:45:34 | 2 | (I |

Figure 2-35 Viewing the Alert Log

Figure 2-36 shows the detailed log information for the constraint. To view the log that was generated by the script, click **View Script Log**.

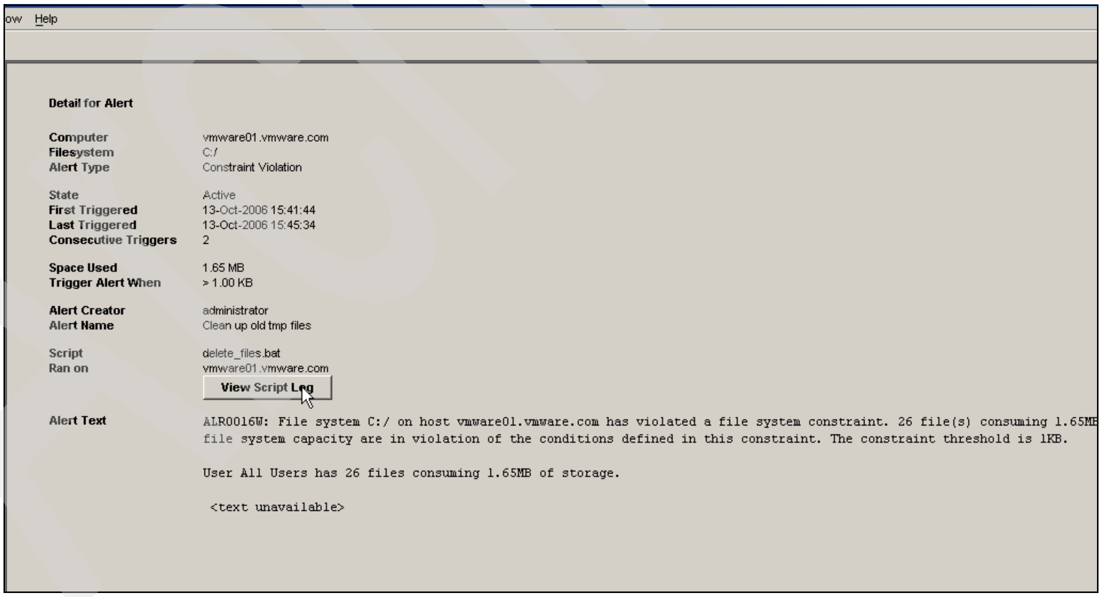

**Detail for Alert**

| | |
|---|---|
| Computer | vmware01.vmware.com |
| Filesystem | C:/ |
| Alert Type | Constraint Violation |
| State | Active |
| First Triggered | 13-Oct-2006 15:41:44 |
| Last Triggered | 13-Oct-2006 15:45:34 |
| Consecutive Triggers | 2 |
| Space Used | 1.65 MB |
| Trigger Alert When | > 1.00 KB |
| Alert Creator | administrator |
| Alert Name | Clean up old tmp files |
| Script | delete_files.bat |
| Ran on | vmware01.vmware.com |

Alert Text

```
ALR0016W: File system C:/ on host vmware01.vmware.com has violated a file system constraint. 26 file(s) consuming 1.65MB
file system capacity are in violation of the conditions defined in this constraint. The constraint threshold is 1KB.

User All Users has 26 files consuming 1.65MB of storage.

 <text unavailable>
```

Figure 2-36 Constraint information

The script log looks similar to Figure 2-37 on page 40.

> **Note:** By default, the sample script lists the file names that were passed to it by TPC, but it does not delete them. This allows you to test the process to ensure that you are satisfied with the results before the sample script deletes the files or does anything destructive.
>
> When you are satisfied with the results, then remove the comment from the last line of the script. This last line contains the delete routine.

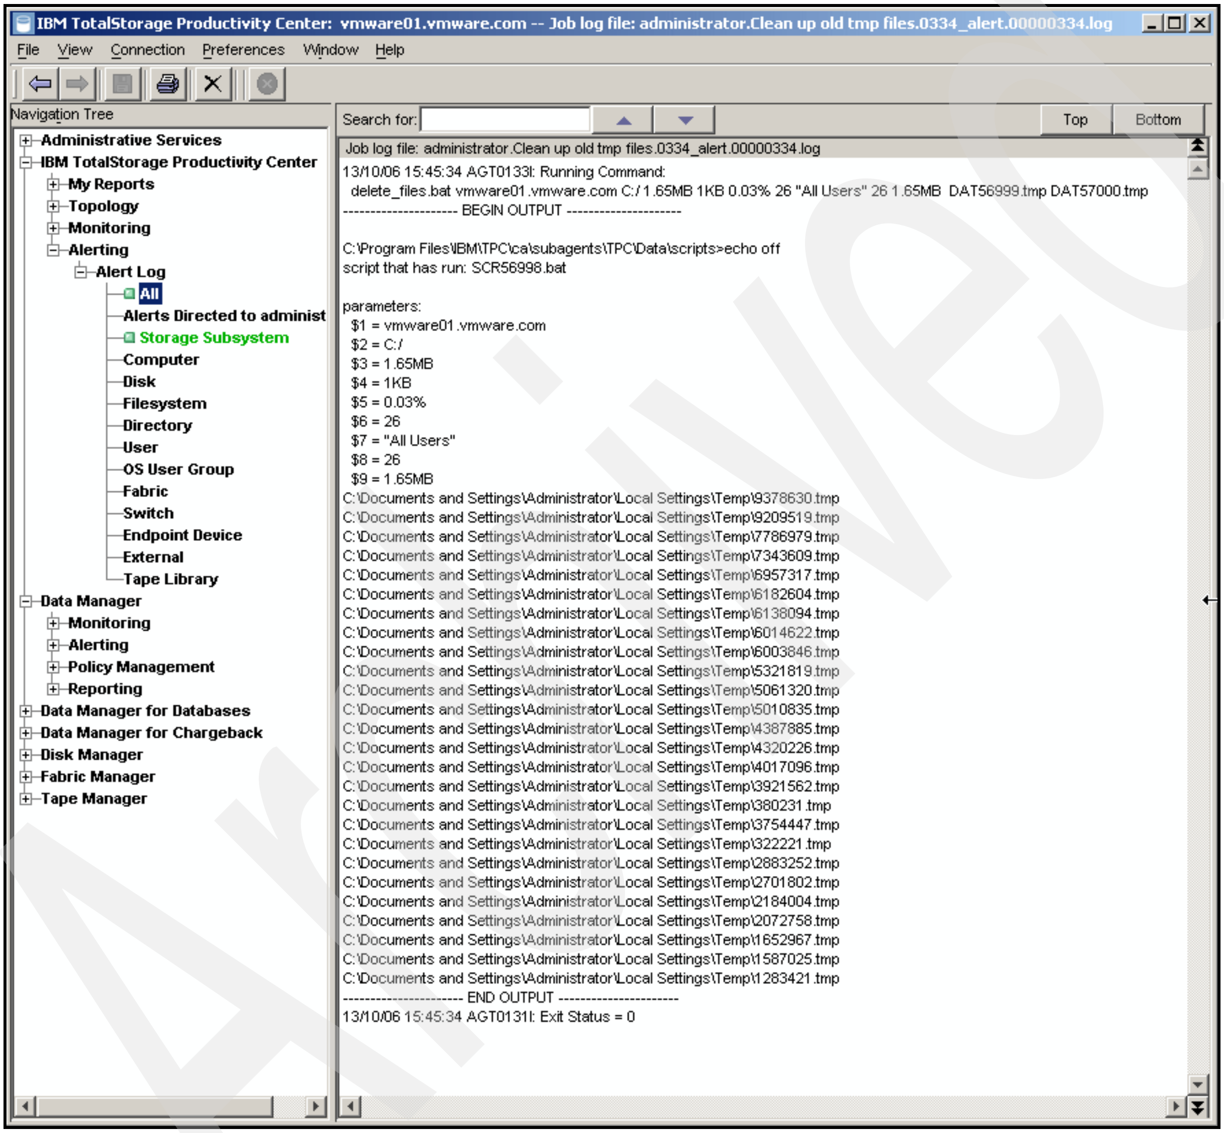

```
Job log file: administrator.Clean up old tmp files.0334_alert.00000334.log
13/10/06 15:45:34 AGT0133I: Running Command:
  delete_files.bat vmware01.vmware.com C:/ 1.65MB 1KB 0.03% 26 "All Users" 26 1.65MB  DAT56999.tmp DAT57000.tmp
--------------------- BEGIN OUTPUT ---------------------

C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\scripts>echo off
script that has run: SCR56998.bat

parameters:
  $1 = vmware01.vmware.com
  $2 = C:/
  $3 = 1.65MB
  $4 = 1KB
  $5 = 0.03%
  $6 = 26
  $7 = "All Users"
  $8 = 26
  $9 = 1.65MB
C:\Documents and Settings\Administrator\Local Settings\Temp\9378630.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\9209519.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\7786979.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\7343609.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\6957317.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\6182604.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\6138094.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\6014622.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\6003846.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\5321819.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\5061320.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\5010835.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\4387885.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\4320226.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\4017096.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\3921562.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\380231.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\3754447.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\322221.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\2883252.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\2701802.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\2184004.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\2072758.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\1652967.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\1587025.tmp
C:\Documents and Settings\Administrator\Local Settings\Temp\1283421.tmp
--------------------- END OUTPUT ---------------------
13/10/06 15:45:34 AGT0131I: Exit Status = 0
```

*Figure 2-37 Viewing the script log*

## 2.4 Spotting unusual data growth incidents

TPC can help you to be proactive in spotting storage-related events that can lead to unexpected storage shortages. Learning from previous events can reduce support diagnosis time for recurring events.

## 2.4.1 Sample data growth scenario

### The incident

A user calls the service desk because the user has trouble logging on to the Windows desktop this morning. The process appears to be taking much longer than usual. This Windows user has a roaming profile that is stored on a file server at a remote location over a wide area link.

Level one support find no problems with the profile server and other users have logged on successfully from the same location. They pass the call to level two support.

Level two support spends some time looking at the issue and then discovers that yesterday this user installed some CD ripping software on their workstation and extracted about 4GB of music. The user saved this music in a folder on the user's desktop. Because this is a roaming profile, all of the data was stored remotely on the profile server. When the user attempted to log in the next day, the user experienced a long delay because the very large profile had to be downloaded over a slow wide area network.

User A was unaware that this type of action can cause a problem.

### The outcome

In this case, level two support told the user not to store personal MP3 files on company equipment and level two support deleted the data.

After the event, the administrators looked at TPC and saw a sudden increase in MP3 data on the previous day.

### Lesson learned

Being able to see what happened after the event adds little value in terms of stopping another user from making a similar mistake a week later. Moreover, the service desk personnel must go through the same support cycle to diagnose and correct the problem again.

Using TPC, this organization expressed in a policy what happened that day, so that TPC in the future spots the sudden growth in a user's profile. TPC can then alert support staff that there is a potential problem and take proactive rather the reactive steps to correct the problem.

Having an alert also shortens the time that it takes to diagnose a problem by giving support staff a heads-up on an issue.

Over time, as these storage related events occur, you can express each event as a policy in TPC. Therefore, TPC becomes a knowledge base of known storage issues and can help storage support become far more proactive in handling problems.

## 2.4.2 Spot data growth policy setup

Using the example scenario in 2.4.1, "Sample data growth scenario" on page 41, this section details how to express this situation in TPC so that TPC can generate an alert if this event recurs.

There are several ways to handle this situation:

- Set a trigger if a user's profile directory reaches a defined size limit (directory alert).
- Look for an increase in a single day in newly created files of any type in a profile directory.
- Look for an increase in a single day of new MP3 files in a profile directory.

This example describes how to look for growth of more than 4 MB of MP3 files in a single day in the user profile directories. The steps are:

1. Start by creating a new constraint as in Figure 2-38.

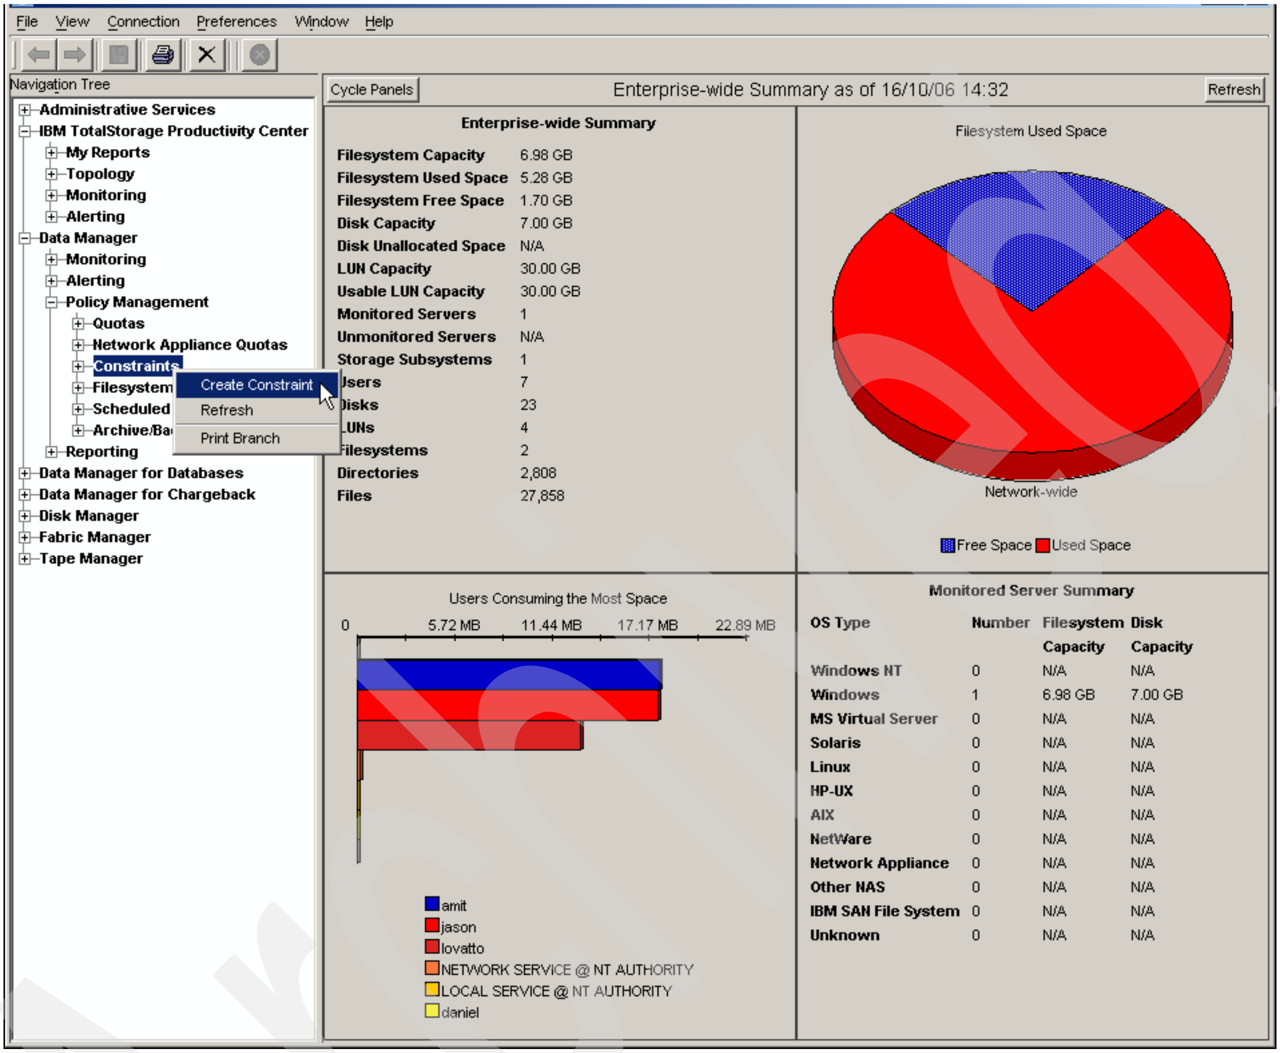

Enterprise-wide Summary as of 16/10/06 14:32

Enterprise-wide Summary

| | |
|---|---|
| Filesystem Capacity | 6.98 GB |
| Filesystem Used Space | 5.28 GB |
| Filesystem Free Space | 1.70 GB |
| Disk Capacity | 7.00 GB |
| Disk Unallocated Space | N/A |
| LUN Capacity | 30.00 GB |
| Usable LUN Capacity | 30.00 GB |
| Monitored Servers | 1 |
| Unmonitored Servers | N/A |
| Storage Subsystems | 1 |
| Users | 7 |
| Disks | 23 |
| LUNs | 4 |
| Filesystems | 2 |
| Directories | 2,808 |
| Files | 27,858 |

Monitored Server Summary

| OS Type | Number | Filesystem Capacity | Disk Capacity |
|---|---|---|---|
| Windows NT | 0 | N/A | N/A |
| Windows | 1 | 6.98 GB | 7.00 GB |
| MS Virtual Server | 0 | N/A | N/A |
| Solaris | 0 | N/A | N/A |
| Linux | 0 | N/A | N/A |
| HP-UX | 0 | N/A | N/A |
| AIX | 0 | N/A | N/A |
| NetWare | 0 | N/A | N/A |
| Network Appliance | 0 | N/A | N/A |
| Other NAS | 0 | N/A | N/A |
| IBM SAN File System | 0 | N/A | N/A |
| Unknown | 0 | N/A | N/A |

*Figure 2-38 Create a constraint*

2. On the first constraint panel, select the filesystems that contain the profile directories of the users to monitor. Then, give the constraint a meaningful name as in Figure 2-39 on page 43.

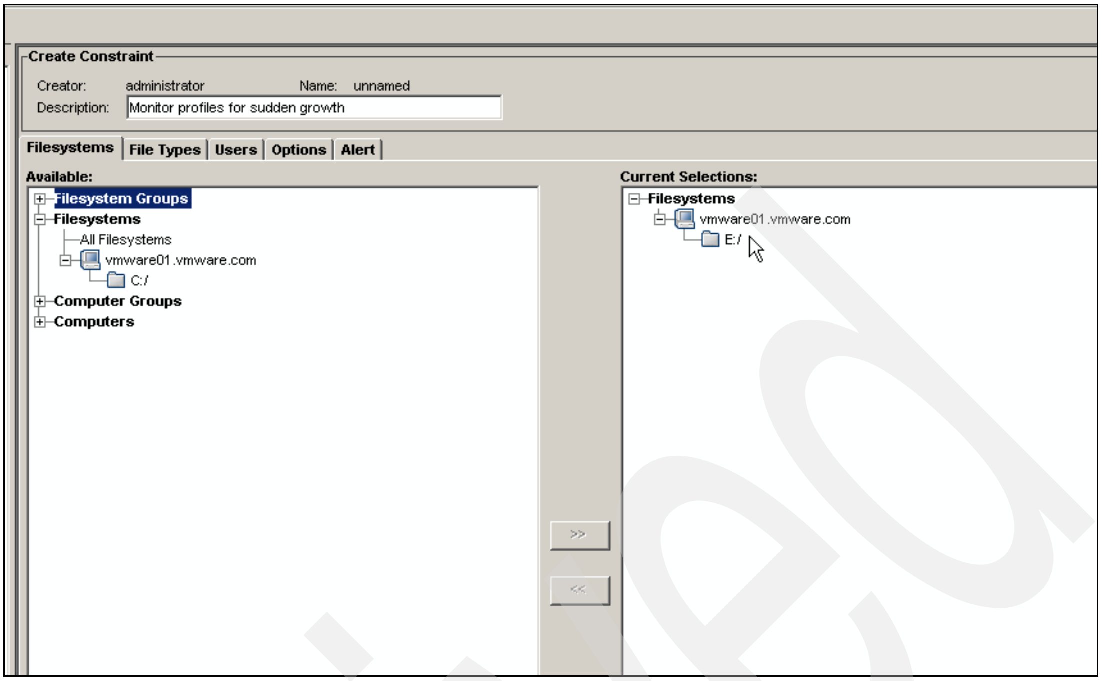

Create Constraint

Creator: administrator Name: unnamed

Description: Monitor profiles for sudden growth

Figure 2-39 Select filesystems for this constraint

3. Click the **Options** tab and click **Edit Filter** as in Figure 2-40.

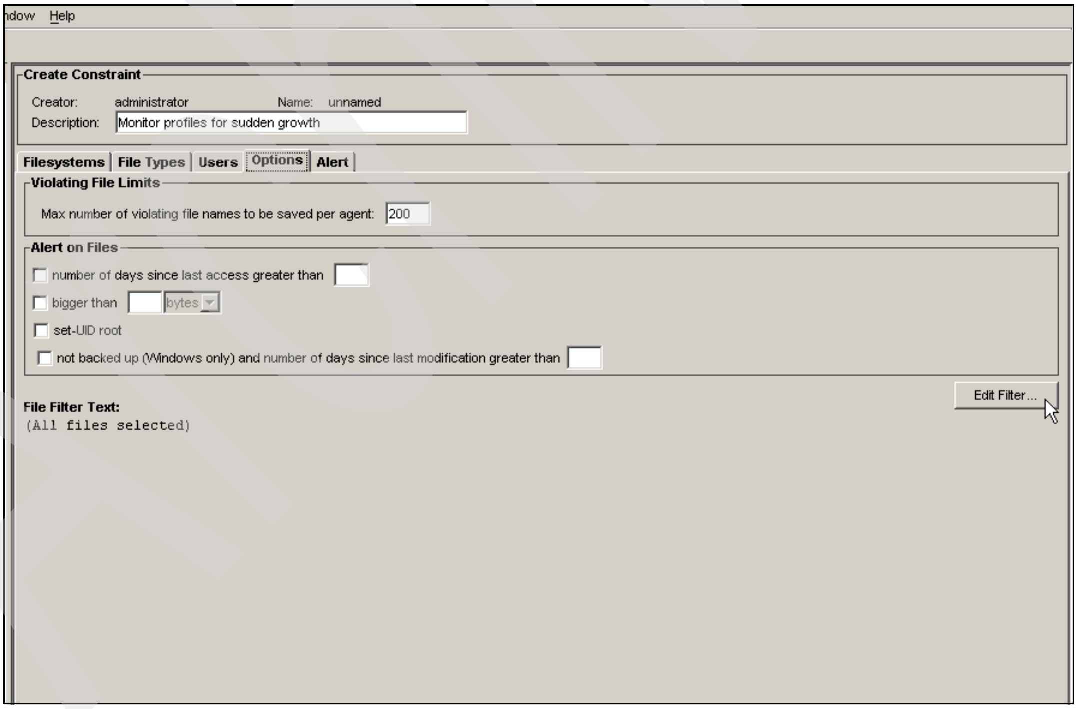

Create Constraint

Creator: administrator Name: unnamed

Description: Monitor profiles for sudden growth

Violating File Limits

Max number of violating file names to be saved per agent: 200

Alert on Files

□ number of days since last access greater than

□ bigger than bytes

□ set-UID root

□ not backed up (Windows only) and number of days since last modification greater than

File Filter Text:
(All files selected)

Figure 2-40 Edit Filter

4. Figure 2-41 on page 44 shows a filter that has the following characteristics:

   – The file names end in mp3 and are in any directory below E:\Profiles, which is the top level directory for the user profiles on this system.

   – Matching files were created within the last day.

**Note:** If you only scan the filesystem in question one time each week, adjust the filter to look for mp3 files created in the last seven days rather than in the last day.

Build your filter in a similar manner to this filter to suit your environment. You might want to change the filter to report on all files that were created in the last seven days in profile directories and not just mp3 files. Adjust the file names and file types that match the filter criteria as required.

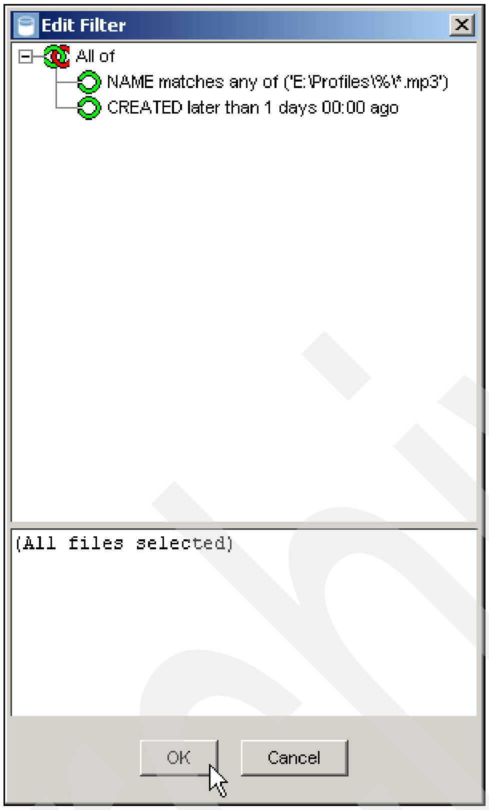

All of
NAME matches any of ('E:\Profiles\%\*.mp3')
CREATED later than 1 days 00:00 ago

```
(All files selected)
```

*Figure 2-41 Edit Filter*

5. Click the **Alert** tab seen in Figure 2-42 on page 45 and add an **Email** recipient to receive the alerts.

   Then, set a value for the trigger condition. In this example, we set the condition to trigger if there are more than 4 MB of newly created files found each day. Adjust this value as appropriate for your environment.

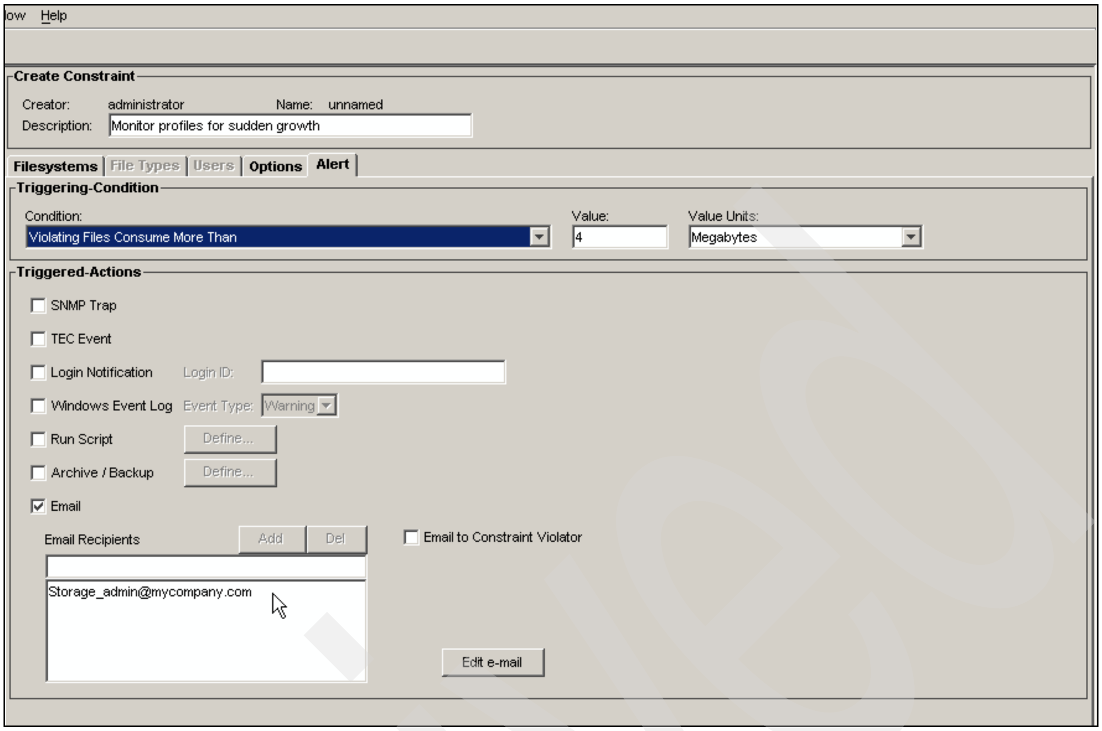

Create Constraint

Creator: administrator Name: unnamed

Description: Monitor profiles for sudden growth

Triggering-Condition

Condition: Violating Files Consume More Than

Value: 4

Value Units: Megabytes

Triggered-Actions

- ☐ SNMP Trap
- ☐ TEC Event
- ☐ Login Notification Login ID:
- ☐ Windows Event Log Event Type: Warning
- ☐ Run Script
- ☐ Archive / Backup
- ☑ Email

Email Recipients

Storage_admin@mycompany.com

☐ Email to Constraint Violator

*Figure 2-42 Define constraint trigger and alerting*

6. Save the new constraint by clicking Ctrl+S and give the new constraint a meaningful name as in Figure 2-43. Click **OK**.

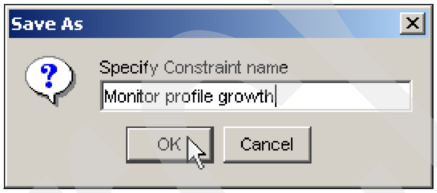

Specify Constraint name

Monitor profile growth

*Figure 2-43 Save Constraint name*

## 2.4.3 Activating the constraint

TPC acts on the newly created data constraint at the next scan of the appropriate filesystems. If you want to test the constraint operation manually, run a scan job against the computer and the filesystem.

When the scan job completes and sends the results to the TPC server, the TPC server triggers an alert if necessary.

## 2.4.4 Constraint reporting

This section describes how to report on the constraint that we set up in the previous section for spotting sudden data growth in the profile directories of the users.

## E-mail alerts

If you configured the constraint to send e-mail alerts to an administrator, this is a sample of the output you can expect. If there are multiple users who violate the constraint on the same day, the TPC server generates multiple e-mail alerts (one for each user).

*Example 2-1 Sample constraint e-mail alert*

```
Alert administrator.Monitor profile growth has been triggered.

File system E:/ on host vmware01.vmware.com has violated a file system constraint. 11 file(s) consuming
48.24MB or 2.37% of the file system capacity are in violation of the conditions defined in this con-
straint. The constraint threshold is 4MB.

User jason has 4 files consuming 17.54MB of storage.

Your largest violating files are:

E:\Profiles\Jason\My Music\track4.mp3
E:\Profiles\Jason\My Music\track3.mp3
E:\Profiles\Jason\My Music\track2.mp3
E:\Profiles\Jason\My Music\Track1.mp3
```

## Constraint reports

To see the constraint report through the TPC GUI, follow these steps:

1. Choose the constraint report type and click **Generate Report** as in Figure 2-44.

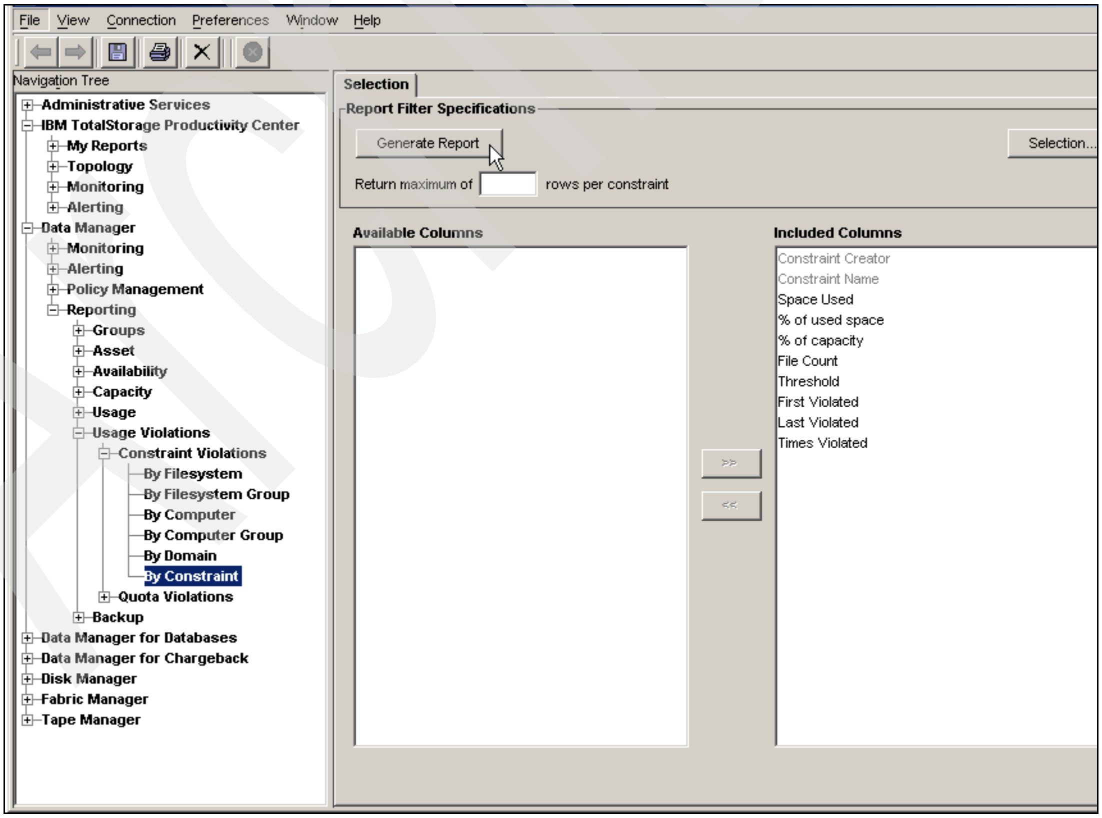

*Figure 2-44 Choose constraint report*

2. This produces a report for all constraint types unless you choose to filter it. Figure 2-45 shows that the Monitor profile growth constraint has been triggered and that there are 48.24 MB of files that are MP3 files that were created in the last day. Right-click the constraint and you can choose either **Show Violating Files** or **Show Violating Owners**.

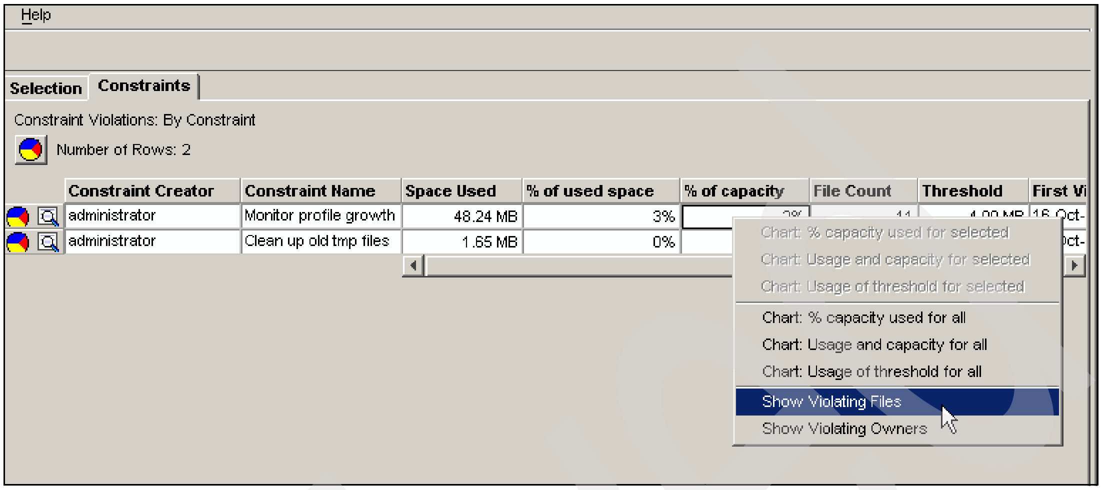
Help
Selection
Constraints
Chart: % capacity used for selected
Chart: Usage and capacity for selected
Chart: Usage of threshold for selected
Chart: % capacity used for all
Chart: Usage and capacity for all
Chart: Usage of threshold for all
Show Violating Files
Show Violating Owners

Constraint Violations: By Constraint

Number of Rows: 2

| Constraint Creator | Constraint Name | Space Used | % of used space | % of capacity | File Count | Threshold | First Vi |
|---|---|---|---|---|---|---|---|
| administrator | Monitor profile growth | 48.24 MB | 3% | [illegible] | 11 | 4.00 MB | 16 Oct- |
| administrator | Clean up old tmp files | 1.65 MB | 0% | | | | ct- |

*Figure 2-45 Constraint reports*

Figure 2-46 show the files that violate the constraint definition.

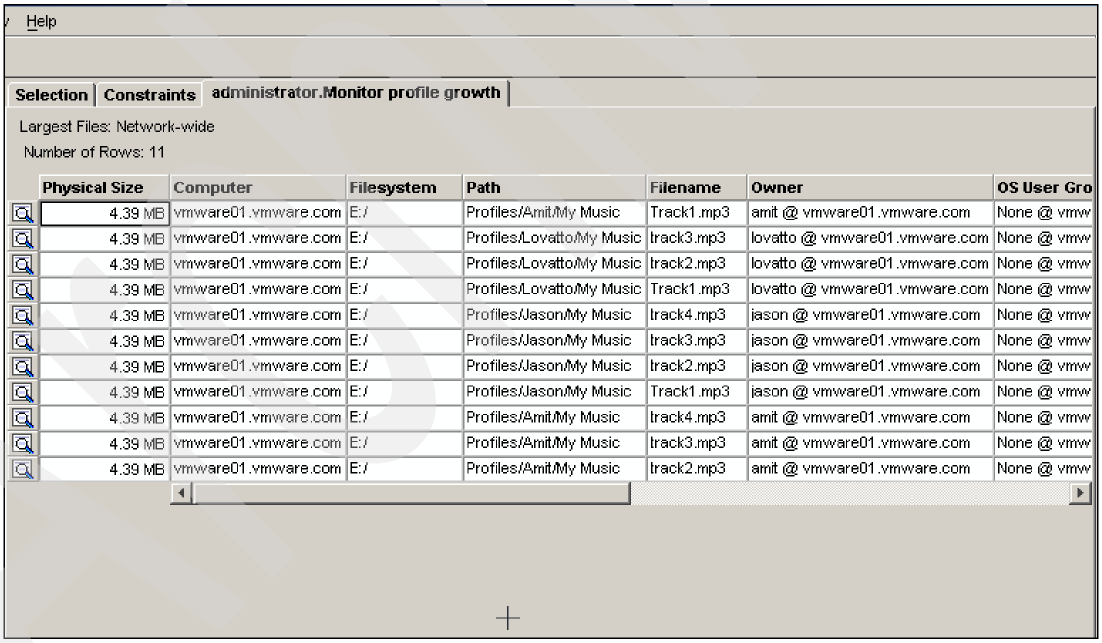

Largest Files: Network-wide

Number of Rows: 11

| Physical Size | Computer | Filesystem | Path | Filename | Owner | OS User Gro |
|---|---|---|---|---|---|---|
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Amit/My Music | Track1.mp3 | amit @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Lovatto/My Music | track3.mp3 | lovatto @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Lovatto/My Music | track2.mp3 | lovatto @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Lovatto/My Music | Track1.mp3 | lovatto @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Jason/My Music | track4.mp3 | jason @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Jason/My Music | track3.mp3 | jason @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Jason/My Music | track2.mp3 | jason @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Jason/My Music | Track1.mp3 | jason @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Amit/My Music | track4.mp3 | amit @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Amit/My Music | track3.mp3 | amit @ vmware01.vmware.com | None @ vmw |
| 4.39 MB | vmware01.vmware.com | E:/ | Profiles/Amit/My Music | track2.mp3 | amit @ vmware01.vmware.com | None @ vmw |

*Figure 2-46 Show violating files*

Figure 2-47 on page 48 shows the violating users.

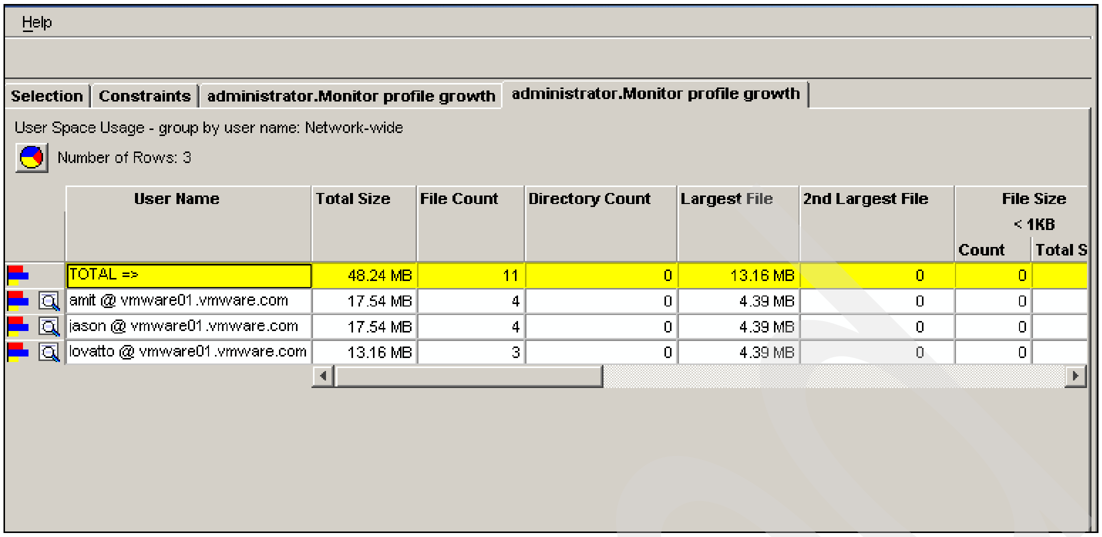
Help
Selection
Constraints
administrator.Monitor profile growth
administrator.Monitor profile growth

User Space Usage - group by user name: Network-wide

Number of Rows: 3

| User Name | Total Size | File Count | Directory Count | Largest File | 2nd Largest File | File Size <1KB | |
|---|---|---|---|---|---|---|---|
| | | | | | | Count | Total S |
| TOTAL => | 48.24 MB | 11 | 0 | 13.16 MB | 0 | 0 | |
| amit @ vmware01.vmware.com | 17.54 MB | 4 | 0 | 4.39 MB | 0 | 0 | |
| iason @ vmware01.vmware.com | 17.54 MB | 4 | 0 | 4.39 MB | 0 | 0 | |
| lovatto @ vmware01.vmware.com | 13.16 MB | 3 | 0 | 4.39 MB | 0 | 0 | |

*Figure 2-47 Violating users*

## 2.5 Tivoli Storage Manager automatic archiving and integration

This section looks at how TPC in combination with Tivoli Storage Manager provides a policy driven file archive function.

Tivoli Storage Manager already contains functions and policies for the controlled retention of data for a defined or indefinite period. TPC can act as an intelligent front-end processor for this capability by using powerful TPC policies to accurately define what gets archived and when it gets archived.

Do not confuse this function with space management tools that leave behind a file stub. The purpose of archiving in this section is to identify data that has business value but no longer justifies the cost of being kept on a server's disk storage.

Here are examples of situations where you might implement file archive:

- Example one

  Organizations that provide their users with a networked home share drive find that users use and abuse this space frequently. Laptop users might treat this area as a backup for data that they keep on their laptop computers.

  Perhaps users create a .zip file of versions of data from their laptop on the network share drive, never use the compressed data, and over time, create more .zip files.

  These .zip files can be large but still contain some valid data for a user. One way to manage the amount of .zip files that are kept on disk storage is to have TPC monitor the home directories for the users and look specifically for .zip files that have not been accessed for a long time, perhaps nine to 12 months.

  When TPC finds .zip files that meet this threshold, it instructs Tivoli Storage Manager to archive and delete the .zip files, thus freeing the disk space.

  With proper configuration, TPC can then e-mail the owner of the files to tell them that their old files have been archived and that the owners can retrieve the old files if they ever need them by placing a server desk request.

  You can set up the Tivoli Storage Manager policy to have a fixed retention of two years, for example. After which, the files expire from Tivoli Storage Manager and are deleted.

- Example two

  Similar to example one, TPC might monitor folder areas that contain many document files for correspondence with clients and suppliers.

  These documents might have a short life in terms of their access but require long term retention as reference material to record transactions with clients and suppliers.

  From the time that a user creates a document, the user often frequently accesses and updates the document for one or two weeks before sending out the final version. After that, the document remains static with perhaps infrequent read access over the next two to three months. After that, the document is not accessed for extended periods, unless an issue arises.

  TPC can monitor these folder areas for doc files that have not been accessed for four months and then archive these files to Tivoli Storage Manager as a group perhaps once a month.

  Not only does this recover the storage space, but it is a method of preventing someone from altering the file at a later date.

## 2.5.1 Defining an archive policy

This section describes how to build a policy in TPC that identifies the data that you want to archive.

For this example, we look for .zip files in users' directories that have not been accessed for six months.

### Understanding the impact of the archive policy

Before implementing a policy that archives and deletes files, you need to understand the impact of the policy before you proceed. If the policy only recovers a few megabytes of storage, the policy is not worth the effort of implementing the process.

To understand how much storage .zip files consume in users home directories, create a profile to collect the specific information for reporting.

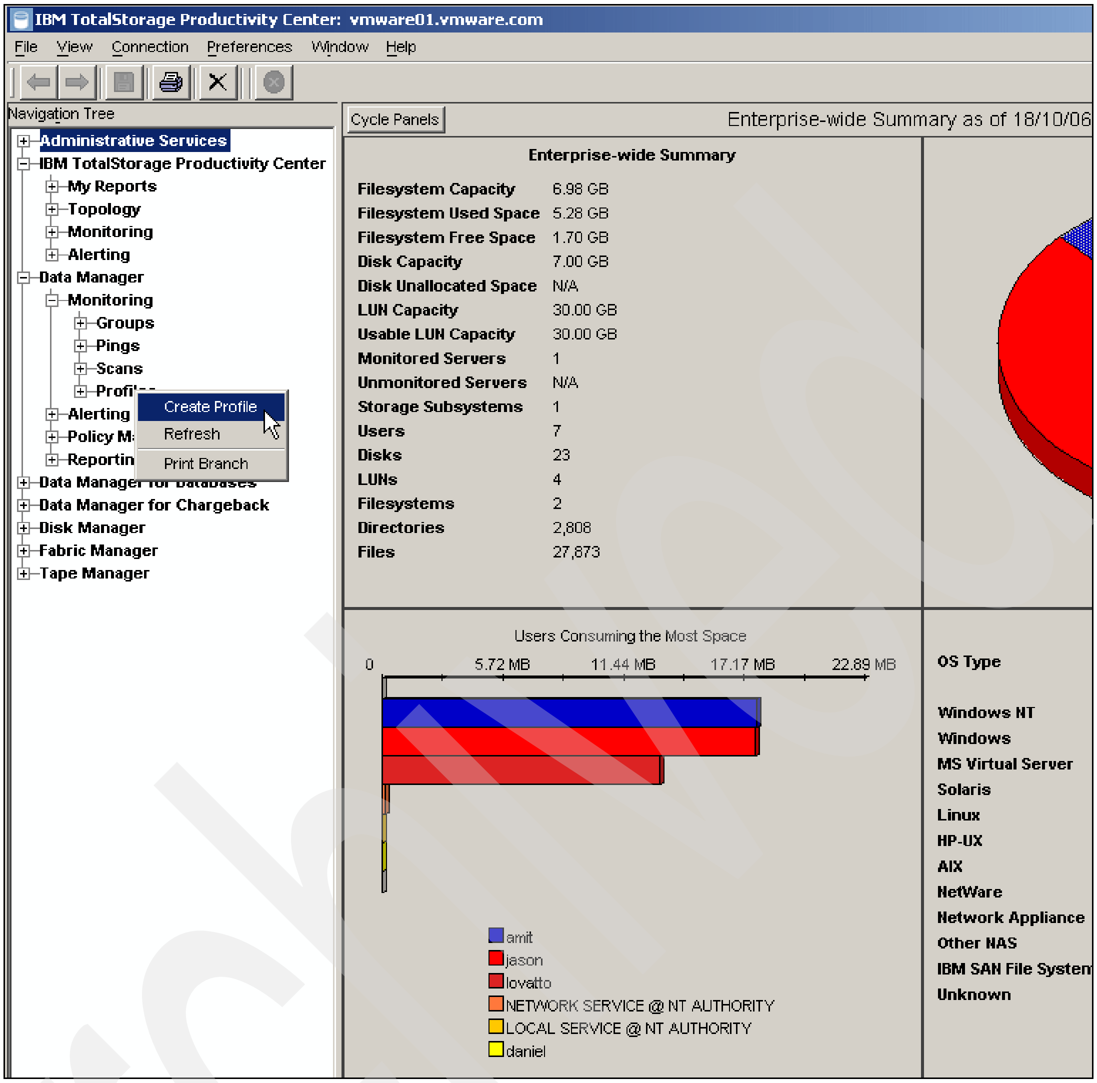

IBM TotalStorage Productivity Center: vmware01.vmware.com
File View Connection Preferences Window Help
Navigation Tree
Cycle Panels
Enterprise-wide Summary as of 18/10/06

- Administrative Services
- IBM TotalStorage Productivity Center
  - My Reports
  - Topology
  - Monitoring
  - Alerting
- Data Manager
  - Monitoring
    - Groups
    - Pings
    - Scans
    - Profi
  - Alerting
  - Policy M
  - Reportin
- Data Manager for Databases
- Data Manager for Chargeback
- Disk Manager
- Fabric Manager
- Tape Manager

Create Profile
Refresh
Print Branch

**Enterprise-wide Summary**

| | |
|---|---|
| Filesystem Capacity | 6.98 GB |
| Filesystem Used Space | 5.28 GB |
| Filesystem Free Space | 1.70 GB |
| Disk Capacity | 7.00 GB |
| Disk Unallocated Space | N/A |
| LUN Capacity | 30.00 GB |
| Usable LUN Capacity | 30.00 GB |
| Monitored Servers | 1 |
| Unmonitored Servers | N/A |
| Storage Subsystems | 1 |
| Users | 7 |
| Disks | 23 |
| LUNs | 4 |
| Filesystems | 2 |
| Directories | 2,808 |
| Files | 27,873 |

Users Consuming the Most Space
0
5.72 MB
11.44 MB
17.17 MB
22.89 MB
amit
jason
lovatto
NETWORK SERVICE @ NT AUTHORITY
LOCAL SERVICE @ NT AUTHORITY
daniel

**OS Type**

Windows NT
Windows
MS Virtual Server
Solaris
Linux
HP-UX
AIX
NetWare
Network Appliance
Other NAS
IBM SAN File Systen
Unknown

*Figure 2-48 Create profile*

This example collects history for the 20 largest and most obsolete .zip files.

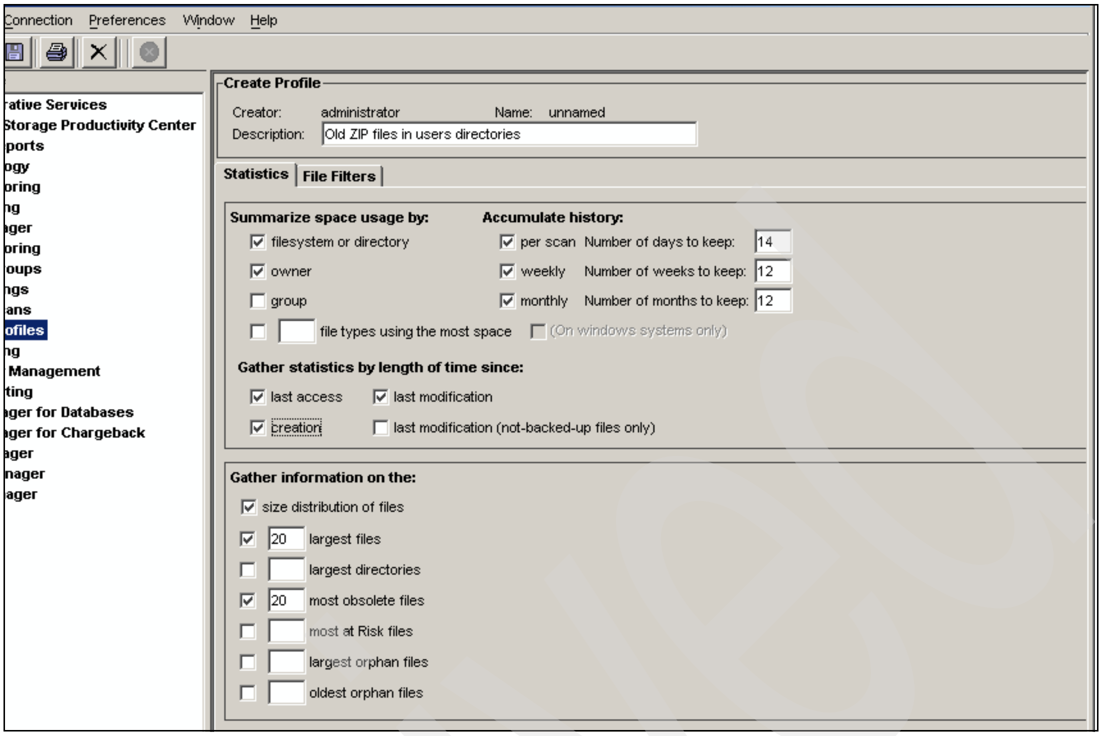

Create Profile

Creator: administrator Name: unnamed

Description: Old ZIP files in users directories

Statistics | File Filters

Summarize space usage by:
- ☑ filesystem or directory
- ☑ owner
- □ group
- □ file types using the most space

Accumulate history:
- ☑ per scan Number of days to keep: 14
- ☑ weekly Number of weeks to keep: 12
- ☑ monthly Number of months to keep: 12
- □ (On windows systems only)

Gather statistics by length of time since:
- ☑ last access
- ☑ last modification
- ☑ creation
- □ last modification (not-backed-up files only)

Gather information on the:
- ☑ size distribution of files
- ☑ 20 largest files
- □ largest directories
- ☑ 20 most obsolete files
- □ most at Risk files
- □ largest orphan files
- □ oldest orphan files

*Figure 2-49 Create profile*

On the File Filter panel, create a filter similar to the filter in Figure 2-50. The filter is targeted at .zip files in the users home directories that have not been accessed for 168 days.

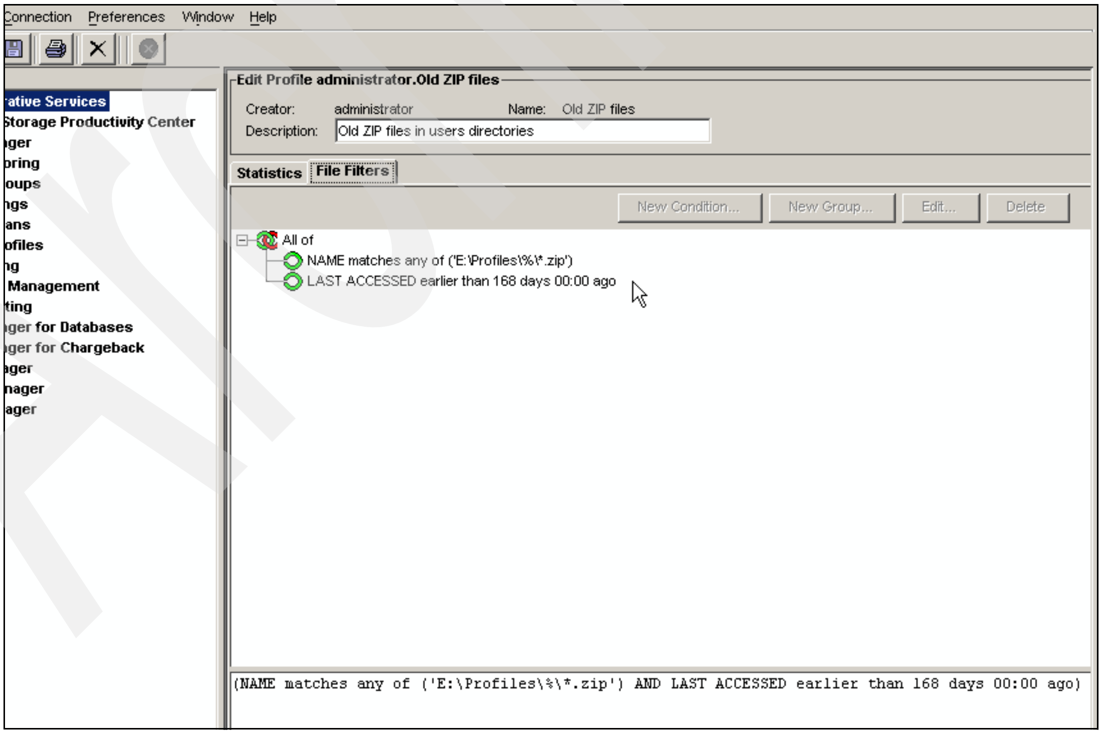

Edit Profile administrator.Old ZIP files

Creator: administrator Name: Old ZIP files

Description: Old ZIP files in users directories

Statistics | File Filters

All of
- NAME matches any of ('E:\Profiles\%\*.zip')
- LAST ACCESSED earlier than 168 days 00:00 ago

```
(NAME matches any of ('E:\Profiles\%\*.zip') AND LAST ACCESSED earlier than 168 days 00:00 ago)
```

*Figure 2-50 Create file filter*

Save the profile and give it a name as in Figure 2-51 on page 52. Click **OK**.

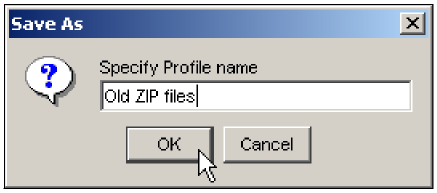

Specify Profile name: Old ZIP files

*Figure 2-51 Save profile*

Next, either add the profile to an existing scan job or create a new scan job. Figure 2-52 shows the Profiles page of the existing default scan. Select the profile from the left panel and add it to the **Profiles to apply to Filesystems** panel on the right. Save the profile by clicking Ctrl+S.

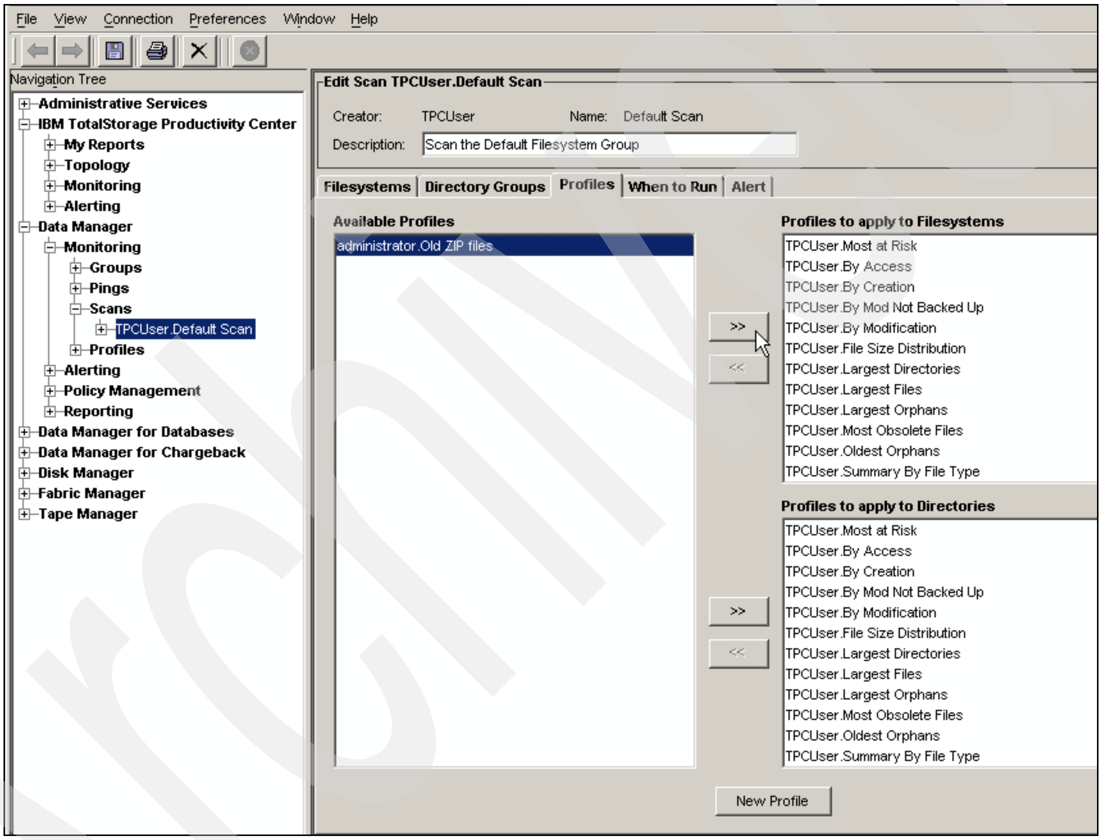

Edit Scan TPCUser.Default Scan

Creator: TPCUser Name: Default Scan

Description: Scan the Default Filesystem Group

Available Profiles

administrator.Old ZIP files

Profiles to apply to Filesystems

TPCUser.Most at Risk
TPCUser.By Access
TPCUser.By Creation
TPCUser.By Mod Not Backed Up
TPCUser.By Modification
TPCUser.File Size Distribution
TPCUser.Largest Directories
TPCUser.Largest Files
TPCUser.Largest Orphans
TPCUser.Most Obsolete Files
TPCUser.Oldest Orphans
TPCUser.Summary By File Type

Profiles to apply to Directories

TPCUser.Most at Risk
TPCUser.By Access
TPCUser.By Creation
TPCUser.By Mod Not Backed Up
TPCUser.By Modification
TPCUser.File Size Distribution
TPCUser.Largest Directories
TPCUser.Largest Files
TPCUser.Largest Orphans
TPCUser.Most Obsolete Files
TPCUser.Oldest Orphans
TPCUser.Summary By File Type

*Figure 2-52 Add profile to an existing scan job*

> **Tip:** If the type of data for which you scan only exists in a small number of file and print servers, consider creating a scan job that only runs against the specific file and print servers. That way, you only apply the profile against machines on which the profile will find this data.
>
> Keeping the number of profiles to a minimum within a scan job improves scan performance.

The next time that the scan job runs, the scan job collects specific information about old .zip files in the targeted directory.

### *Running a targeted report*

Reporting data is available after the scan job that was configured in “Understanding the impact of the archive policy” on page 49 has run against the appropriate servers.

Figure 2-53 shows an Access Time report. Notice that the profile has been set to **Old ZIP files**. This means that the statistic you are given only counts those files that are .zip files and older than 168 days.

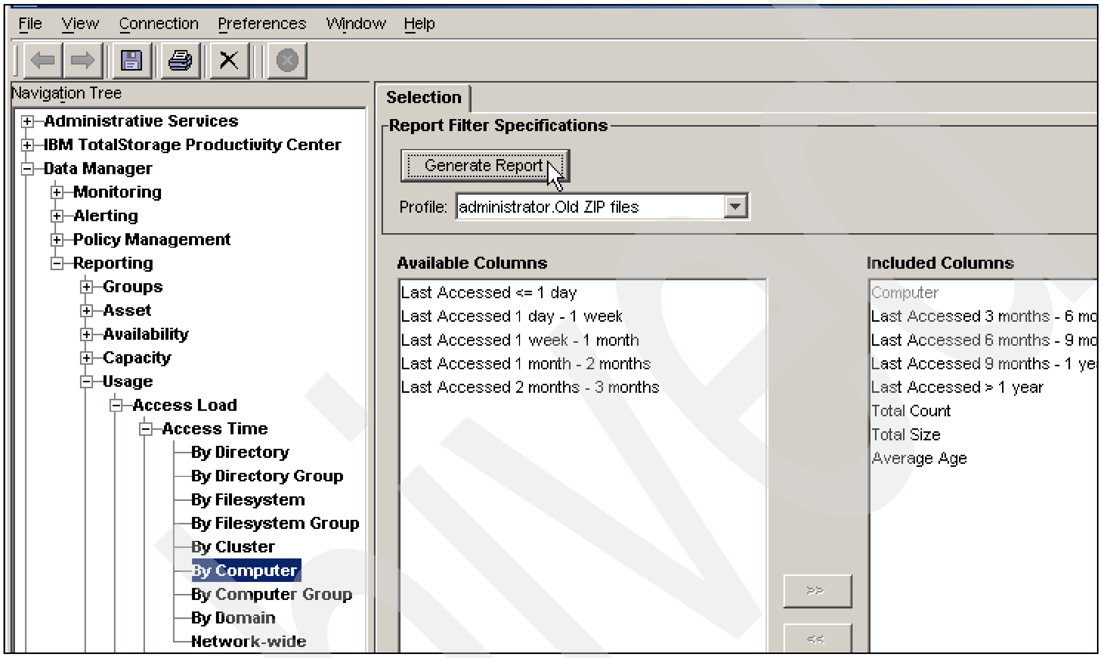

Report Filter Specifications

Profile: administrator.Old ZIP files

Available Columns

Last Accessed <= 1 day
Last Accessed 1 day - 1 week
Last Accessed 1 week - 1 month
Last Accessed 1 month - 2 months
Last Accessed 2 months - 3 months

Included Columns

Computer
Last Accessed 3 months - 6 mo
Last Accessed 6 months - 9 mo
Last Accessed 9 months - 1 ye
Last Accessed > 1 year
Total Count
Total Size
Average Age

*Figure 2-53 Running a report on old .zip files*

Figure 2-54 shows the results. The analysis of the data on our test machine shows 12 files totalling 71.96 MB. Therefore, if we implement a policy to archive these files, this is the maximum space that you can recover. Based on the results in your environment, you can make a decision about the viability of implementing an automatic archiving process for this set of data.

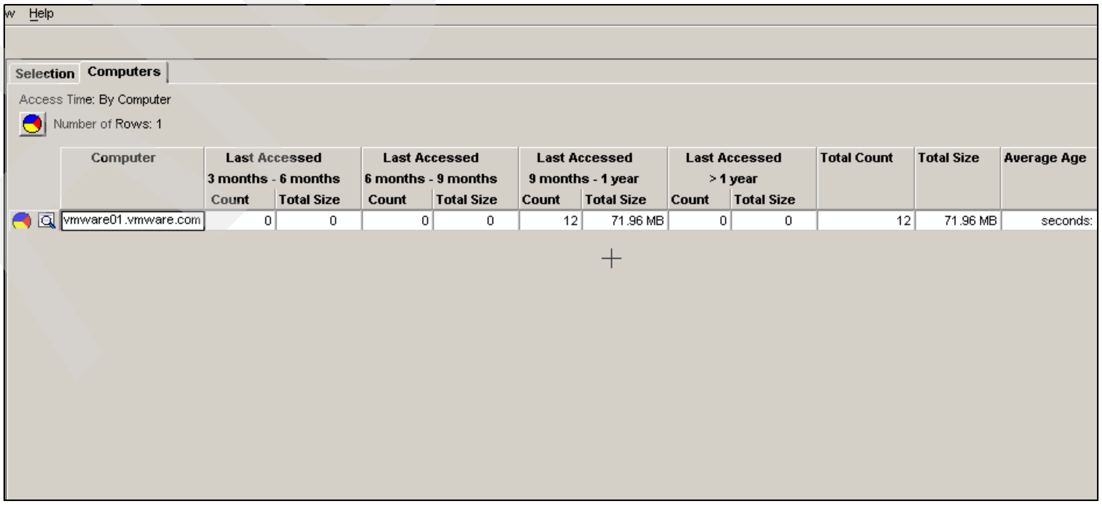

Access Time: By Computer

Number of Rows: 1

| Computer | Last Accessed<br>3 months - 6 months | | Last Accessed<br>6 months - 9 months | | Last Accessed<br>9 months - 1 year | | Last Accessed<br>> 1 year | | Total Count | Total Size | Average Age |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | Count | Total Size | Count | Total Size | Count | Total Size | Count | Total Size | | | |
| vmware01.vmware.com | 0 | 0 | 0 | 0 | 12 | 71.96 MB | 0 | 0 | 12 | 71.96 MB | seconds: |

*Figure 2-54 Access Time by Computer report*

## 2.5.2 Configuring automatic archiving to Tivoli Storage Manager

This section describes the steps that are required to configure TPC to archive files to Tivoli Storage Manager based on a policy.

We assume the following configuration:

- We have already installed and configured a Tivoli Storage Manager Backup/Archive client on the machines that contain the files to archive.
- The Tivoli Storage Manager client does not need to enter a password to send files. If a password is required in your environment, you need to enter an extra command line option into TPC.
- We have configured a suitable archive management class in Tivoli Storage Manager to receive the files.

Work with the Tivoli Storage Manager administrator in your environment to establish this configuration.

To create the constraint that selects old .zip files and archives them to Tivoli Storage Manager, use the following steps:

1. Right-click **Constraint** to create a new constraint as in Figure 2-55.

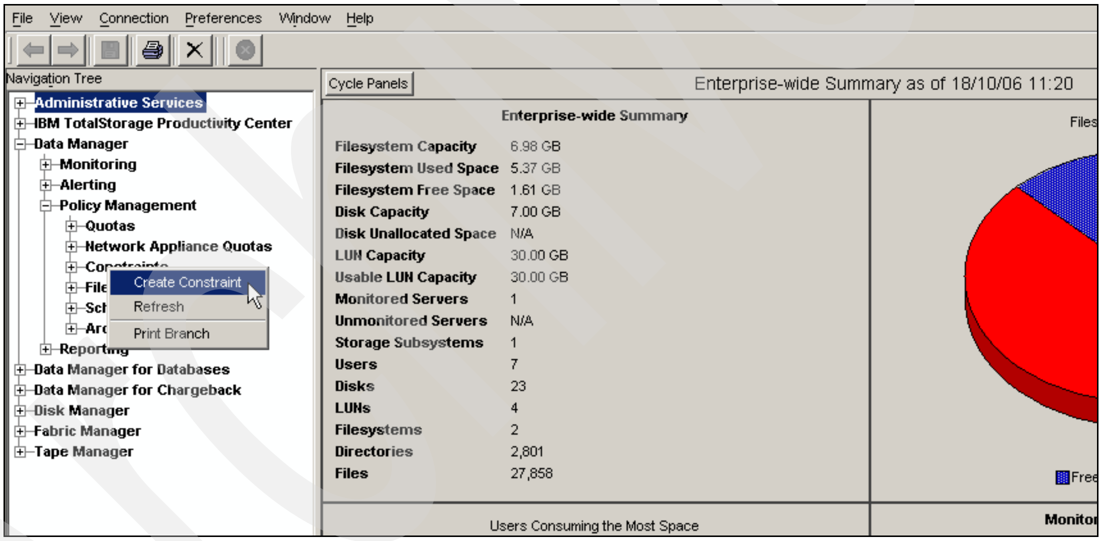

Enterprise-wide Summary as of 18/10/06 11:20

Enterprise-wide Summary

| | |
|---|---|
| Filesystem Capacity | 6.98 GB |
| Filesystem Used Space | 5.37 GB |
| Filesystem Free Space | 1.61 GB |
| Disk Capacity | 7.00 GB |
| Disk Unallocated Space | N/A |
| LUN Capacity | 30.00 GB |
| Usable LUN Capacity | 30.00 GB |
| Monitored Servers | 1 |
| Unmonitored Servers | N/A |
| Storage Subsystems | 1 |
| Users | 7 |
| Disks | 23 |
| LUNs | 4 |
| Filesystems | 2 |
| Directories | 2,801 |
| Files | 27,858 |

*Figure 2-55 Create a constraint*

2. Select the specific machines and filesystems to which this constraint applies. In Figure 2-56 on page 55, we selected only the E: drive, because this drive is where the user profile directories exist. You can select multiple machines and filesystems on this panel.

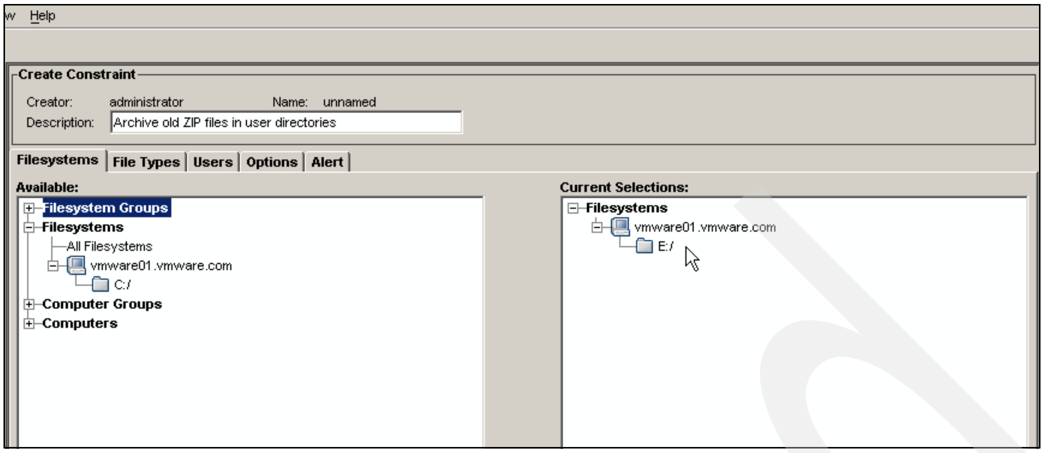

Create Constraint

Creator: administrator Name: unnamed

Description: Archive old ZIP files in user directories

Available:

Current Selections:

*Figure 2-56 Select the filesystems*

3. Select the **Options** tab as in Figure 2-57. Create a filter that specifies the path, the file type, and the access time requirements. In our example, we want to include both:
   - All .zip files in any directory below E:\Profiles.
   - All .zip files that have not been accessed in 168 days.
4. Click **Edit Filter** to build the filter as in Figure 2-58 on page 56.

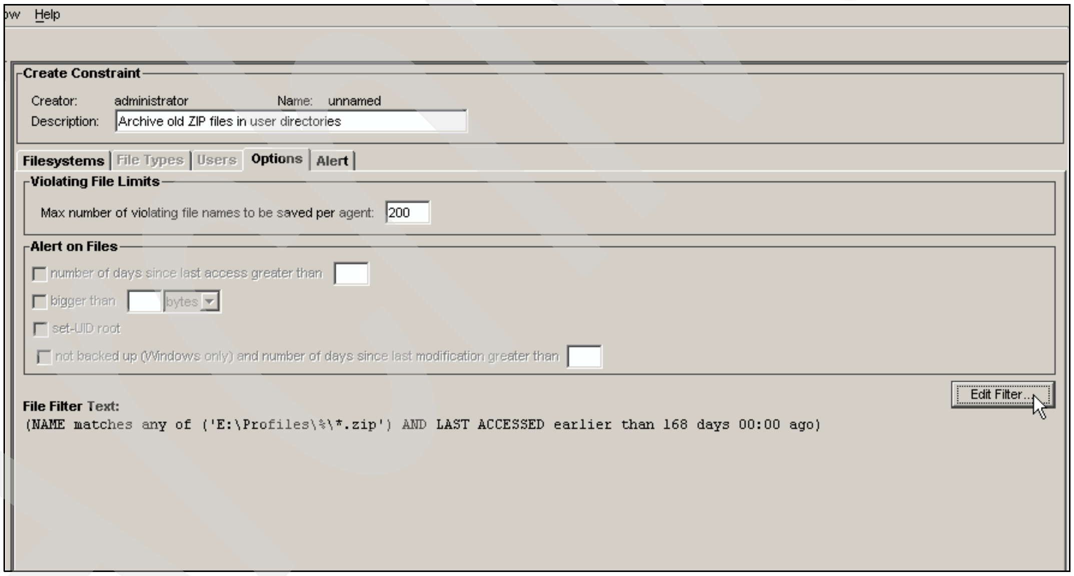

Create Constraint

Creator: administrator Name: unnamed

Description: Archive old ZIP files in user directories

Violating File Limits

Max number of violating file names to be saved per agent: 200

Alert on Files

□ number of days since last access greater than

□ bigger than bytes

□ set-UID root

□ not backed up (Windows only) and number of days since last modification greater than

File Filter Text:

```
(NAME matches any of ('E:\Profiles\%\*.zip') AND LAST ACCESSED earlier than 168 days 00:00 ago)
```

*Figure 2-57 Edit constraint filter*

5. Build a file filter to your requirements.

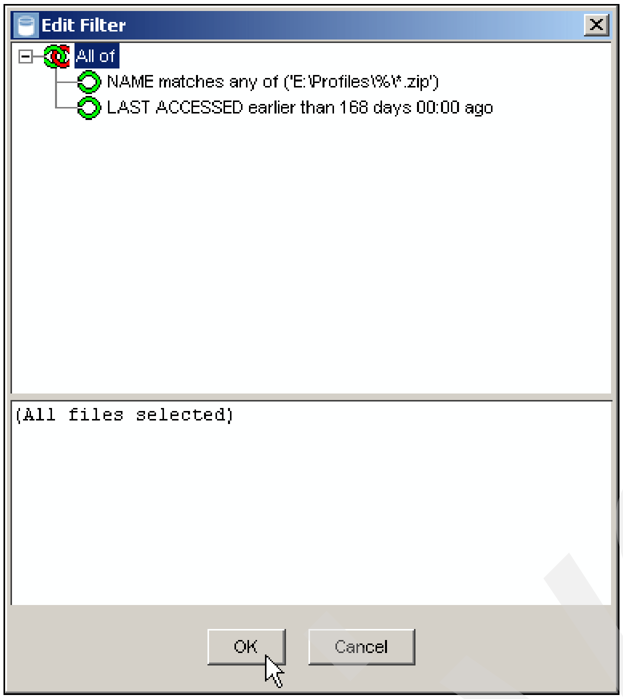
Edit Filter

All of
- NAME matches any of ('E:\Profiles\%\*.zip')
- LAST ACCESSED earlier than 168 days 00:00 ago

(All files selected)

*Figure 2-58 Edit Filter*

6. Select the **Alert** tab to configure Tivoli Storage Manager archiving and e-mail alerting as in Figure 2-59. Set the **Triggering Condition** to **Violating Files Consume More Than 1 Kilobytes** so that all files that meet the filter get archived.

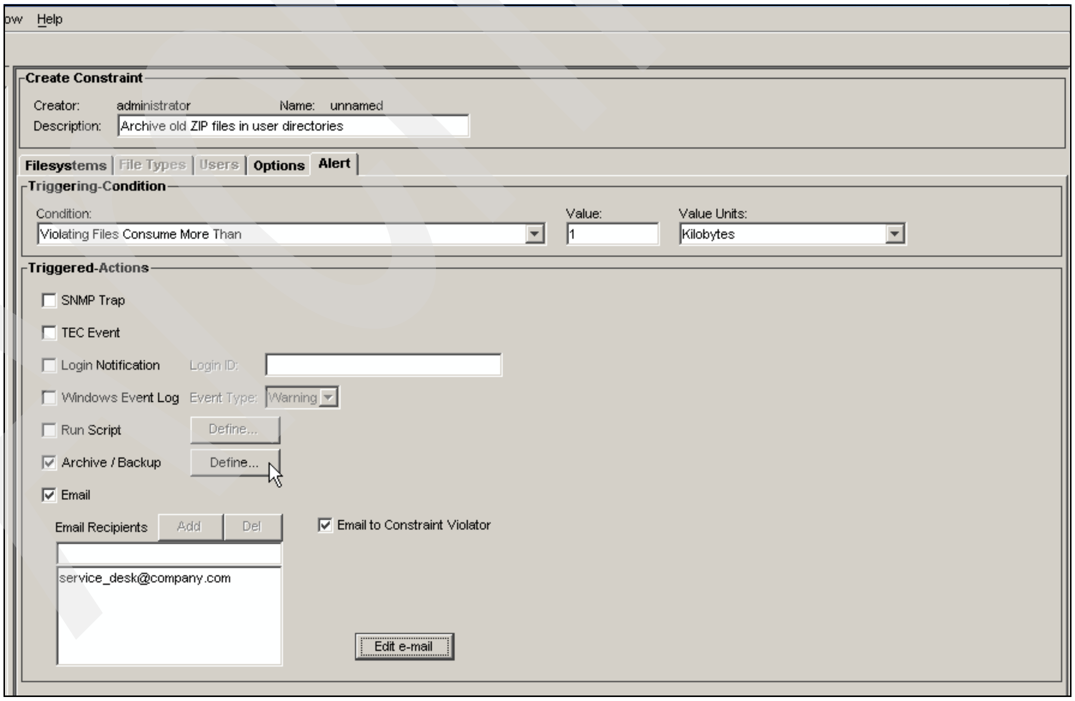

**Create Constraint**

Creator: administrator  Name: unnamed

Description: Archive old ZIP files in user directories

**Triggering-Condition**

Condition: Violating Files Consume More Than  Value: 1  Value Units: Kilobytes

**Triggered-Actions**

- ☐ SNMP Trap
- ☐ TEC Event
- ☐ Login Notification  Login ID:
- ☐ Windows Event Log  Event Type: Warning
- ☐ Run Script
- ☑ Archive / Backup
- ☑ Email

Email Recipients  ☑ Email to Constraint Violator

service_desk@company.com

*Figure 2-59 Configure Tivoli Storage Manager and e-mail alerting*

7. Check **Archive / Backup** and click **Define** to configure the Tivoli Storage Manager archiving.

8. In Figure 2-60, click **Archive** and also **Delete After Successful Archive**. This way, Tivoli Storage Manager deletes the files from storage when Tivoli Storage Manager is sure that it has a good copy of the files. Click **OK**.

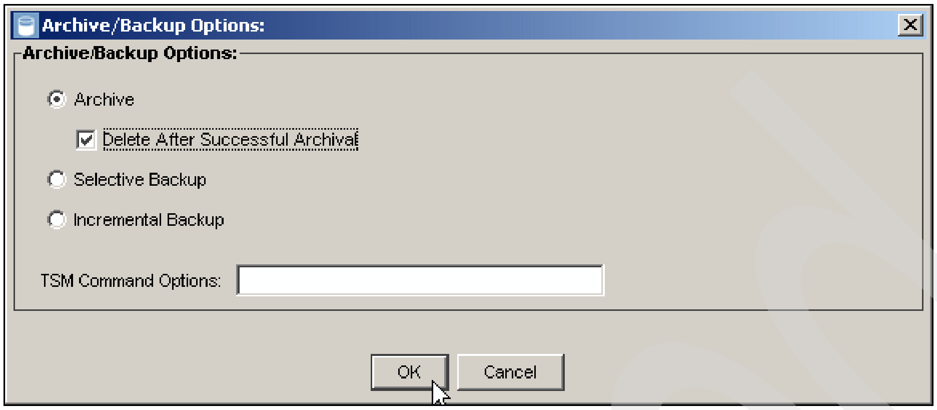

*Figure 2-60 Tivoli Storage Manager configuration panel*

> **Note:** If you need to specify extra parameters for Tivoli Storage Manager client passwords or non-default archive management classes in Tivoli Storage Manager, enter these command options on this panel.

9. In Figure 2-59 on page 56, add an e-mail recipient. This recipient is a storage administrator or monitored account. To send an e-mail to each user that details the user's files that got archived, check **Email to Constraint Violator**. Next, click **Edit e-mail** to create a meaningful e-mail to send to the users who have violated the constraint. Figure 2-61 shows the default system-generated e-mail, which is too generic and unsuitable for informing users about what has happened to their files.

10. Change the text of the **Subject** and **Text** to something more informative, such as Figure 2-62 on page 58.

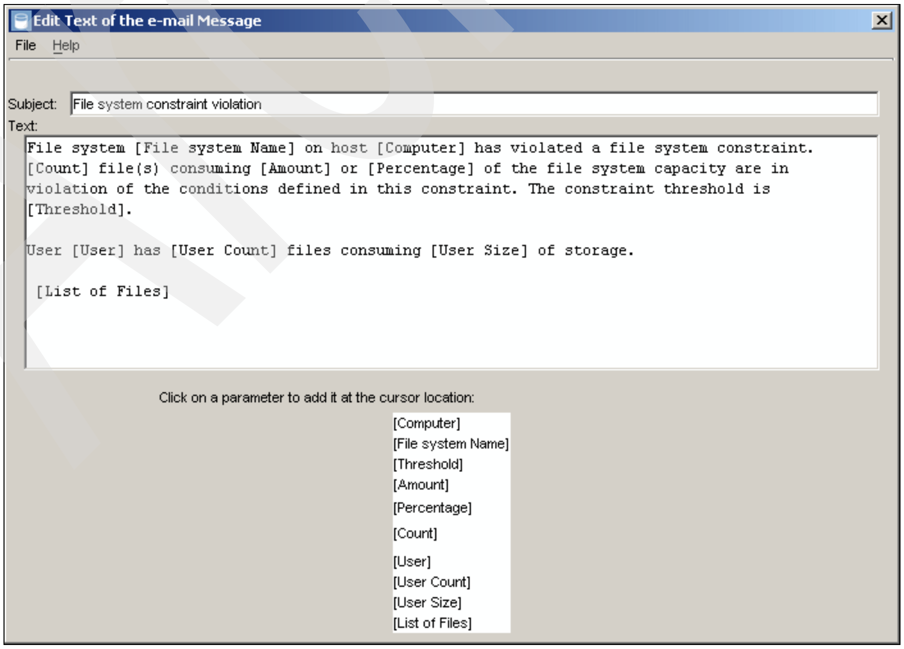

Subject: File system constraint violation

Text:

```
File system [File system Name] on host [Computer] has violated a file system constraint.
[Count] file(s) consuming [Amount] or [Percentage] of the file system capacity are in
violation of the conditions defined in this constraint. The constraint threshold is
[Threshold].

User [User] has [User Count] files consuming [User Size] of storage.

 [List of Files]
```

*Figure 2-61 Default system-generated e-mail*

11. Figure 2-62 shows a more informative e-mail for the user. It addresses the user with the [*User*] name, tells the user what happened, why, and how to get the user's files back in the future. This e-mail message also gives a list of the files that were affected by this action.

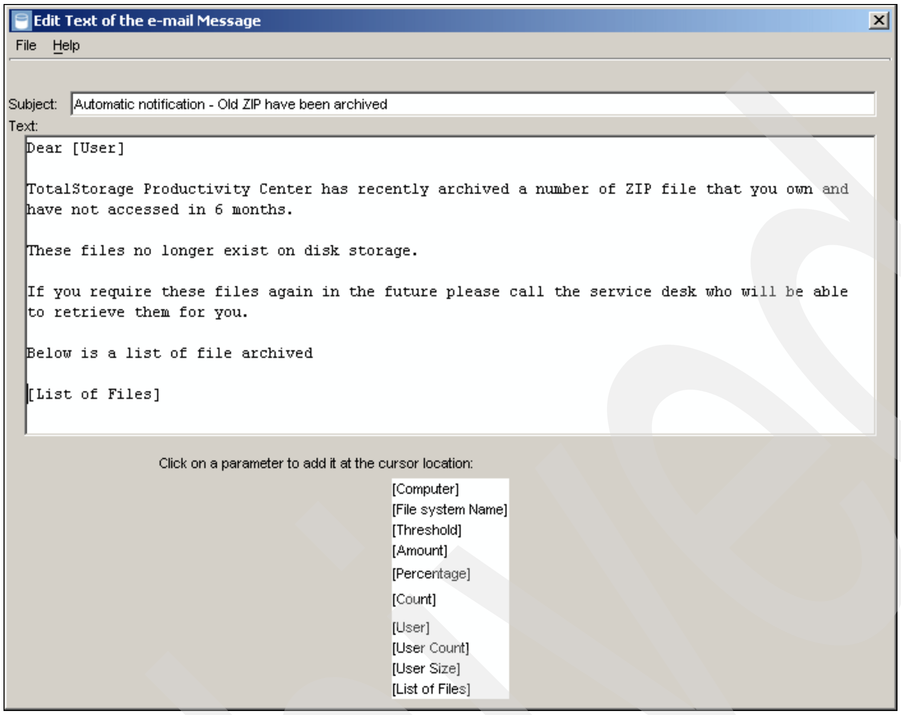

Subject: Automatic notification - Old ZIP have been archived

Text:

```
Dear [User]

TotalStorage Productivity Center has recently archived a number of ZIP file that you own and
have not accessed in 6 months.

These files no longer exist on disk storage.

If you require these files again in the future please call the service desk who will be able
to retrieve them for you.

Below is a list of file archived

[List of Files]
```

Click on a parameter to add it at the cursor location:

[Computer]
[File system Name]
[Threshold]
[Amount]
[Percentage]
[Count]
[User]
[User Count]
[User Size]
[List of Files]

*Figure 2-62 Customized user e-mail*

12. Save the new constraint with Ctrl+S and give it a name as in Figure 2-63. Click **OK**.

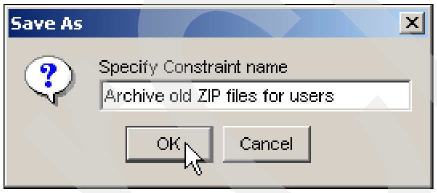

Specify Constraint name: Archive old ZIP files for users

*Figure 2-63 Save archive constraint*

### 2.5.3 Viewing the results of an archive event

TPC processes the constraint each time that a scan runs against the machines to which the constraint applies. If TPC finds files that meet the constraint, the scan process automatically calls Tivoli Storage Manager and archives the files.

The administrator and users receive an e-mail as in Example 2-2 on page 59 if files are archived.

*Example 2-2 Sample e-mail sent to a user*

```
Dear amit

TotalStorage Productivity Center has recently archived a number of .zip files that you own
and have not accessed in 6 months.

These files no longer exist on disk storage.

If you require these files again in the future, please call the service desk, which will be
able to retrieve them for you.

Below is a list of files archived.

Your largest violating files are:

E:\Profiles\Amit\my_zips.zip
E:\Profiles\Amit\backup_my_laptop.zip
E:\Profiles\Amit\A_loadoffiles.zip
```

To see in detail what the scan did and to see any messages from Tivoli Storage Manager, view the scan log for each machine as in Figure 2-64 on page 60. This shows you the standard Tivoli Storage Manager output for the archive process.

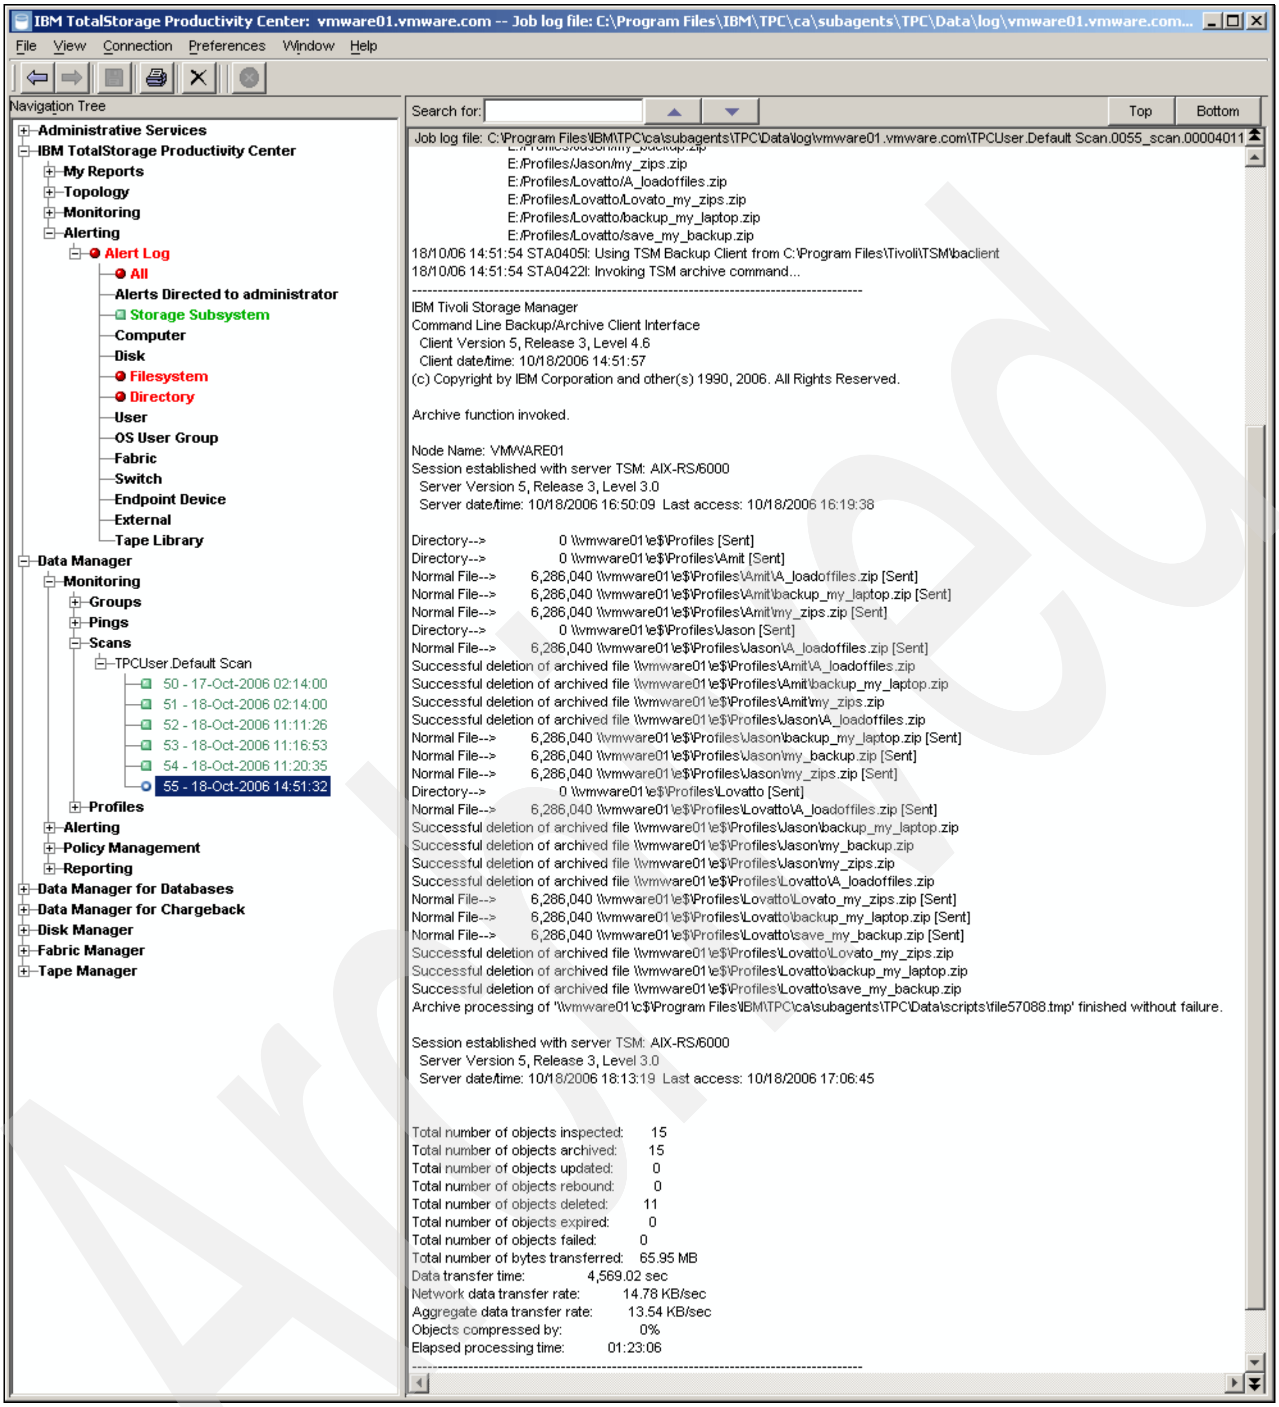
IBM TotalStorage Productivity Center: vmware01.vmware.com -- Job log file: C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\vmware01.vmware.com...
File View Connection Preferences Window Help
Navigation Tree
Administrative Services
IBM TotalStorage Productivity Center
My Reports
Topology
Monitoring
Alerting
Alert Log
All
Alerts Directed to administrator
Storage Subsystem
Computer
Disk
Filesystem
Directory
User
OS User Group
Fabric
Switch
Endpoint Device
External
Tape Library
Data Manager
Monitoring
Groups
Pings
Scans
TPCUser.Default Scan
50 - 17-Oct-2006 02:14:00
51 - 18-Oct-2006 02:14:00
52 - 18-Oct-2006 11:11:26
53 - 18-Oct-2006 11:16:53
54 - 18-Oct-2006 11:20:35
55 - 18-Oct-2006 14:51:32
Profiles
Alerting
Policy Management
Reporting
Data Manager for Databases
Data Manager for Chargeback
Disk Manager
Fabric Manager
Tape Manager
Search for:
Top
Bottom

```
Job log file: C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\vmware01.vmware.com\TPCUser.Default Scan.0055_scan.00004011
                      E:/Profiles/Jason/my_backup.zip
                      E:/Profiles/Jason/my_zips.zip
                      E:/Profiles/Lovatto/A_loadoffiles.zip
                      E:/Profiles/Lovatto/Lovato_my_zips.zip
                      E:/Profiles/Lovatto/backup_my_laptop.zip
                      E:/Profiles/Lovatto/save_my_backup.zip
18/10/06 14:51:54 STA0405I: Using TSM Backup Client from C:\Program Files\Tivoli\TSM\baclient
18/10/06 14:51:54 STA0422I: Invoking TSM archive command...
--------------------------------------------------------------------------------
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 3, Level 4.6
  Client date/time: 10/18/2006 14:51:57
(c) Copyright by IBM Corporation and other(s) 1990, 2006. All Rights Reserved.

Archive function invoked.

Node Name: VMWARE01
Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 3, Level 3.0
  Server date/time: 10/18/2006 16:50:09  Last access: 10/18/2006 16:19:38

Directory-->                 0 \\vmware01\e$\Profiles [Sent]
Directory-->                 0 \\vmware01\e$\Profiles\Amit [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Amit\A_loadoffiles.zip [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Amit\backup_my_laptop.zip [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Amit\my_zips.zip [Sent]
Directory-->                 0 \\vmware01\e$\Profiles\Jason [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Jason\A_loadoffiles.zip [Sent]
Successful deletion of archived file \\vmware01\e$\Profiles\Amit\A_loadoffiles.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Amit\backup_my_laptop.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Amit\my_zips.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Jason\A_loadoffiles.zip
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Jason\backup_my_laptop.zip [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Jason\my_backup.zip [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Jason\my_zips.zip [Sent]
Directory-->                 0 \\vmware01\e$\Profiles\Lovatto [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Lovatto\A_loadoffiles.zip [Sent]
Successful deletion of archived file \\vmware01\e$\Profiles\Jason\backup_my_laptop.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Jason\my_backup.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Jason\my_zips.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Lovatto\A_loadoffiles.zip
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Lovatto\Lovato_my_zips.zip [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Lovatto\backup_my_laptop.zip [Sent]
Normal File-->       6,286,040 \\vmware01\e$\Profiles\Lovatto\save_my_backup.zip [Sent]
Successful deletion of archived file \\vmware01\e$\Profiles\Lovatto\Lovato_my_zips.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Lovatto\backup_my_laptop.zip
Successful deletion of archived file \\vmware01\e$\Profiles\Lovatto\save_my_backup.zip
Archive processing of '\\vmware01\c$\Program Files\IBM\TPC\ca\subagents\TPC\Data\scripts\file57088.tmp' finished without failure.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 3, Level 3.0
  Server date/time: 10/18/2006 18:13:19  Last access: 10/18/2006 17:06:45


Total number of objects inspected:      15
Total number of objects archived:       15
Total number of objects updated:         0
Total number of objects rebound:         0
Total number of objects deleted:        11
Total number of objects expired:         0
Total number of objects failed:          0
Total number of bytes transferred:    65.95 MB
Data transfer time:                4,569.02 sec
Network data transfer rate:          14.78 KB/sec
Aggregate data transfer rate:        13.54 KB/sec
Objects compressed by:                  0%
Elapsed processing time:           01:23:06
--------------------------------------------------------------------------------
```

*Figure 2-64 Detailed scan log showing Tivoli Storage Manager messages*

3

# IBM Total Productivity Center database considerations

This chapter explains how to plan for backing up and restoring the TotalStorage Productivity Center (TPC) databases that reside on the TPC server in DB2® databases. The chapter covers both offline backup (cold backup) and online backup (hot backup) of the databases along with the merits of each type.

The topics include:

- Repository database backup and recovery
- Repository size planning for:
  - Disk subsystem performance collection
  - SAN Fabric performance collection
  - Statistical data: TPC for Data
- History aggregation
- Simple database tuning tips

The TPC product does not provide any extra backup and recovery tools over and above those tools already provided with the DB2 product. This chapter is not intended to be a comprehensive guide to all functions of backup and recovery built into DB2. Refer to the manual *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC09-4831, for detailed information about this subject.

This chapter also discusses TPC database repository growth and how you consolidate and eventually prune the amount of history stored in the TPC database repository. This chapter also covers basic tips for improving DB2 performance, which will be beneficial for those organizations that are planning large scale deployments of TPC.

## 3.1 Scripts provided

This book provides the Windows scripts listed in Table 3-1 as-is for your convenience. They are simple scripts to give you an understanding of how the process of backup works for filesystem type backups and Tivoli Storage Manager backups. Use these scripts as a basis for your own processes, and modify them as necessary.

All examples given in the chapter are based on a Windows platform TPC server.

*Table 3-1 Example scripts provided*

| Script name | Function |
|---|---|
| TPC_start.bat | Starts all the TPC services |
| TPC_stop.bat | Stops all the TPC services |
| TPC_backup_offline_file.bat | Backs up the TPC databases offline to a filesystem |
| TPC_backup_offline_tsm.bat | Backs up the TPC databases offline to Tivoli Storage Manager |
| TPC_backup_online_file.bat | Backs up the TPC databases online to a filesystem |
| TPC_backup_online_tsm.bat | Backs up the TPC databases online to Tivoli Storage Manager |
| database_list_offline_file.txt | Lists the databases to back up for an offline filesystem |
| database_list_offline_tsm.txt | Lists the databases to back up for offline to Tivoli Storage Manager |
| database_list_online_file.txt | Lists the databases to back up for online filesystem |
| database_list_online_tsm.txt | Lists the databases to back up for online to Tivoli Storage Manager |

## 3.2 Database backup

This section describes the high level points that you need to understand about DB2 UDB before you can plan a backup strategy for securing TotalStorage Productivity Center (TPC) Version 3.

### Backup Types

There are two primary methods of backing up DB2 databases:

- **Offline backup** (sometimes known as cold backup) is when all database access is terminated, and the database is closed. The backup then runs standalone before the database is restarted and access-enabled. This is the simplest type of backup to set up, configure, and maintain.
- **Online backup** (sometimes known as hot backup) is when all user and application database access continues to run while the backup process takes place. This type of backup provides for continuous availability of the database and the applications that require it. This is a more complex type of backup to set up, configure, and maintain.

### Backup output destination

You can direct database backup output to a number of destinations from within DB2:

- **Filesystem:** Direct output to normal filesystem structure flat files. Then, you can copy these files to removable tape for added security or back them up with products, such as Tivoli Storage Manager or other widely available similar tools.
- **Tape:** Send output directly to tape as long as the tape device is directly attached to the server hardware.
- **Tivoli Storage Manager:** Sends output directly to Tivoli Storage Manager through direct integration between the two products. If a Tivoli Storage Manager environment exists within your organization, you can back up directly to it by installing the Tivoli Storage Manager Backup/Archive client and client API on the same machine that hosts the TPC DB2 databases.
- **XBSA:** Directs output to an X/Open Backup Services Application (X/BSA) compliant application, such as Legato NetWorker.
- **Vendor DLL:** Directs output to a third-party vendor written interface API.

> **Note:** This chapter focuses on the filesystem and Tivoli Storage Manager backup destinations.

### Database logging

DB2 UDB uses log files to keep a sequential record of all database changes. The log files are specific to DB2 UDB activity. The logs record the database activity in transactions. If there is a crash, you use logs to play back or redo committed transactions during recovery.

There are two types of logging:

- **Circular logging** (default): This is the simplest method and the default logging type that TPC uses.
- **Archive logging**: This type of logging enables online backup as well as roll-forward recovery of a database to a point-in-time. It is, however, more complex to manage.

# 3.3 Database backup method considerations

This section considers the merits of offline backup methods compared to online backup methods for the TPC databases. The default method of backup for TPC is to use offline backup.

## 3.3.1 Offline backup advantages and disadvantages

The advantages and disadvantages of offline backup are:

- Advantages:
  - **Simple:** You can perform offline backup with DB2 logging set to the default circular method.
  - **DB2 skills:** Offline backup requires a minimum amount of DB2 skills DB2 to perform, because it is the simplest method of backup.
  - **Logging:** Circular logs are the simplest to manage and maintain.

- Disadvantages:
  - **Stopped TPC server services:** The offline method entails stopping all of the TPC server services on a regular basis (typically daily) to perform the backup. This regular outage might not be acceptable to all organizations that want to use TPC.
  - **Missed performance data collection:** If you have set up TPC to continuously collect disk subsystem and SAN fabric performance statistics, you lose data points for the duration that TPC is down each day for backup. You can minimize the impact of this loss by scheduling the backup at a time when the monitored equipment statistics are of little importance from a reporting perspective. This loss of data points might not be acceptable to all organizations wanting to use TPC.
  - **Missed events:** TPC monitors the infrastructure and alerts you about events, such as failures within a SAN fabric. You run the risk that you can miss critical events if the events occur when you stop the TPC server services for the backup process.

#### *Online backup advantages and disadvantages*

The advantages and disadvantages of online backup are:

- Advantages:
  - **Greater availability:** You do not need to stop and start the The TPC server services on a daily basis for the backup operation. Online backups do not interrupt user access to the database while the backup operation is in progress.
  - **No missed events:** TPC monitors the infrastructure and alerts you about events, such as failures within a SAN fabric. Using online backup ensures that TPC is able to respond quickly to critical events at any time of the day.
  - **Uninterrupted performance collection:** You experience no interruption or missing data points in the collection of performance data from disk subsystems and SAN fabrics.
- Disadvantages:
  - **More DB2 skills required:** Archive logging is a more advanced method of DB2 operation, and administering archive logging requires more skills.
  - **TPC software update process can fail:** Our testing found that TPC software updates that alert the database layout can fail. You need to revert to circular logging to perform updates. Then, switch back to archive logging.

## 3.4 Common backup setup steps

This section describes the first setup steps that you need to perform for both filesystem and Tivoli Storage Manager backups:

1. Configure the DB2 history file to retain the number of backup versions that you want to retain. Your organization might already have a policy for how many versions you need to keep.

   You need to change the DB2 parameter **num_db_backups** and set the value to the number of backup versions that you require. You also need to set the **rec_his_retentn** parameter to a value of `-1`. By setting this value to -1, **rec_his_retentn** follows the value set in **num_db_backups.**

**Important:** Changing this value requires a stop and a start of Agent Manager and TPC services to take effect. This restart does not necessarily need to happen directly after you change the parameter.

2. Start a DB2 command line processor window as shown in Figure 3-1.

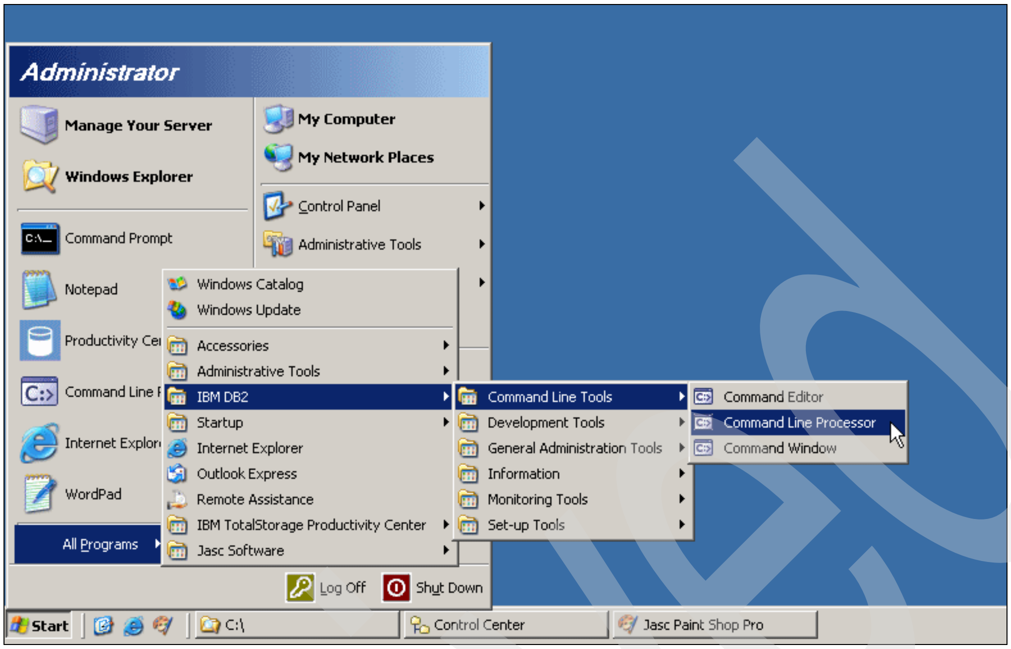

*Figure 3-1 Launch the DB2 command line processor*

3. A command line processor window appears as in Figure 3-2.

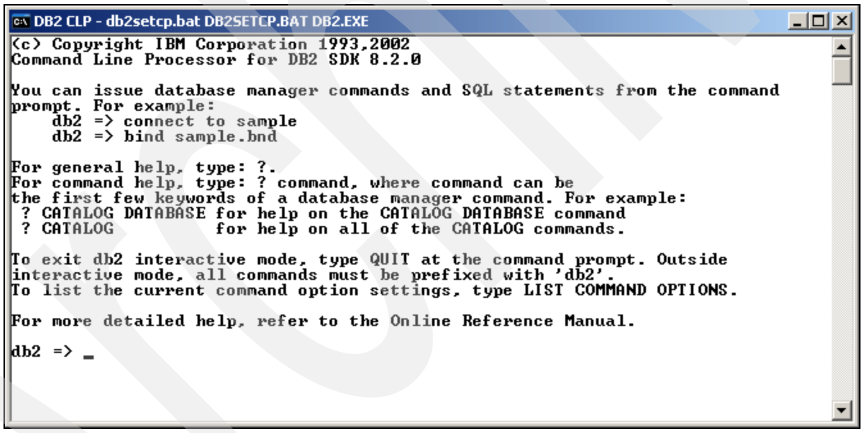

```
(c) Copyright IBM Corporation 1993,2002
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _
```

*Figure 3-2 DB2 command line processor*

4. Example 3-1 on page 66 shows how to set the **num_db_backups** value to `4` versions and **rec_his_retentn** to `-1` for both the Agent Manager and the TPC databases.

5. Issue the following commands at the **db2 =>** prompt in the command line processor window.

*Example 3-1   DB2 commands to configure how many backup versions to keep*

```
connect to TPCDB
update db cfg using num_db_backup 4
update db cfg using rec_his_retentn -1
disconnect TPCDB

connect to IBMCDB
update db cfg using num_db_backup 4
update db cfg using rec_his_retentn -1
disconnect IBMCDB

exit
```

> **Important:** When you set new values for **num_db_backups** and **rec_his_retentn**, the new values are not effective until you stop all of the database connections.

6. Restart TPC and Agent Manager to make the changes effective. You can either reboot the server, or alternatively stop and start the services.

   You can either stop and start the services through the Windows Services interface or open a command prompt window and issue the commands in Example 3-2. This process applies to Windows servers only.

*Example 3-2   Windows commands to stop and start TPC services*

```
net stop “IBM TotalStorage Productivity Center - Data Server”
net stop “IBM WebSphere Application Server V5 - Device Server”
net stop “IBM WebSphere Application Server V5 - Tivoli AGENT MANAGER”

net start “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
net start “IBM WebSphere Application Server V5 - Device Server”
net start “IBM TotalStorage Productivity Center - Data Server”
```

## 3.5  Offline backup to filesystem setup steps

This section describes how to set up offline backup for the TPC server databases to flat files in a filesystem. Because the offline backup method is the default method for TPC, there is little DB2 configuration needed before you can perform a backup.

> **Note:** Ensure that you perform the steps in 3.4, “Common backup setup steps” on page 64 as well as these steps.

The steps are:

1. Choose a location to use for the DB2 backup output. Choose a directory that has enough free space to hold the number of backups that you plan to retain. We advise that you use a separate filesystem rather than the filesystem that contains the DB2 database.

   You can choose to use a location that is a remotely mounted CIFS or NFS filesystem, so that the backup data is secured to another server, perhaps at another location in your organization.

   This example uses **E:\TPC_database_backups**

**Important:** DB2 does not create this directory for you. Create this directory before you attempt a backup.

2. Create a batch script to control the backup process.

   We based our example on a TPC installation on Windows.

   There are two files. The script, which is shown in Example 3-3, to run the backup is:

   **C:\scripts\TPC_backup_offline_file.bat**

*Example 3-3 File C:\scripts\TPC_backup_offline_file.bat*

```
@echo on

@REM This is a sample backup script
@REM To backup TPC offline
@REM To disk filesystems


@REM Stopping TotalStorage Productivity Center services
@REM --------------------------------------------------

net stop "IBM TotalStorage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V5 - Device Server"
net stop "IBM WebSphere Application Server V5 - Tivoli Agent Manager"

@REM Starting backup of the DB2 databases
@REM ------------------------------------

C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf C:\scripts\database_list_offline_file.txt

@REM Restarting TotalStorage Productivity Center services
@REM ----------------------------------------------------

net start "IBM WebSphere Application Server V5 - Tivoli Agent Manager"
net start "IBM TotalStorage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V5 - Device Server"

@REM Offline backup process complete
@REM -------------------------------
```

3. The DB2 scripted list of databases ( shown in Example 3-4) to back up is:

   **C:\scripts\database_list_offline_file.txt**

*Example 3-4 File C:\scripts\database_list_offline_file.txt*

```
backup database IBMCDB to “E:\TPC_database_backups” without prompting;
backup database TPCDB to “E:\TPC_database_backups” without prompting;
```

See 3.9.1, “Performing an offline backup to a filesystem” on page 81 to run an offline backup.

## 3.6 Offline backup to Tivoli Storage Manager setup steps

This section describes the steps necessary to set up an offline backup of the TPC server databases to a Tivoli Storage Manager server. The backup to Tivoli Storage Manager is a little more complex to set up but does not require you to set aside large amounts of local disk space for backup versions on the TPC server.

This section assumes:

- You have a basic working knowledge of Tivoli Storage Manager.
- An operational Tivoli Storage Manager server already exists to which you can send backup data.
- Your Tivoli Storage Manager administrator has defined storage, which will receive the backups, to the policies.
- You have already installed a Tivoli Storage Manager Backup/Archive client on the TPC server, and you have configured it to do standard file backups.
- You have installed the Tivoli Storage Manager API Client on the TPC server.
- You used default installation paths for Tivoli Storage Manager.

> **Note:** You need to stop TPC and DB2 as part of this configuration process. We recommend that you reboot the TPC server to complete the configuration process, because this process also adds operating system environment variables. Plan this exercise at a time when you can reboot the TPC server.

Follow these steps to configure DB2 to Tivoli Storage Manager integration.

## 3.6.1 Add new variables to Windows

Table 3-2 shows a list of Tivoli Storage Manager API environment variables to add to Windows. The values shown assume a default installation of Tivoli Storage Manager on the TPC server.

*Table 3-2 System environment variables*

| Environment variable name | Value |
|---|---|
| DSMI_DIR | C:\Program Files\Tivoli\TSM\baclient |
| DSMI_CONFIG | C:\Program Files\Tivoli\TSM\baclient\dsm.opt |
| DSMI_LOG | C:\tsm |

The steps to add new variables to Windows are:

1. Figure 3-3 on page 69 shows the Windows System Properties panel. Click **Environment Variables** to proceed to the next step.

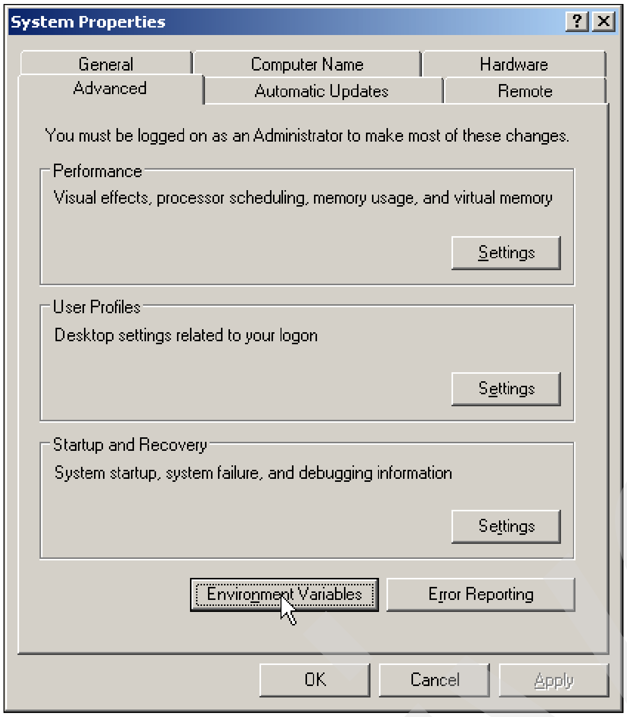

You must be logged on as an Administrator to make most of these changes.

Performance: Visual effects, processor scheduling, memory usage, and virtual memory

User Profiles: Desktop settings related to your logon

Startup and Recovery: System startup, system failure, and debugging information

*Figure 3-3 Windows System Properties*

2. Select **New** from the Environment Variables panel as in Figure 3-4.

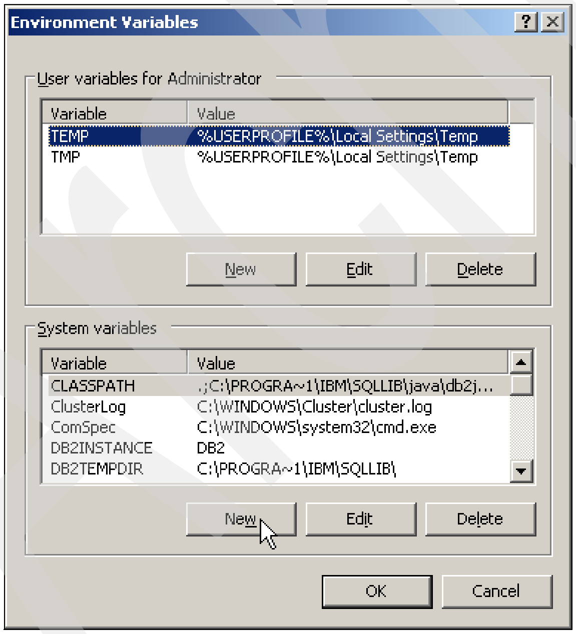

User variables for Administrator

| Variable | Value |
|---|---|
| TEMP | %USERPROFILE%\Local Settings\Temp |
| TMP | %USERPROFILE%\Local Settings\Temp |

System variables

| Variable | Value |
|---|---|
| CLASSPATH | .;C:\PROGRA~1\IBM\SQLLIB\java\db2j... |
| ClusterLog | C:\WINDOWS\Cluster\cluster.log |
| ComSpec | C:\WINDOWS\system32\cmd.exe |
| DB2INSTANCE | DB2 |
| DB2TEMPDIR | C:\PROGRA~1\IBM\SQLLIB\ |

*Figure 3-4 Windows Environment Variables*

3. Add all three new system variables that are listed in Table 3-2 on page 68. Repeat the add process as seen in Figure 3-5 on page 70 for each variable.

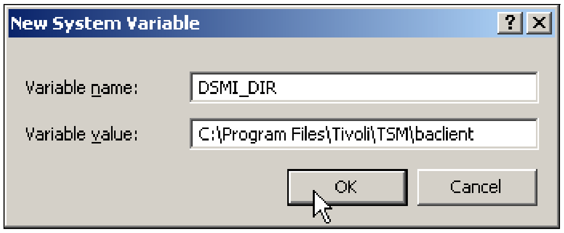

Variable name: DSMI_DIR

Variable value: C:\Program Files\Tivoli\TSM\baclient

*Figure 3-5 Adding a new System Variable*

## 3.6.2 Configure Tivoli Storage Manager option file and password

This section describes the steps necessary to configure the Tivoli Storage Manager option file dsm.opt and then set the Tivoli Storage Manager password so that the DB2 backup process can communicate with the Tivoli Storage Manager API. The steps are:

1. Edit the dsm.opt, which is located in **C:\Program Files\Tivoli\TSM\baclient** by default.
2. Set the client option PASSWORDACCESS GENERATE as shown in Figure 3-6 and save the file.

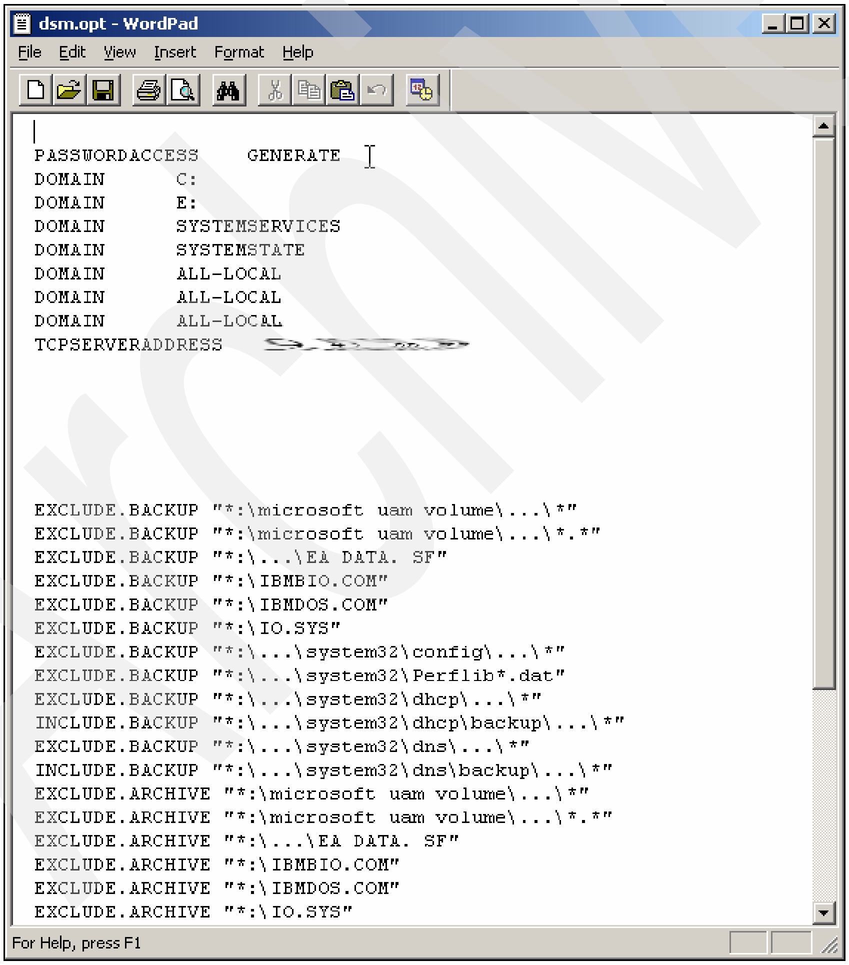

```
PASSWORDACCESS     GENERATE
DOMAIN       C:
DOMAIN       E:
DOMAIN       SYSTEMSERVICES
DOMAIN       SYSTEMSTATE
DOMAIN       ALL-LOCAL
DOMAIN       ALL-LOCAL
DOMAIN       ALL-LOCAL
TCPSERVERADDRESS    [redacted]




EXCLUDE.BACKUP "*:\microsoft uam volume\...\*"
EXCLUDE.BACKUP "*:\microsoft uam volume\...\*.*"
EXCLUDE.BACKUP "*:\...\EA DATA. SF"
EXCLUDE.BACKUP "*:\IBMBIO.COM"
EXCLUDE.BACKUP "*:\IBMDOS.COM"
EXCLUDE.BACKUP "*:\IO.SYS"
EXCLUDE.BACKUP "*:\...\system32\config\...\*"
EXCLUDE.BACKUP "*:\...\system32\Perflib*.dat"
EXCLUDE.BACKUP "*:\...\system32\dhcp\...\*"
INCLUDE.BACKUP "*:\...\system32\dhcp\backup\...\*"
EXCLUDE.BACKUP "*:\...\system32\dns\...\*"
INCLUDE.BACKUP "*:\...\system32\dns\backup\...\*"
EXCLUDE.ARCHIVE "*:\microsoft uam volume\...\*"
EXCLUDE.ARCHIVE "*:\microsoft uam volume\...\*.*"
EXCLUDE.ARCHIVE "*:\...\EA DATA. SF"
EXCLUDE.ARCHIVE "*:\IBMBIO.COM"
EXCLUDE.ARCHIVE "*:\IBMDOS.COM"
EXCLUDE.ARCHIVE "*:\IO.SYS"
```

*Figure 3-6 Example of the dsm.opt*

3. Now, set the Tivoli Storage Manager password so that DB2 can authenticate with the Tivoli Storage Manager server when DB2 performs a backup or restore operation:

   a. Open a Windows Command prompt window.

   b. Change to **C:\Program Files\IBM\SQLLIB\adsm**

   a. Run the **dsmapipw** command as shown in Figure 3-7.

   b. Enter the current and new Tivoli Storage Manager password. You can reuse the existing Tivoli Storage Manager password.

> **Important:** You must run the **dsmapipw** command even if you do not intend to change the Tivoli Storage Manager password. Running this command registers it with the Tivoli Storage Manager API. Registering this password in the setup phase means that a DB2 operator can perform backup and restore operations without needing to know the Tivoli Storage Manager client password. If a Tivoli Storage Manager administrator changes or resets the Tivoli Storage Manager password, you need to run the **dsmapipw** command again.

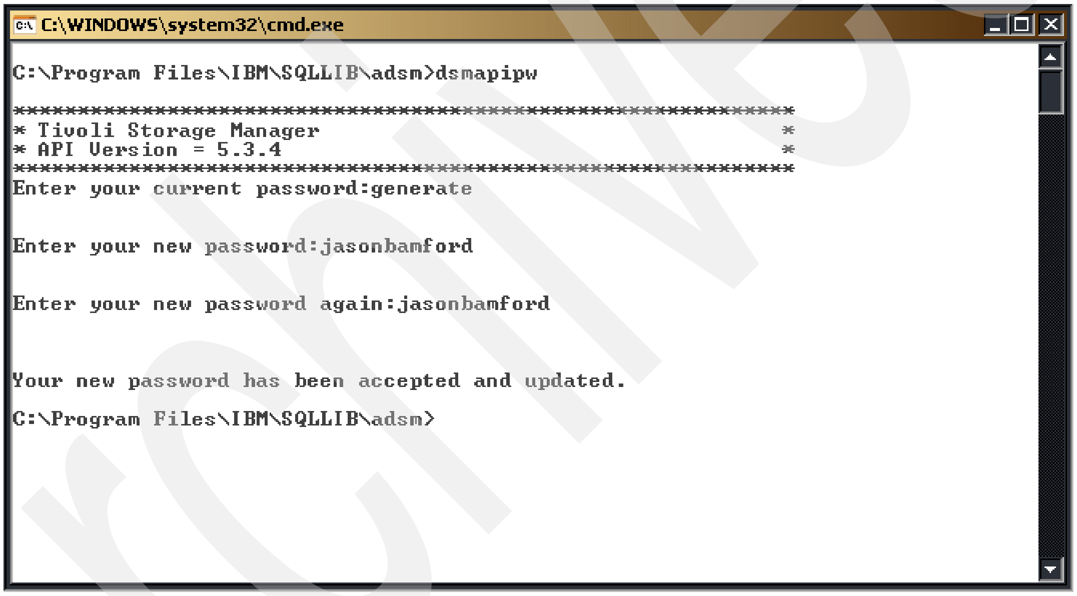

```
C:\Program Files\IBM\SQLLIB\adsm>dsmapipw

*************************************************************
* Tivoli Storage Manager                                    *
* API Version = 5.3.4                                       *
*************************************************************
Enter your current password:generate


Enter your new password:jasonbamford


Enter your new password again:jasonbamford



Your new password has been accepted and updated.

C:\Program Files\IBM\SQLLIB\adsm>
```

*Figure 3-7 Running the dsmapipw command*

> **Important:** The **dsmapipw** command displays both the old and new passwords on the window in plain text. Ensure that you perform this task in a secure area to prevent password exposure.

## 3.6.3 Reboot the TPC server

Now that you have completed the configuration steps, reboot the TPC server to ensure that the environment variables are picked up by DB2.

### 3.6.4 Create an offline backup to Tivoli Storage Manager script

We based the script in Example 3-5 on a TPC installation on Windows.

Create two files. The first is the script that you run:

**C:\scripts\TPC_backup_offline_tsm.bat**

*Example 3-5 File C:\scripts\TPC_backup_offline_tsm.bat*

```
echo on

REM This is a sample backup script
REM To backup TPC offline
REM To Tivoli Storage Manager


REM Stopping TotalStorage Productivity Center services
REM --------------------------------------------------

net stop "IBM TotalStorage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V5 - Device Server"
net stop "IBM WebSphere Application Server V5 - Tivoli Agent Manager"

REM Starting backup of the DB2 databases
REM ------------------------------------

C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf C:\scripts\database_list_offline_tsm.txt

REM Restarting TotalStorage Productivity Center services
REM ----------------------------------------------------

net start "IBM WebSphere Application Server V5 - Tivoli Agent Manager"
net start "IBM TotalStorage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V5 - Device Server"

REM Offline backup process complete
REM -------------------------------
```

The second file is the DB2 scripted list of databases to backup:

**C:\scripts\database_list_offline_tsm.txt**

*Example 3-6 File C:\scripts\database_list_offline_tsm.txt*

```
backup database IBMCDB use tsm without prompting;
backup database TPCDB use tsm without prompting;
```

## 3.7 Online backup to Tivoli Storage Manager setup steps

This section describes the steps that are necessary to configure the TPC databases to enable for online backup to Tivoli Storage Manager. The significant difference between online and offline backup is the need to enable archive logging on the databases. As we discussed in 3.3, “Database backup method considerations” on page 63, operating in this mode provides many backup and recovery benefits at the expense of increased complexity in the database operation.

Take time to consider the advantages and disadvantages of archive logging before continuing with this setup. For full details of DB2 logging methods, refer to the DB2 product manuals. See *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC09-4831, for detailed information about this subject.

**Note:** You need to stop the TPC and Agent Manager to perform these tasks.

DB2 requires a full backup of each database before you can start the TPC and Agent Manager databases again after these reconfiguration steps. We include the instructions to perform a full backup of each database. Allow time in your outage planning for the backups to complete.

Also, complete the steps in 3.4, “Common backup setup steps” on page 64 to set the number of backup versions that you want to retain in the history file.

**Important:** If you set up DB2 for online backup to Tivoli Storage Manager, you cannot easily change to an online backup to filesystem. You need to choose between these methods, because you are setting the destination for the archive logging process. If you decide in the future to change to the online filesystem method, you will need to reconfigure DB2 to send the archive logs to filesystem. This reconfiguration requires a TPC restart to complete the task.

It is possible to perform an online backup to filesystem and have the archive logs going to Tivoli Storage Manager. However, we do not recommend that you do this because the difficulty of managing and tracking information makes this a poor practice.

Set up and test DB2 to Tivoli Storage Manager integration before you attempt this section. Use 3.6, “Offline backup to Tivoli Storage Manager setup steps” on page 67. When you are satisfied that DB2 is communicating with Tivoli Storage Manager and you have performed at least one successful offline backup, return to this section.

### 3.7.1 DB2 parameter changes for archive logging to Tivoli Storage Manager

To set up archive logging to Tivoli Storage Manager, complete the following tasks:

1. You need to make a number of parameter choices for the configuration of archive logging as seen in Table 3-3 on page 74. These parameters determine where DB2 keeps its log files, the number of log files, and the size of the log files.

*Table 3-3 DB2 parameters*

| DB2 parameter | Example value | Comment |
|---|---|---|
| Primary log path | C:\DB2_active_logs | This is the location where DB2 keeps the current logs for the database. For best performance, place these logs on a separate volume than the volume that holds the data. |
| Failed log path | E:\DB2_failed_log | This is the location where DB2 put log files if the archive process fails. This can happen if Tivoli Storage Manager is down or unreachable when DB2 tries to send a log file to Tivoli Storage Manager. |
| Number of primary logs | 8 | This is the number of primary log files that DB2 creates in the primary log path. This example uses a low value. If your TPC installation will manage many agents and disk subsystems, you need to increase this number significantly. |
| Number of secondary logs | 16 | This is the number of log files that DB2 can use if it runs out of primary log space. This can happen if there are very long running transactions in the database. |
| Log file size | 2500 | This is the size of each primary and secondary log file. The value is in 4 K blocks. This example is, therefore, 10 MB log files. |

2. Stop TPC and Agent Manager services using the commands in Example 3-7. You can also perform this task through the Windows Services interface.

*Example 3-7 Windows commands to stop TPC*

```
net stop “IBM TotalStorage Productivity Center - Data Server”
net stop “IBM WebSphere Application Server V5 - Device Server”
net stop “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
```

3. Launch a DB2 command line processor as in Figure 3-8 on page 75 and issue the commands as in Example 3-8 on page 76.

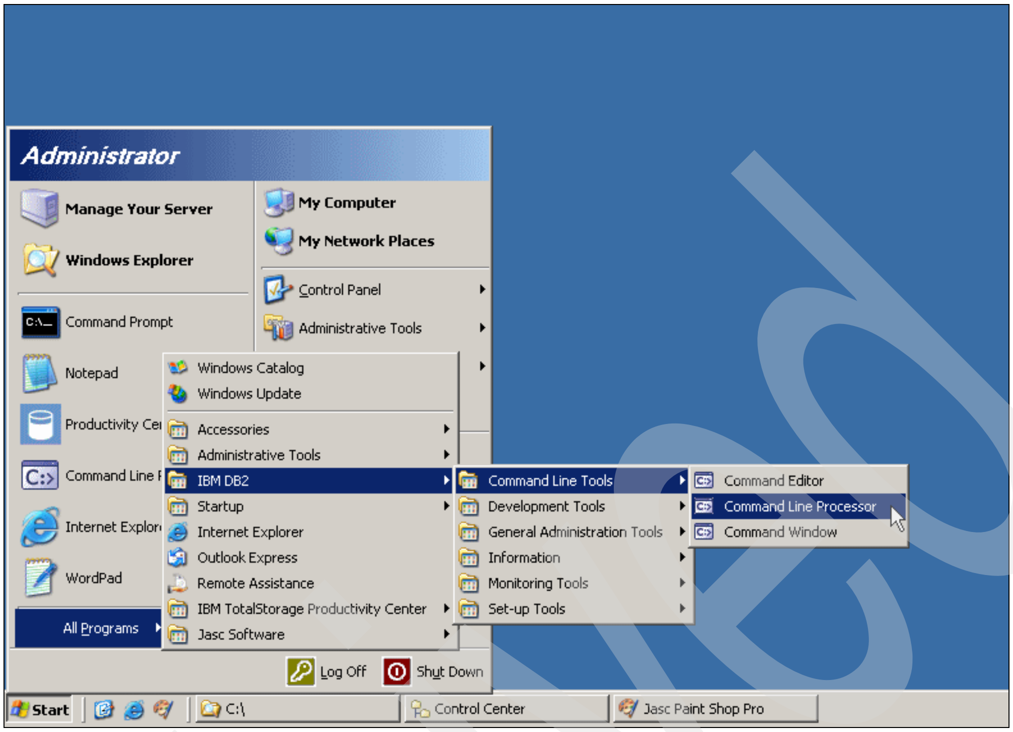

*Figure 3-8 Launch a DB2 command line processor*

A command line processor appears as in Figure 3-9.

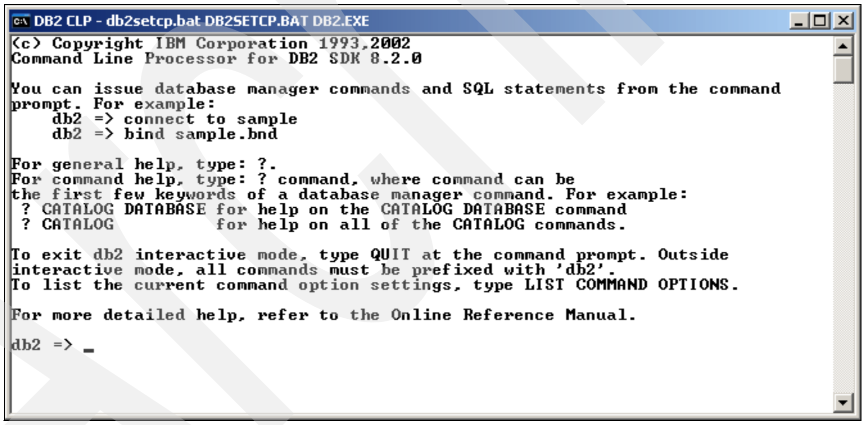

```
(c) Copyright IBM Corporation 1993,2002
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _
```

*Figure 3-9 DB2 command line processor*

4. Issue the commands from Example 3-8 on page 76 in the command line processor window. Substitute your chosen values for the parameters that form part of the **UPDATE DB CFG** command. See Table 3-3 on page 74. Note that the final two commands perform an offline backup of both databases.

> **Important:** The database backups are required after this reconfiguration, and the DB2 databases will not open again until the database backups are completed.

*Example 3-8 DB2 command to configure archive logging to Tivoli Storage Manager*

```
CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 TSM failarchpath "E:\DB2_failed_logs" logprimary
8 logsecond 16 logfilsiz 2500 newlogpath C:\DB2_active_logs\TPCD

DISCONNECT TPCDC

CONNECT TO IBMCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR IBMCDB USING logarchmeth1 TSM failarchpath "E:\DB2_failed_logs"
logprimary 8 logsecond 16 logfilsiz 2500 newlogpath C:\DB2_active_logs\IBMCDB

DISCONNECT IBMCDB

BACKUP DATABASE TPCDB USE TSM

BACKUP IBMCDB USE TSM
```

5. When both of the database backups are complete, you can restart Agent Manager and TPC. Either use the Windows Services interface or issue the commands shown in Example 3-9 in a command prompt window.

*Example 3-9 Start Agent Manager and TPC*

```
net start “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
net start “IBM TotalStorage Productivity Center - Data Server”
net start “IBM WebSphere Application Server V5 - Device Server”
```

### 3.7.2 Create online backup script for Tivoli Storage Manager

We based this example on a TPC installation on Windows.

Create two files. The first file is the script (Example 3-10 on page 77) that you run to start the backup:

**C:\scripts\TPC_backup_online_tsm.bat**

*Example 3-10 File C:\scripts\TPC_backup_online_tsm.bat*

```
echo on

REM This is a sample backup script
REM To backup TPC online
REM To Tivoli Storage Manager

REM Starting backup of the DB2 databases
REM ------------------------------------

C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf C:\scripts\database_list_online_tsm.txt

REM Offline backup process complete
REM -------------------------------
```

The second file (Example 3-11) is the DB2 scripted list of databases to back up:

**C:\scripts\database_list_online_file.txt**

*Example 3-11 File C:\scripts\database_list_online_tsm.txt*

```
backup database IBMCDB online use tsm without prompting;
backup database TPCDB online use tsm without prompting;
```

## 3.8 Online backup to a filesystem setup steps

Performing online backups to a filesystem requires you to set up archive logging to a filesystem also. When operating with this method, DB2 does not clean up old and no longer necessary archive log files. Therefore, you need to put processes in place to clean up old log files after a specific amount of time to prevent the system from filling up. You also need to plan for this amount of space. The log space required for a TPC database can become many times larger than the database over a number of weeks.

To be able to restore an online DB2 database taken two weeks ago, for example, you need to have log files going back to that same date that you can restore. An online DB2 database backup is not standalone, because you cannot restore the online DB2 database backup without at least some logs for it to roll forward to a consistent state.

> **Important:** Although it is straightforward to switch between a backup destination of online to a filesystem and online to Tivoli Storage Manager, it is not so easy to switch the logging path. To switch the logging from Tivoli Storage Manager to a filesystem requires a stop and a start of the database and, therefore, a stop and a start of TPC.
>
> We recommend that you choose either Tivoli Storage Manager backup or a filesystem backup and stay with that specific method.

### 3.8.1 Set up DB2 archive logging to a filesystem

Set up DB2 archive logging to a filesystem using these steps:

1. You need to make a number of parameter choices to configure archive logging. See Table 3-4 on page 78. These parameters determine where DB2 will keep its log files, how many log files to keep, and the size of the log files.

*Table 3-4 DB2 parameters for archive logging to a filesystem*

| DB2 parameter | Example value | Comment |
|---|---|---|
| Primary log path | C:\DB2_active_logs | The location where DB2 will keep the current logs for the database. *For best performance, place the logs on a separate volume than the data.* |
| Archive log path | C:\DB2_archive_logs\TPCD<br>and<br>C:\DB2_archive_logs\IBMCDB | The location where DB2 will archive log files for both the TPCD and IBMCDB databases. |
| Failed log path | E:\DB2_failed_log | This is a location where DB2 will put log files if the archive process fails, which can happen if the filesystem for the primary logs fills up. *Choose a location that is NOT on the same filesystem as the archive logs.* |
| Number of primary logs | 8 | This is the number of primary log files DB2 will create in the primary log path. This example uses a low value. If your TPC installation will manage many agents and disk subsystems, you will need to increase this number significantly. |
| Number of secondary logs | 16 | This is the number of logs that DB2 can use if it runs out of primary log space. This can happen if there are very long running transactions in the database. |
| Log file size | 2500 | This is the size of each primary and secondary log file. The value is in 4 k blocks. This example is, therefore, 10 MB log files. |

2. Choose a filesystem path to store the DB2 database backups.

*Table 3-5 Filesystem location for database backups*

| Database backup path |
|---|
| E:\TPC_database_backups |

3. Stop TPC and Agent Manager services by using the commands in Example 3-12. You can also perform this task through the Windows Services interface.

*Example 3-12 Windows commands to stop TPC*

```
net stop “IBM TotalStorage Productivity Center - Data Server”
net stop “IBM WebSphere Application Server V5 - Device Server”
net stop “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
```

4. Launch a DB2 command line processor as shown in Figure 3-10 on page 79 and issue the commands in the command line processor.

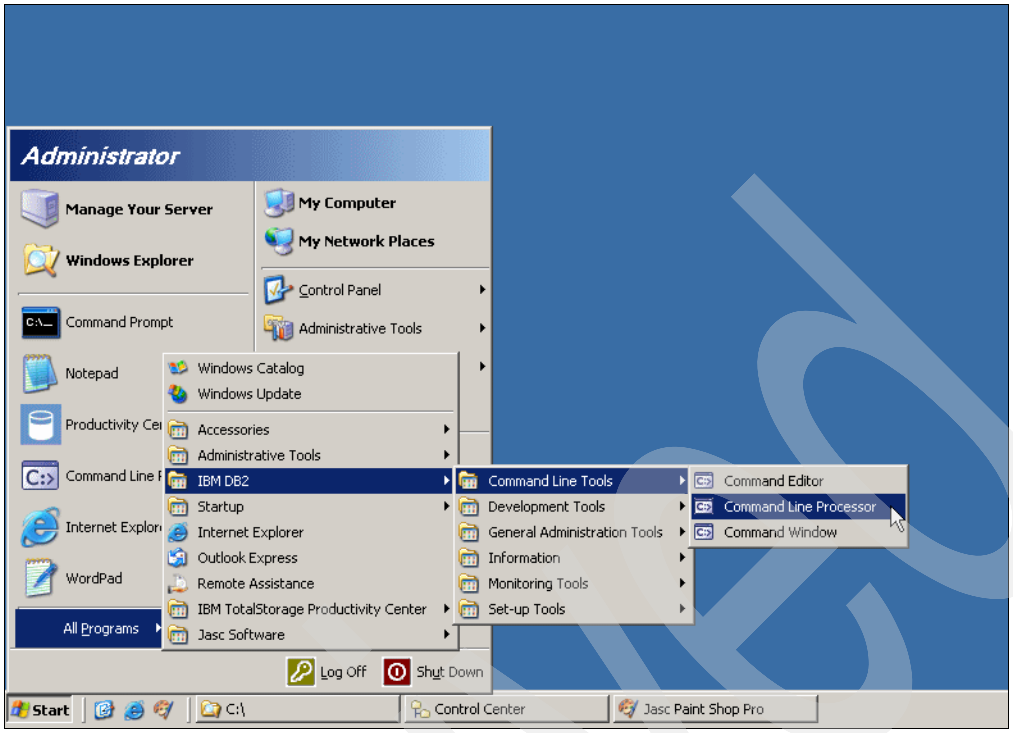

*Figure 3-10 Launch a DB2 command line processor*

5. A DB2 command line processor appears as in Figure 3-11.

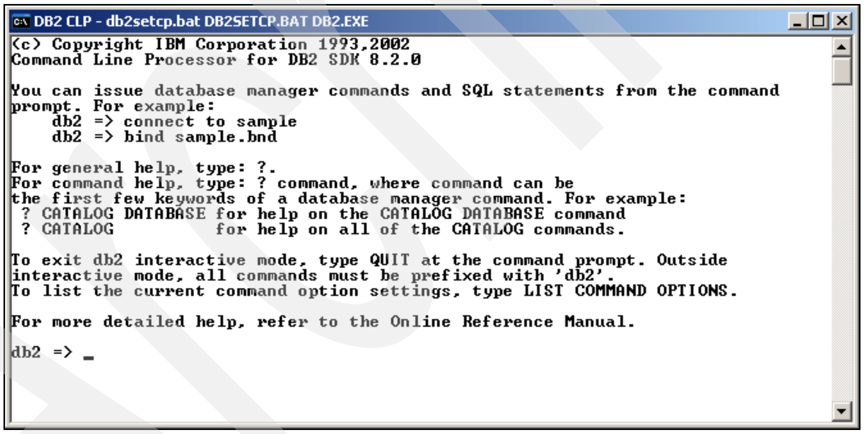

```
(c) Copyright IBM Corporation 1993,2002
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _
```

*Figure 3-11 DB2 command line processor*

6. Issue the commands from Example 3-13 on page 80 in the command line processor window. Substitute your chosen values for the parameters that form part of the **UPDATE DB CFG** command. See Table 3-4 on page 78. Note that the final two commands perform an offline backup of both databases.

> **Important:** The offline backup of both databases is required after the reconfiguration; the DB2 databases will not open until the backups are complete.

*Example 3-13   DB2 command to configure archive logging to a filesystem*

```
CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 “DISK:C:\DB2_archive_logs” failarchpath
"E:\DB2_failed_logs" logprimary 8 logsecond 16 logfilsiz 2500 newlogpath
C:\DB2_active_logs\TPCD

DISCONNECT TPCDC

CONNECT TO IBMCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR IBMCDB USING logarchmeth1 “DISK:C:\DB”_archive_logs” failarchpath
"E:\DB2_failed_logs" logprimary 8 logsecond 16 logfilsiz 2500 newlogpath
C:\DB2_active_logs\IBMCDB

DISCONNECT IBMCDB

BACKUP DATABASE TPCDB TO “E:\TPC_database_backups”
BACKUP IBMCDB IBMCDB TO “E:\TPC_database_backups
```

7. When both of the database backups complete, you can restart Agent Manager and TPC. Either use the Windows Services interface or issue the commands shown in Example 3-14 in a command prompt window.

*Example 3-14   Start Agent Manager and TPC*

```
net start “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
net start “IBM TotalStorage Productivity Center - Data Server”
net start “IBM WebSphere Application Server V5 - Device Server”
```

## 3.8.2 Create online backup script to filesystem

Create these files to control the backup process for online backup to filesystem output. We based this example on a TPC installation on Windows.

Create two files. The first file is the script that you run to start the backup.

**C:\scripts\TPC_backup_online_file.bat**

*Example 3-15 File C:\scripts\TPC_backup_online_file.bat*

```
echo on

REM This is a sample backup script
REM To backup TPC online
REM To filesystem

REM Starting backup of the DB2 databases
REM ------------------------------------

C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf C:\scripts\database_list_online_file.txt

REM Offline backup process complete
REM -------------------------------
```

The second file is the DB2 scripted list of databases to backup.

**C:\scripts\database_list_online_file.txt** - The DB2 scripted list of databases to backup.

*Example 3-16 File C:\scripts\database_list_online_file.txt*

```
backup database IBMCDB online to “E:\TPC_database_backups” without prompting;
backup database TPCDB online to “E:\TPC_database_backups” without prompting;
```

## 3.9 Performing offline database backups

This section describes how to perform the offline backup of the TPC databases.

Running an offline DB2 database backup takes TPC out of service for the period of the backup. This impacts data collections from CIMOMs, and you might miss other infrastructure events.

Make sure that you understand the impact of stopping TPC in your environment before proceeding. If your environment cannot tolerate regular stoppages of TPC for a backup operation, consider configuring and using online backup.

### 3.9.1 Performing an offline backup to a filesystem

**Important:** You must complete the initial steps as detailed in 3.4, “Common backup setup steps” on page 64 before you can start to perform offline backups.

To perform an offline backup to a filesystem, run the **`TPC_backup_offline_file.bat`** script in a command window as seen in Figure 3-12 on page 82. This script stops the TPC and Agent Manager processes, performs a backup of the two DB2 databases to the location specified in the scripts, and then restarts the services.

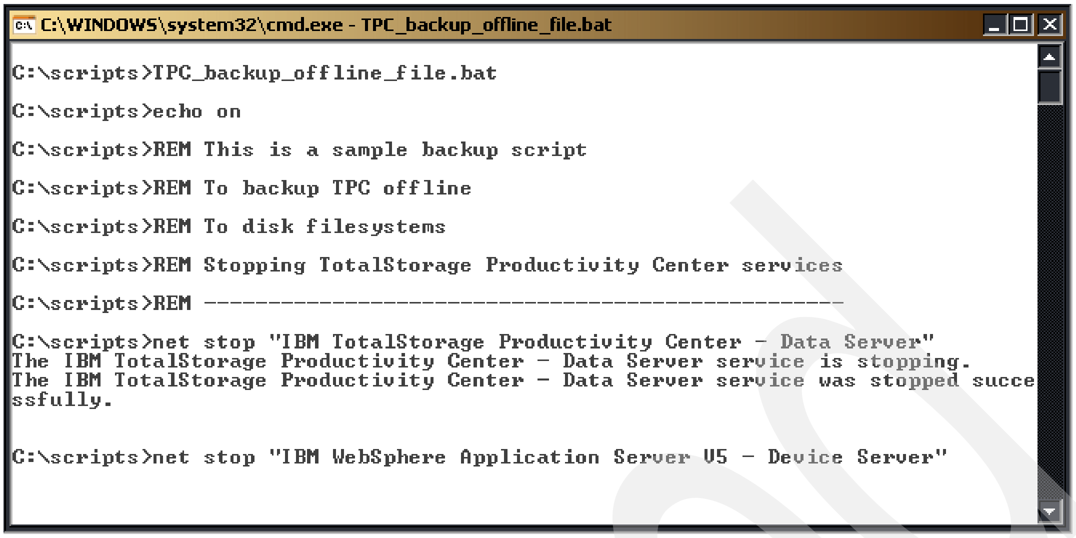

```
C:\scripts>TPC_backup_offline_file.bat

C:\scripts>echo on

C:\scripts>REM This is a sample backup script

C:\scripts>REM To backup TPC offline

C:\scripts>REM To disk filesystems

C:\scripts>REM Stopping TotalStorage Productivity Center services

C:\scripts>REM --------------------------------------------------

C:\scripts>net stop "IBM TotalStorage Productivity Center - Data Server"
The IBM TotalStorage Productivity Center - Data Server service is stopping.
The IBM TotalStorage Productivity Center - Data Server service was stopped succe
ssfully.


C:\scripts>net stop "IBM WebSphere Application Server V5 - Device Server"
```

*Figure 3-12 Running an offline backup to a filesystem*

## 3.9.2 Performing an offline backup to Tivoli Storage Manager

To perform an offline backup to Tivoli Storage Manager, run the **`TPC_backup_offline_tsm.bat`** script in a command window as seen in Figure 3-13.

**Note:** You must complete the initial setup steps that we detail in 3.6, “Offline backup to Tivoli Storage Manager setup steps” on page 67 before you can start to perform offline backups.

Running an offline DB2 database backup takes TPC out of service for the period of the backup. Make sure it is acceptable to take TPC out of service before you proceed.

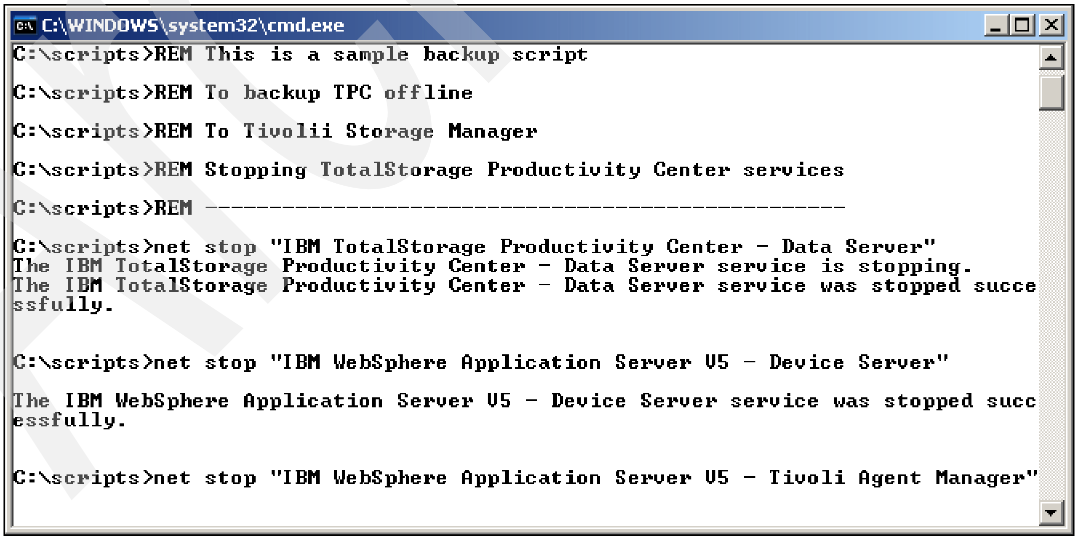

```
C:\scripts>REM This is a sample backup script

C:\scripts>REM To backup TPC offline

C:\scripts>REM To Tivolii Storage Manager

C:\scripts>REM Stopping TotalStorage Productivity Center services

C:\scripts>REM --------------------------------------------------

C:\scripts>net stop "IBM TotalStorage Productivity Center - Data Server"
The IBM TotalStorage Productivity Center - Data Server service is stopping.
The IBM TotalStorage Productivity Center - Data Server service was stopped succe
ssfully.


C:\scripts>net stop "IBM WebSphere Application Server V5 - Device Server"

The IBM WebSphere Application Server V5 - Device Server service was stopped succ
essfully.


C:\scripts>net stop "IBM WebSphere Application Server V5 - Tivoli Agent Manager"
```

*Figure 3-13 Running an offline backup to Tivoli Storage Manager*

# 3.10 Performing online database backup

This section describes how to run online backups of the TPC databases. By running the backups online (or *hot*), it is not necessary to stop the TPC services, which means that the backup process does not interrupt processes, such as performance data collection.

> **Note:** Before you can perform online DB2 backup, you must configure Tivoli Storage Manager to DB2 communication and DB2 archive logging to Tivoli Storage Manager or DB2 to filesystem logging.
>
> If you plan to use online backup to Tivoli Storage Manager, perform the setup steps in 3.6, “Offline backup to Tivoli Storage Manager setup steps” on page 67 and 3.7, “Online backup to Tivoli Storage Manager setup steps” on page 72.

## 3.10.1 Performing an online database backup to Tivoli Storage Manager

This section describes how to run an online backup of the TPC databases to Tivoli Storage Manager and assumes that you have run the appropriate setup steps.

To perform the backup, run the **`C:\scripts\TPC_backup_online_TSM.bat`** script in a command window as seen in Figure 3-14.

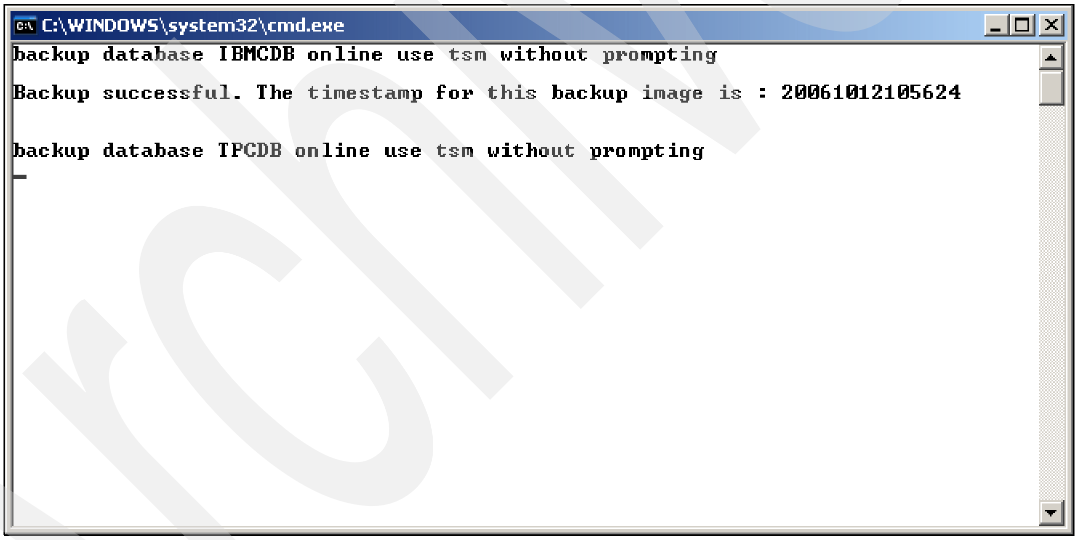

```
backup database IBMCDB online use tsm without prompting

Backup successful. The timestamp for this backup image is : 20061012105624


backup database TPCDB online use tsm without prompting
_
```

*Figure 3-14 Running an online backup to Tivoli Storage Manager*

## 3.10.2 Performing an online backup to a filesystem

This section describes how to run an online backup of the TPC databases to filesystem output files. It assumes that you have already completed the necessary setup steps detailed in 3.8, “Online backup to a filesystem setup steps” on page 77.

To perform the backup, run the **`C:\scripts\TPC_backup_online_file.bat`** script in a command window as seen in Figure 3-15 on page 84.

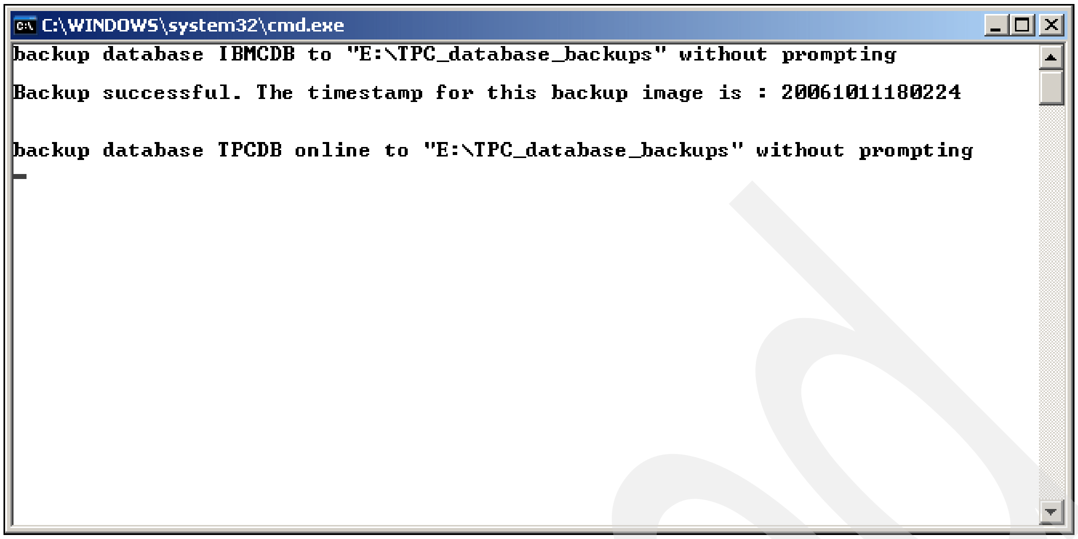

```
backup database IBMCDB to "E:\TPC_database_backups" without prompting

Backup successful. The timestamp for this backup image is : 20061011180224


backup database TPCDB online to "E:\TPC_database_backups" without prompting
_
```

Figure 3-15 Performing an online backup to filesystem output

## 3.11 Other backup considerations

Apart from the DB2 databases, there are a number of important files and directories to back up to preserve the state of a TPC server installation.

We recommend that you back up all files under the TPC and Agent Manager install directories. Additionally, the Tivoli GUID is stored in the Registry on Windows. Ensure that you also preserve this Tivoli GUID.

Important Agent Manager files to secure are:

- Agent_Manager_install_dir>/os.guid
- Agent_Manager_install_dir>/certs/CARootKeyRing.jks
- Agent_Manager_install_dir>/certs/CARootKey.pwd
- Agent_Manager_install_dir>/agentManagerKeys.jks
- Agent_Manager_install_dir>/agentManagerTrust.jks

Important TPC server directories to secure are:

- TPC_Server_install_dir/config/
- TPC_Server_install_dir/data/config/
- TPC_Server_install_dir/device/conf/

These directories contain the various configuration files for your installation. It is important to save these directories, because they might be customized configurations and not the defaults configurations.

## 3.12 Managing database backup versions

In this section, we intend to give you an overview of the maintenance processes for which you need to plan. You need to maintain the desired number of TPC database backup versions on a filesystem or Tivoli Storage Manager. DB2 does not prune older versions automatically.

### 3.12.1 Managing backup versions for a filesystem

This section describes what you need to know to manage DB2 backups that were performed to disk. DB2 does not automatically manage the deletion of the unwanted database backups or archive logs from the filesystem. You need to create a maintenance plan to delete the old backups. If you plan to use online backup to a filesystem, you need to create a plan to delete the old archive logs.

#### How does DB2 organize backups on a filesystem

When you perform a backup to a filesystem, you supply the backup script with a path to use. DB2 creates a path structure under this directory to organize the backup data. Figure 3-16 shows a backup of the TPCDB database.

The example in Figure 3-16 shows a backup that was taken on September 2nd 2006 at 15:30 and 29 seconds. DB2 timestamps all backups in this way. If multiple database backups were performed on the same day, there are multiple files in the same date directory, for example, the 20061002 directory in this case. A similar structure is created for the IBMCDB database.

Tomorrow, when a backup is made, another directory under CATN0000 is created that is made up of the date in *YYYYMMDD* format (20061003) and so on.

Plan to delete old backup directories and contents to suit the requirements of your backup and recovery policy.

> **Note:** If you plan to use another tool to copy this disk backup to tape, you can delete the entire directory from TPCDB.0 downwards each day when the tape backup is complete. DB2 will recreate it again at the next backup cycle. Be sure to restore the database backup files to the same path that they came from before you attempt a DB2 restore.

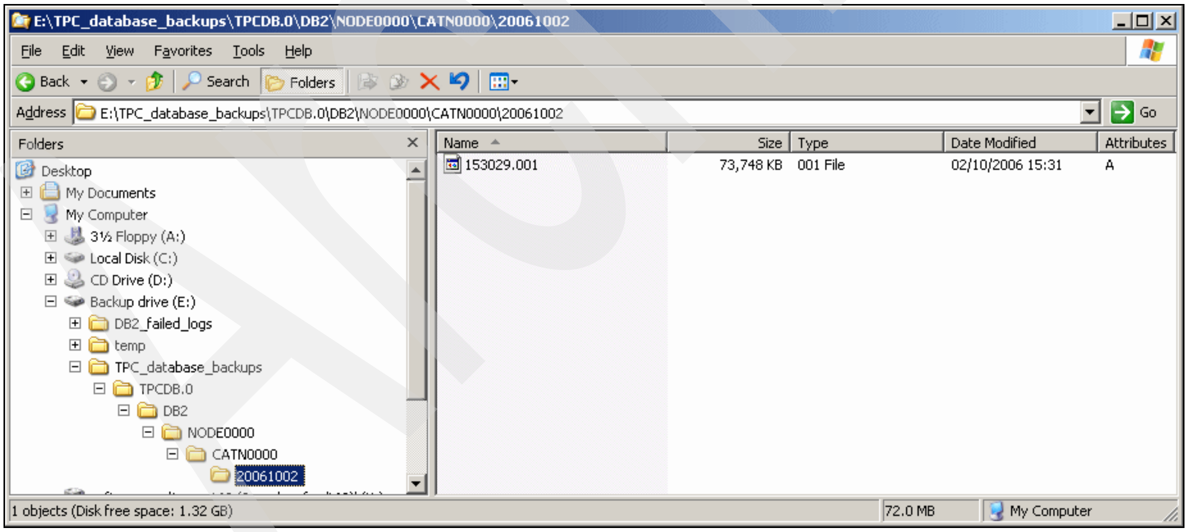

E:\TPC_database_backups\TPCDB.0\DB2\NODE0000\CATN0000\20061002

Address: E:\TPC_database_backups\TPCDB.0\DB2\NODE0000\CATN0000\20061002

| Name | Size | Type | Date Modified | Attributes |
|---|---|---|---|---|
| 153029.001 | 73,748 KB | 001 File | 02/10/2006 15:31 | A |

1 objects (Disk free space: 1.32 GB) | 72.0 MB | My Computer

Figure 3-16 DB2 backup filesystem structure

### 3.12.2 Managing archive log files on a filesystem

It is necessary to configure DB2 to use archive logging if you plan to perform online backups. If you plan to perform online backups to disk, you also need to maintain the archive logs directory on a regular basis.

Figure 3-17 shows the directory structure for logging on to the TPCDB database. Over time, this directory will fill up with logs. If your recovery policy is to keep backup versions for five days, you must keep logs in this directory for at least the same period of time, because you cannot restore an online backup without logs from the same date and time in order to make the recovery valid.

Notice that the directory that holds the logs is named *C0000000*. This is the log cycle number. If you restore the database, the cycle number increments by one and starts in *C0000001* and so on. Ensure that any automated deletion process that you implement can handle this numbering.

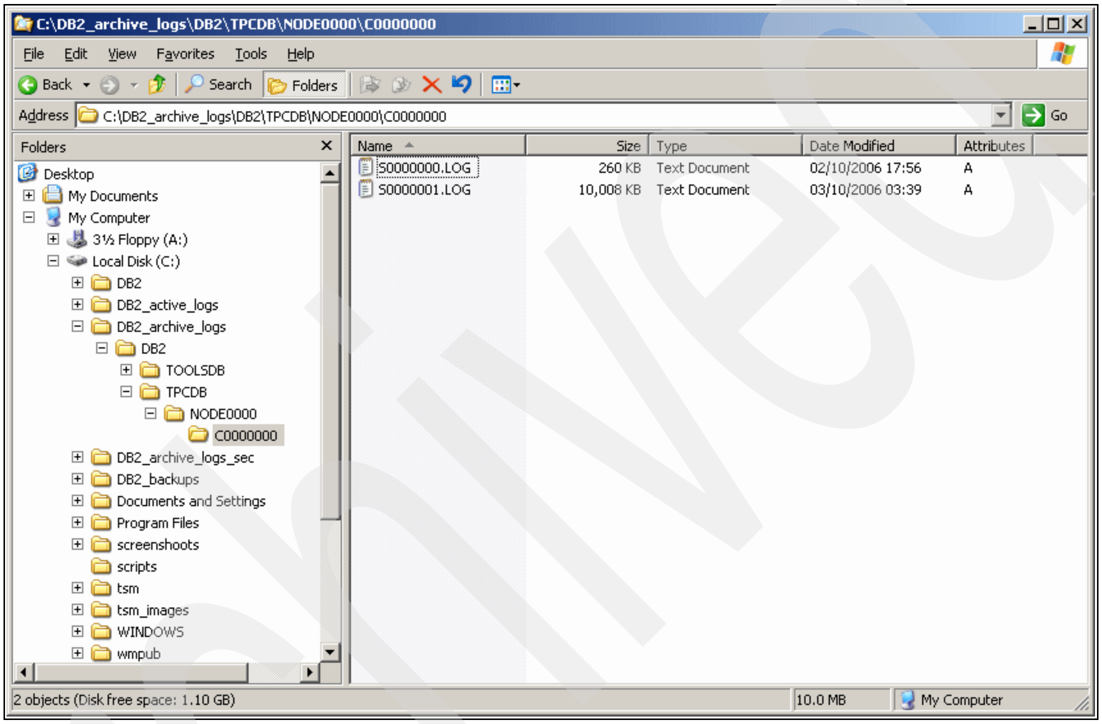

| Name | Size | Type | Date Modified | Attributes |
|---|---|---|---|---|
| S0000000.LOG | 260 KB | Text Document | 02/10/2006 17:56 | A |
| S0000001.LOG | 10,008 KB | Text Document | 03/10/2006 03:39 | A |

*Figure 3-17 Cleaning out unwanted DB2 archive logs from the filesystem*

## 3.12.3 Managing backup versions that you store in Tivoli Storage Manager

This section describes how to maintain, view, and delete backup data and archive logs that you have sent to Tivoli Storage Manager. DB2 does not automatically prune backup versions and log files from Tivoli Storage Manager. You need to use the **db2adutl** tool to perform these housekeeping functions.

> **Note:** This section is not intended to be a comprehensive guide to the **db2adutl** tool. The intent here is to detail the commands that you likely need to maintain the data that is held in Tivoli Storage Manager on a regular basis.

### What is the db2adutl command

The command line tool, **db2adutl**, communicates with Tivoli Storage Manager through its API interface. Use this tool to interrogate the backup and archive log data that is stored in Tivoli Storage Manager at any one time, verify that you no longer require old backups, and delete unnecessary old backups.

## Why do I need to use db2adutl

When DB2 stores a backup session in Tivoli Storage Manager, DB2 always stores the backup session with a unique file name, which is the timestamp when the backup was made. This means that these backup versions never get superseded by a new version with the same file name. The backup files remain “active” versions in Tivoli Storage Manager, and, therefore, Tivoli Storage Manager never deletes the backup versions. Use the command, **db2adutl**, to select unwanted backup versions and tell Tivoli Storage Manager to flag them as “inactive.” This way, Tivoli Storage Manager then deletes them over time based on the standard policy rules that the Tivoli Storage Manager administrator set.

You handle DB2 archive logs differently. They are stored in Tivoli Storage Manager as “archive” data, which means Tivoli Storage Manager retains them for a set period of time based on its policies. You can use **db2adutl** to explicitly remove DB2 archive logs, but if Tivoli Storage Manager archive retention policy is set appropriately, this is not necessary.

**Important:** Make sure that the Tivoli Storage Manager archive retention policy that you use to store the DB2 logs is set for a sufficient period of time to allow recovery of your oldest database backup. However, you also want to make sure that the policy for the retention period is not so long that it wastes storage space in Tivoli Storage Manager.

## How to query backups held in Tivoli Storage Manager

Next, we discuss how to query backups that are held in Tivoli Storage Manager.

**Note:** Database names are case sensitive in these commands. Make sure that they are uppercase. You invoke **db2adutl** from a standard Windows CMD window. You have already set the path information for this command by the DB2 installation process. This might not be true for UNIX platforms. This command is normally in the SQLLIB\bin directory of DB2.

The versions of **db2adutl** to query database backup versions are:

- This command lists all the database versions and the logs that are held for all databases stored in Tivoli Storage Manager, TPCDB, and IBMCDB, in this case.

  ```
  db2adutl query
  ```

- This command lists all database versions and logs for the TPCDB database. Note that the database name is case sensitive and is in capital letters.

  ```
  db2adutl query database TPCDB
  ```

  Figure 3-18 on page 88 shows the sample output from this command. The output shows that two database backups are stored in Tivoli Storage Manager as well as six archive logs.

- This command has a shorter output. It lists only the database backup versions and the archive logs.

  ```
  db2adutl query full
  ```

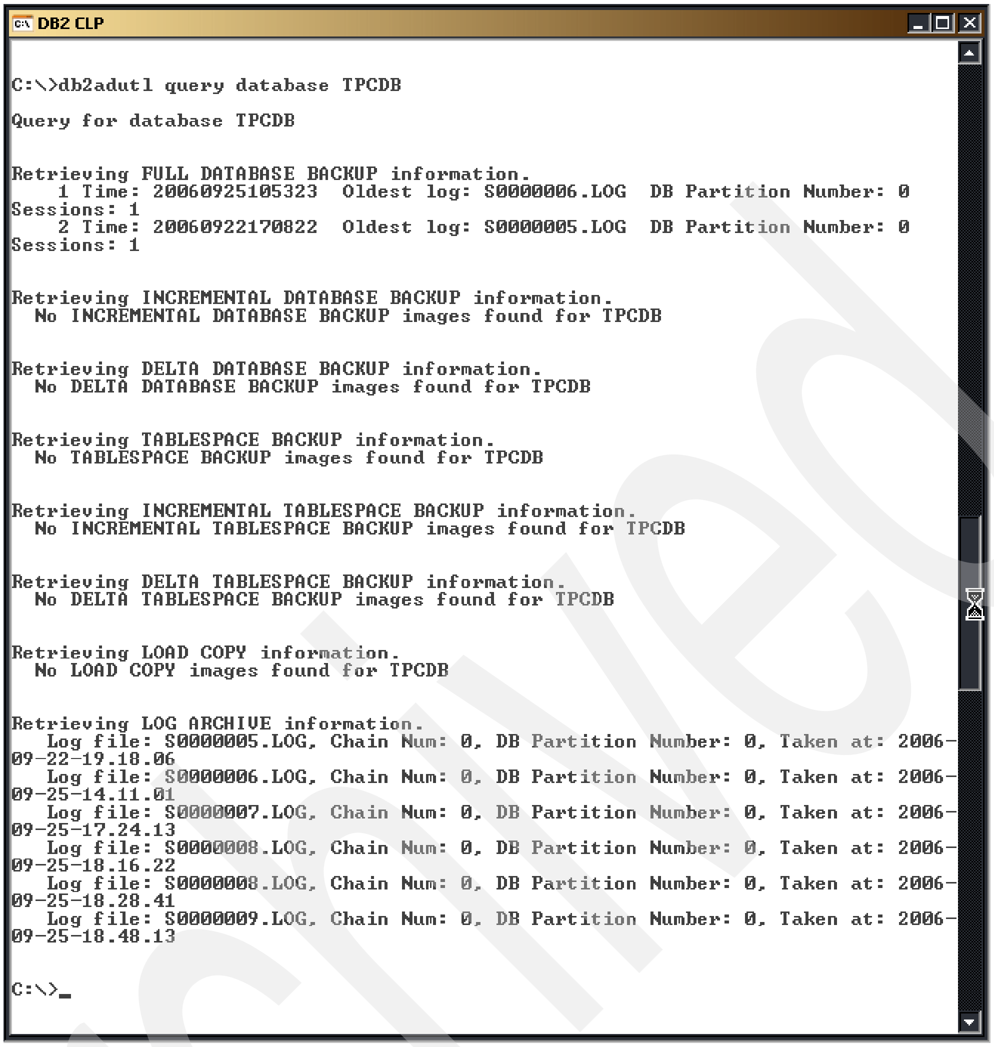

```
C:\>db2adutl query database TPCDB

Query for database TPCDB


Retrieving FULL DATABASE BACKUP information.
    1 Time: 20060925105323  Oldest log: S0000006.LOG  DB Partition Number: 0
Sessions: 1
    2 Time: 20060922170822  Oldest log: S0000005.LOG  DB Partition Number: 0
Sessions: 1


Retrieving INCREMENTAL DATABASE BACKUP information.
  No INCREMENTAL DATABASE BACKUP images found for TPCDB


Retrieving DELTA DATABASE BACKUP information.
  No DELTA DATABASE BACKUP images found for TPCDB


Retrieving TABLESPACE BACKUP information.
  No TABLESPACE BACKUP images found for TPCDB


Retrieving INCREMENTAL TABLESPACE BACKUP information.
  No INCREMENTAL TABLESPACE BACKUP images found for TPCDB


Retrieving DELTA TABLESPACE BACKUP information.
  No DELTA TABLESPACE BACKUP images found for TPCDB


Retrieving LOAD COPY information.
  No LOAD COPY images found for TPCDB


Retrieving LOG ARCHIVE information.
   Log file: S0000005.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-22-19.18.06
   Log file: S0000006.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-25-14.11.01
   Log file: S0000007.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-25-17.24.13
   Log file: S0000008.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-25-18.16.22
   Log file: S0000008.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-25-18.28.41
   Log file: S0000009.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-25-18.48.13


C:\>_
```

*Figure 3-18   Sample output from a db2adutl query database TPCDB command*

## Deleting backup versions held in Tivoli Storage Manager

The following commands and examples show how to delete database backup versions that are held in Tivoli Storage Manager:

- This command deletes backup versions from Tivoli Storage Manager that are older than three days. This type of command is useful, because you can easily script it to run each day to remove the next oldest backup.

  **`db2adutl delete full older than 3 days`**

  Or specify a database name:

  **`db2adutl delete full older than 3 days database IBMCDB`**

  Figure 3-19 on page 89 gives you an example of running this command.

- This command deletes all backup versions from Tivoli Storage Manager, except for the last three versions. Again, this command is useful when scripting an automatic process.

  **`db2adutl delete full keep 3`**

  O specify a database name:

  **`db2adutl delete full keep 3 database IBMCDB`**

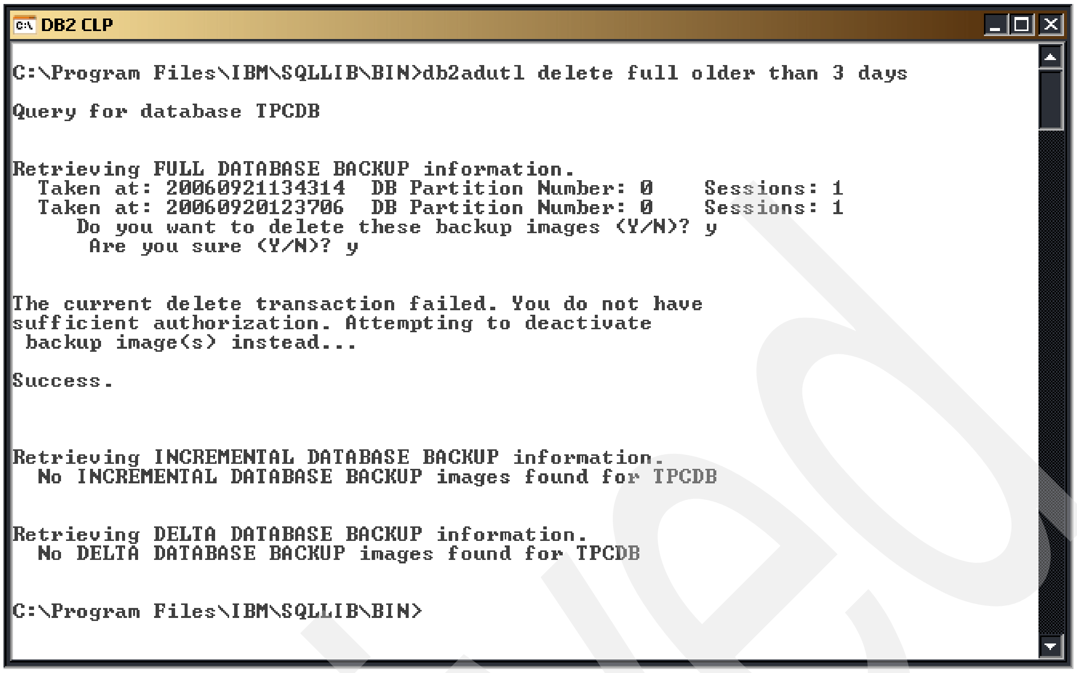

```
C:\Program Files\IBM\SQLLIB\BIN>db2adutl delete full older than 3 days

Query for database TPCDB


Retrieving FULL DATABASE BACKUP information.
  Taken at: 20060921134314  DB Partition Number: 0    Sessions: 1
  Taken at: 20060920123706  DB Partition Number: 0    Sessions: 1
     Do you want to delete these backup images (Y/N)? y
      Are you sure (Y/N)? y


The current delete transaction failed. You do not have
sufficient authorization. Attempting to deactivate
 backup image(s) instead...

Success.



Retrieving INCREMENTAL DATABASE BACKUP information.
  No INCREMENTAL DATABASE BACKUP images found for TPCDB


Retrieving DELTA DATABASE BACKUP information.
  No DELTA DATABASE BACKUP images found for TPCDB


C:\Program Files\IBM\SQLLIB\BIN>
```

*Figure 3-19 Example of a db2adutl delete full older than 3 days command*

### Managing DB2 archive log files in Tivoli Storage Manager

The following commands are an example of how to delete database archive logs from Tivoli Storage Manager.

You invoke the **db2adutl** command from a standard Windows CMD window.

> **Important:** Be careful when you delete archive log files. If you delete logs that are still needed for some of your backup versions, you render those backups *useless*.
>
> Archive logs only exist in Tivoli Storage Manager if you have configured archive logging so that online backup is possible.
>
> Ask the Tivoli Storage Manager administrator to configure Tivoli Storage Manager to delete the archive logs on a regular basis by configuring the Tivoli Storage Manager "archive copy group" that DB2 uses. Set a retention period that suits your needs. If you use a general purpose archive copy group, Tivoli Storage Manager might keep all archive logs for several years causing unnecessary usage of the storage in your Tivoli Storage Manager environment.

- To delete archive logs, first query the Tivoli Storage Manager server to establish which logs you want to delete. Figure 3-18 on page 88 shows example output.

  To query the Tivoli Storage Manager server for the TPCDB database, issue the command:

  ```
  db2adutl query database TPCDB
  ```

  Look at the "oldest log" number against the oldest backup version. In the case shown in Figure 3-18 on page 88, it is *S0000005.log*.

  Then, look at the list of log files from the same output to see if there are any earlier logs. If there are earlier logs and you do not want to wait for Tivoli Storage Manager to expire them, use the follow command to delete them. See Figure 3-20 on page 90.

  ```
  db2adutl delete logs between S0000001 and S0000004 database TPCDB
  ```

**Tip:** When specifying log numbers, you need to add the “S” at the start of the number but not the “.LOG” at the end.

Use the same process for deleting archive logs from the IBMCDB database by changing the database name in the **`db2adutl`** commands.

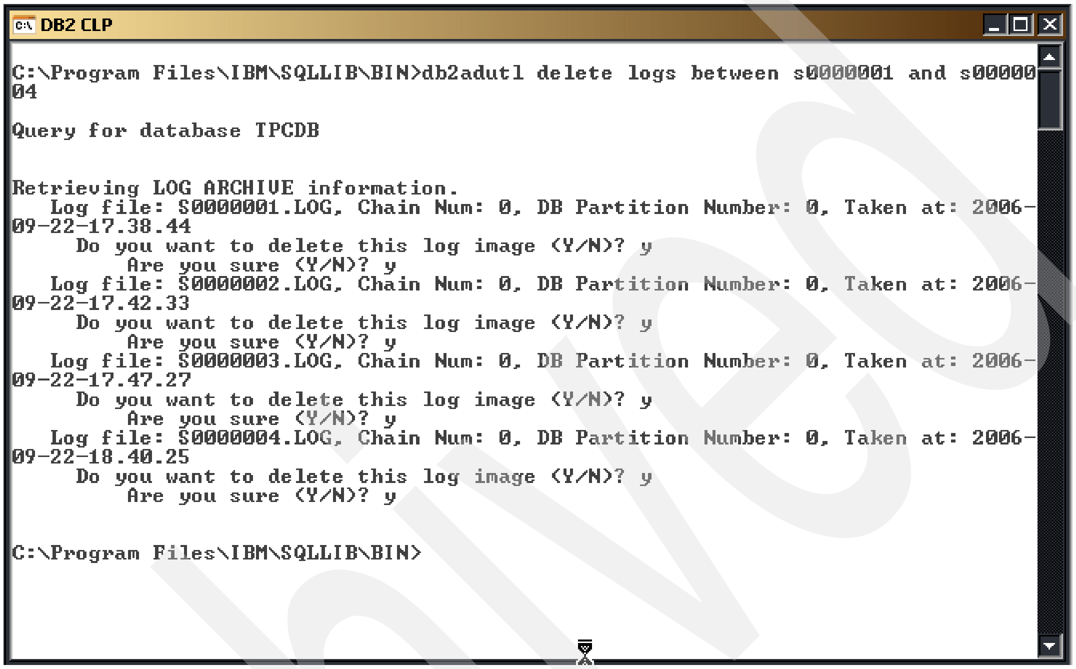

```
C:\Program Files\IBM\SQLLIB\BIN>db2adutl delete logs between s0000001 and s00000
04

Query for database TPCDB


Retrieving LOG ARCHIVE information.
   Log file: S0000001.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-22-17.38.44
     Do you want to delete this log image (Y/N)? y
         Are you sure (Y/N)? y
   Log file: S0000002.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-22-17.42.33
     Do you want to delete this log image (Y/N)? y
         Are you sure (Y/N)? y
   Log file: S0000003.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-22-17.47.27
     Do you want to delete this log image (Y/N)? y
         Are you sure (Y/N)? y
   Log file: S0000004.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2006-
09-22-18.40.25
     Do you want to delete this log image (Y/N)? y
         Are you sure (Y/N)? y


C:\Program Files\IBM\SQLLIB\BIN>
```

*Figure 3-20 Example command to delete DB2 archive logs*

## 3.13 Restoring the TPC databases

This section describes the steps necessary to restore the DB2 repository databases for TPC. As with the backup process, restoring from an online backup is more complex than restoring from an offline backup.

Restoring from an offline backup is a simple point-in-time exercise. Because the database was stopped at the time of the offline backup, it is logically consistent and you can restore the data “as is.” However, circular logging does not offer the ability to roll forward through database changes using the logs to recover to an exact point-in-time. Therefore, if you take a database backup on a 24 hour cycle, you lose updates to the TPC repository that were made between these points.

When you configure archive logging, you have the ability to restore a backup and then roll forward through the logs to any point-in-time to minimize data loss. This gives you an enhanced level of protection to the TPC repository data at the expense of more complexity in the process. You cannot simply restore a backup taken online as is, because an online backup is not logically consistent in its own right. Following an online restore, some roll forward is necessary to bring the restored database to a consistent and usable state.

Finally, we do not intend for this section to be a comprehensive guide to the DB2 restore commands. We intend to give you the basic restore functions that you need to recover a database from both filesystem and Tivoli Storage Manager backups. See *IBM DB2 Universal*

*Database™ Data Recovery and High Availability Guide and Reference*, SC09-4831, for detailed information about this subject.

## 3.13.1 Restoring from offline backups

Restoring from an offline backup is the most simple type of restore. It brings the database back to the specific point-in-time that the backup was taken. You can then restart TPC.

### Restoring an offline backup from a filesystem or Tivoli Storage Manager

This is the basic set of steps to perform a restore from an offline backup:

1. Stop the TPC services if they are still running.
2. Choose the backup image from which to restore.
3. Restore both the TPCDB and IBMCDB databases.
4. Restart the TPC services.
5. Resolve potential agent issues after you restore. For more information, see "Potential agent issues after the restore" on page 100.

#### *Stop the TPC services*

Stop the TPC services on Windows using the commands in Example 3-17.

*Example 3-17 Windows commands to stop TPC*

```
net stop “IBM TotalStorage Productivity Center - Data Server”
net stop “IBM WebSphere Application Server V5 - Device Server”
net stop “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
```

#### *Choose the backup image to restore from filesystem*

If the backup image that you require is stored in Tivoli Storage Manager, skip to the next step.

Use either Windows Explorer to look at the filesystem where you stored the backups and choose a backup image from which to restore, or use the DB2 command **`list history backup all for TPCDB`** (in a DB2 command window) to see a list of the backup versions that are available.

Figure 3-21 on page 92 shows that a backup image for the TPCDB database from the date directory 20061002 with a backup time of 153029 has been selected. This translates to 2 October 2006 at 15:30:29.

Repeat this exercise for the IBMCDB database.

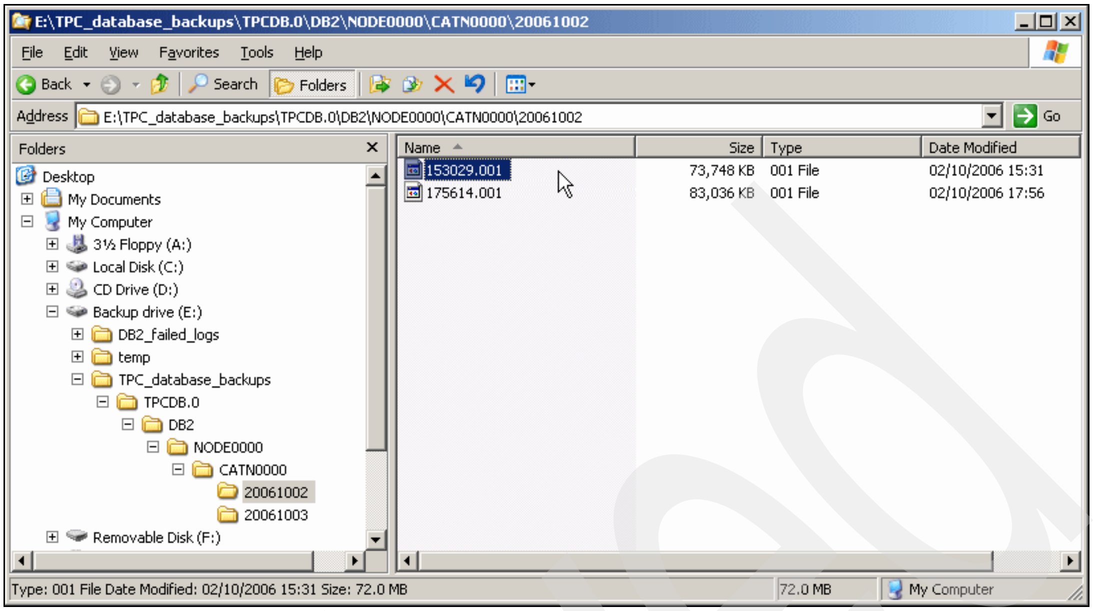

E:\TPC_database_backups\TPCDB.0\DB2\NODE0000\CATN0000\20061002

Address E:\TPC_database_backups\TPCDB.0\DB2\NODE0000\CATN0000\20061002

| Name | Size | Type | Date Modified |
|---|---|---|---|
| 153029.001 | 73,748 KB | 001 File | 02/10/2006 15:31 |
| 175614.001 | 83,036 KB | 001 File | 02/10/2006 17:56 |

Type: 001 File Date Modified: 02/10/2006 15:31 Size: 72.0 MB 72.0 MB My Computer

*Figure 3-21 Viewing backup versions available for restore*

You will have two backup image timestamps. For example:

- TPCDB database - `20061002153029`
- IBMCDB database - `20061002152837`

You need these timestamp numbers for the next step, "Restore TPCDB and IBMCDB databases (offline)" on page 93.

> **Note:** The timestamps from the two databases are not the same, because they are backed up sequentially.

### *Choose a backup image to restore from Tivoli Storage Manager*

If you have chosen a backup image from the filesystem, skip this step and move on to "Restore TPCDB and IBMCDB databases (offline)" on page 93.

To search for a backup image in Tivoli Storage Manager, use the **`db2adutl`** command:

- For the TPCDB database, issue **`db2adutl query full database TPCDB`**
- For the IBMCDB database, issue **`db2adutl query full database IBMCDB`**

Figure 3-22 on page 93 shows example output from the **`db2adutl`** command for the TPCDB database.

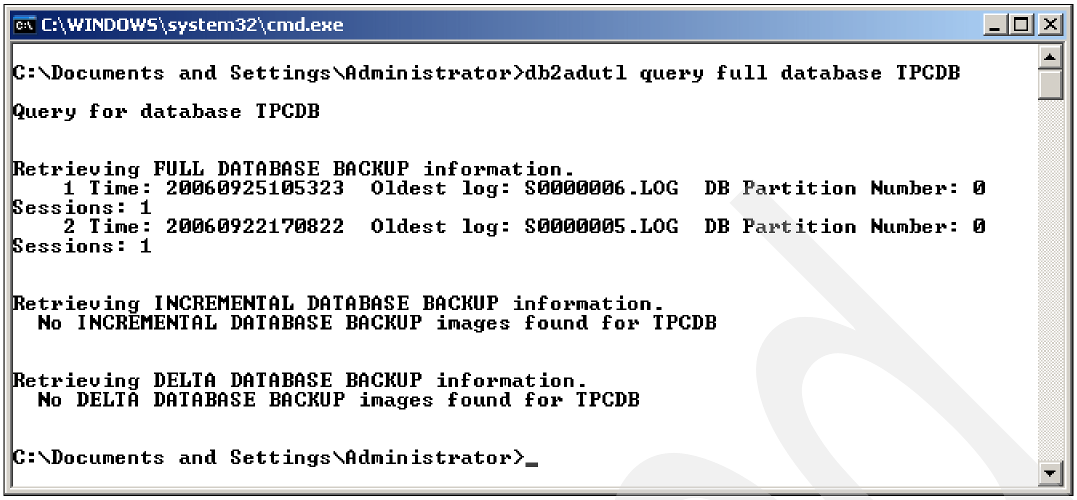

```
C:\Documents and Settings\Administrator>db2adutl query full database TPCDB

Query for database TPCDB


Retrieving FULL DATABASE BACKUP information.
    1 Time: 20060925105323  Oldest log: S0000006.LOG  DB Partition Number: 0
Sessions: 1
    2 Time: 20060922170822  Oldest log: S0000005.LOG  DB Partition Number: 0
Sessions: 1


Retrieving INCREMENTAL DATABASE BACKUP information.
  No INCREMENTAL DATABASE BACKUP images found for TPCDB


Retrieving DELTA DATABASE BACKUP information.
  No DELTA DATABASE BACKUP images found for TPCDB


C:\Documents and Settings\Administrator>_
```

*Figure 3-22 Command db2adutl example to query backup versions available*

You end up with two backup image timestamps. For example:

- TPCDB database - 20060925105323
- IBMCDB database - 20060925105100

You need these timestamp numbers for the next step.

### *Restore TPCDB and IBMCDB databases (offline)*

To restore the databases, launch a DB2 command line processor window.

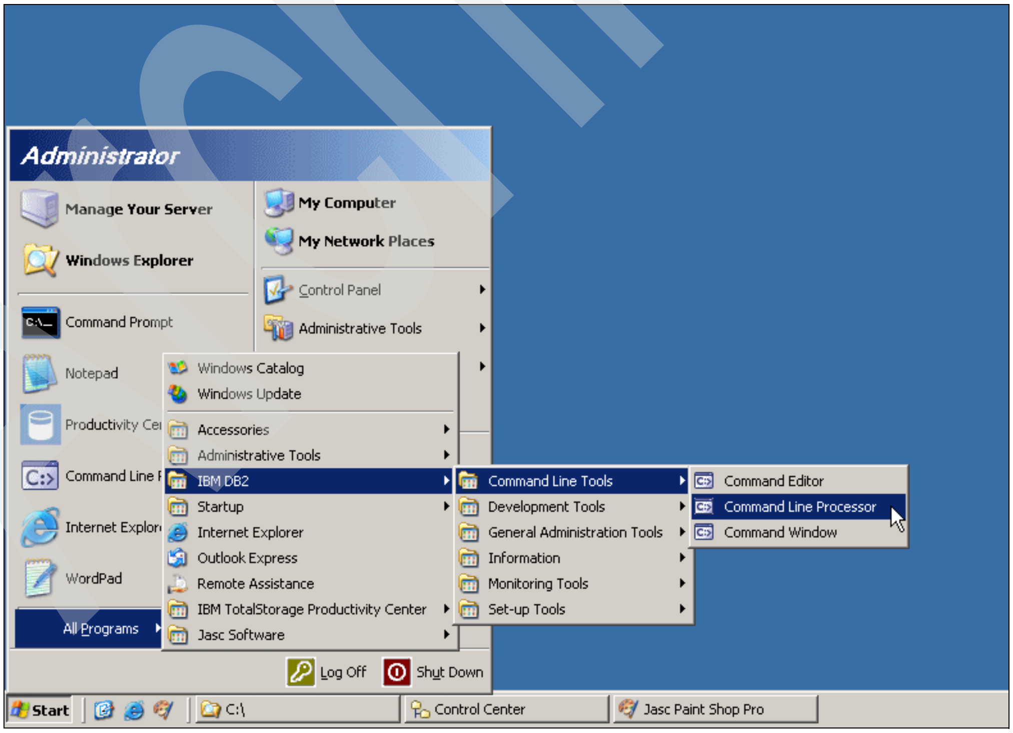

*Figure 3-23 Launch a DB2 command line processor*

A command line processor appears as in Figure 3-24 on page 94.

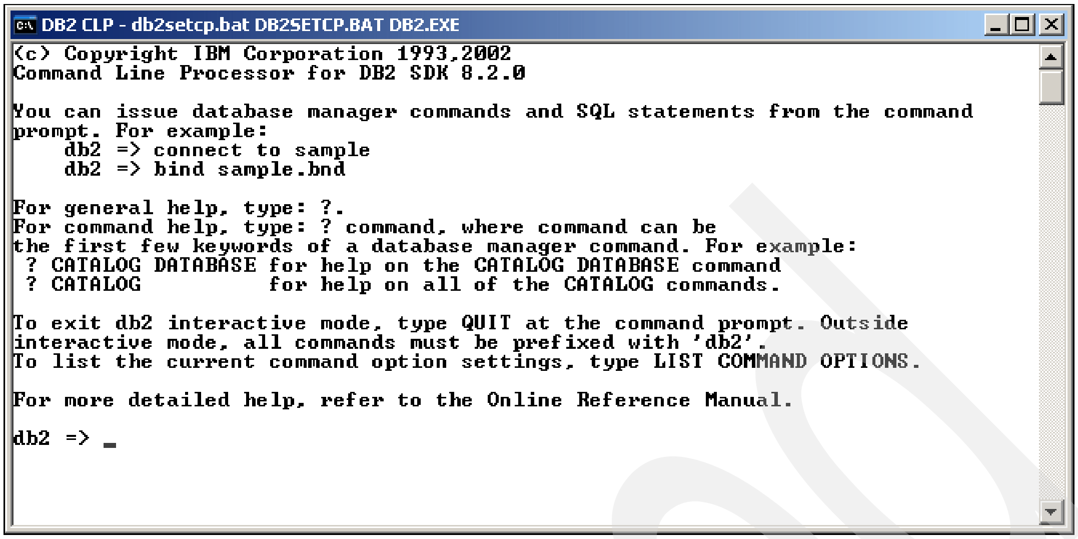

```
(c) Copyright IBM Corporation 1993,2002
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _
```

*Figure 3-24 DB2 command line processor*

To restore from filesystem backups, issue the commands in Example 3-18 in the DB2 command line processor using the timestamps that you have selected.

*Example 3-18 Restore command from filesystem backups*

```
restore database TPCDB from “E:\TPC_database_backups” taken at 20061002153029
restore database IBMCDB from “E:\TPC_database_backups” taken at 20061002152837
```

If you restore from Tivoli Storage Manager, use the commands that are shown in Example 3-19.

*Example 3-19 Restore command from Tivoli Storage Manager backups*

```
restore database TPCDB use TSM taken at 20060925105323
restore database IBMCDB use TSM taken at 20060925105100
```

Figure 3-25 on page 95 shows an example of the restore process dialog for the TPCDB database restore process from a filesystem.

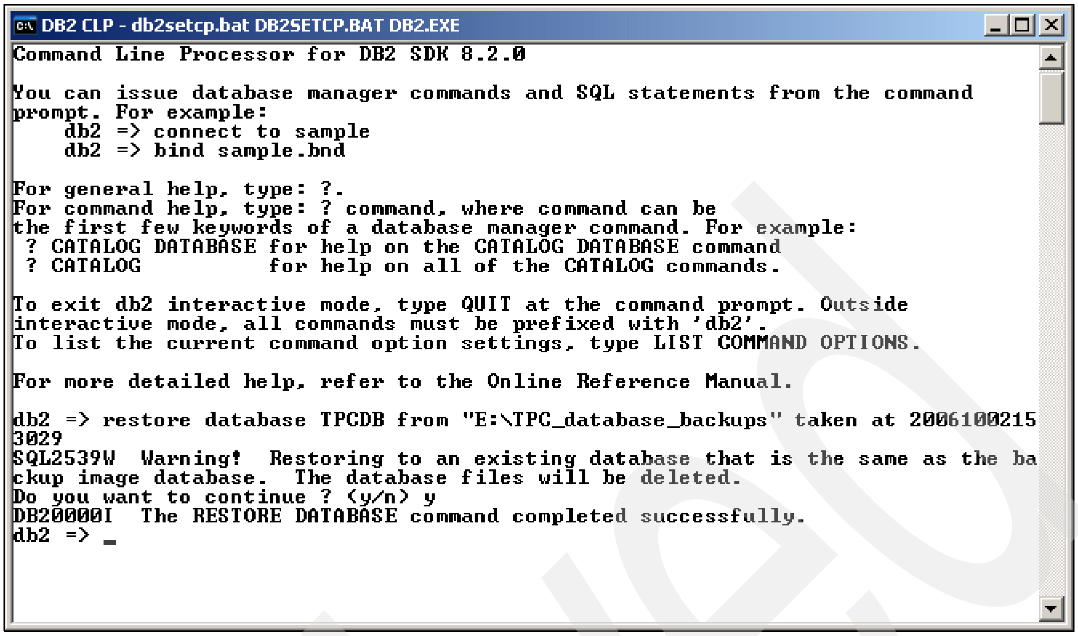

```
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => restore database TPCDB from "E:\TPC_database_backups" taken at 2006100215
3029
SQL2539W  Warning!  Restoring to an existing database that is the same as the ba
ckup image database.  The database files will be deleted.
Do you want to continue ? (y/n) y
DB20000I  The RESTORE DATABASE command completed successfully.
db2 => _
```

*Figure 3-25 Example of offline restore of TPCDB from a filesystem*

### Restart the TPC services

When you have restored the TPC and IBMCDB databases, restart the TPC services to bring the server back online. To do this on a Windows TPC server, issue the commands, which are shown in Example 3-17 on page 91, in a command window.

*Example 3-20 Example 3-17 on page 91 Windows commands to start TPC*

```
net start "IBM WebSphere Application Server V5 - Tivoli Agent Manager"
net start "IBM TotalStorage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V5 - Device Server"
```

## 3.13.2 Restoring from online backups

Restoring from an online backup can be more complex than restoring from an offline backup, because there are more choices about what to do after you restore the backup image.

You might restore to a backup image from a week ago, because you actually want your TPC environment put back to that point. You might want to restore from the last known good backup and roll forward through the archive logs to get your TPC databases as close as possible to the point before the problem occurred that triggered the need to restore.

### Restoring an online backup from a filesystem or Tivoli Storage Manager

This is the basic set of steps to perform a restore from an online backup:

1. Stop the TPC services if they are not already stopped.
2. Choose the backup image from which to restore.
3. Restore the TPCDB and IBMCDB databases.
4. Roll forward the database.
5. Restart the TPC services.
6. Resolve any new agent issues after you restore.

### Stop the TPC services

Stop the TPC services on Windows using the commands in Example 3-21. The services might already be stopped if something is broken.

*Example 3-21 Windows commands to stop TPC*

```
net stop “IBM TotalStorage Productivity Center - Data Server”
net stop “IBM WebSphere Application Server V5 - Device Server”
net stop “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
```

### Choose the backup image from which to restore

Choose a backup image from which to restore using the same process as you use for offline backups.

See “Choose the backup image to restore from filesystem” on page 91 or “Choose a backup image to restore from Tivoli Storage Manager” on page 92.

### Restore the TPCDB and IBMCDB databases (online)

The initial process of restoring a database that was taken online is the same as the offline process. However, when you complete the restore, you are *not* ready to use the database. After you restore the backup, the database status is “Roll-Forward Pending.” The next section explains how to proceed from this point.

To restore the databases, launch a DB2 command line processor window as seen in Figure 3-26.

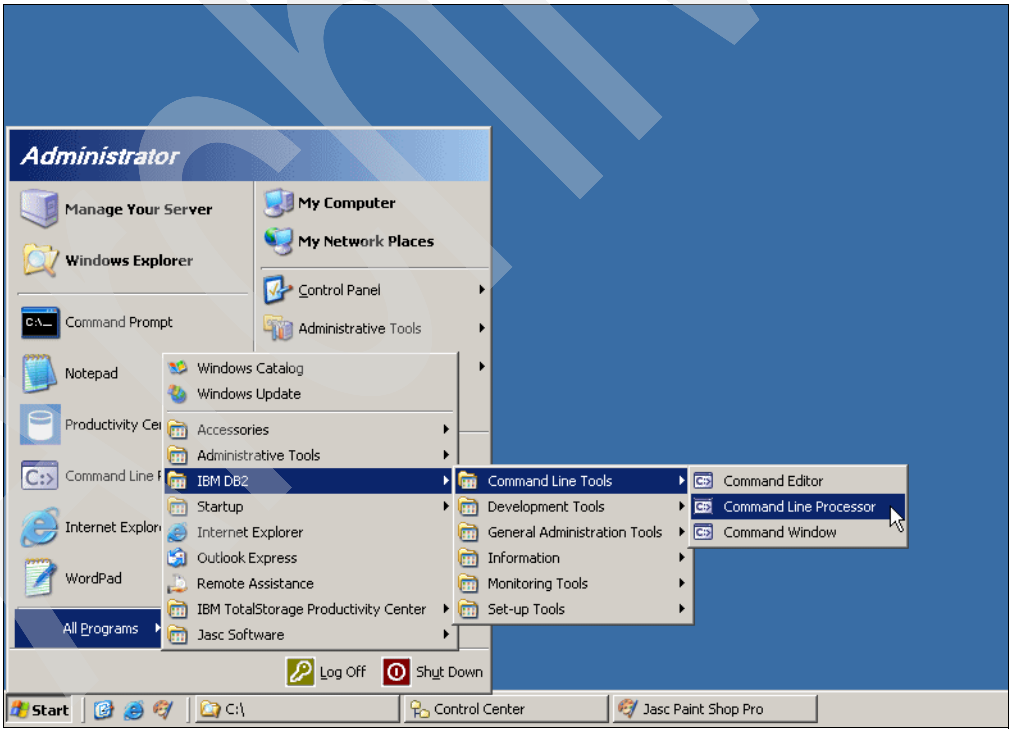

*Figure 3-26 Launch a DB2 command line processor*

A command line processor appears as in Figure 3-27 on page 97.

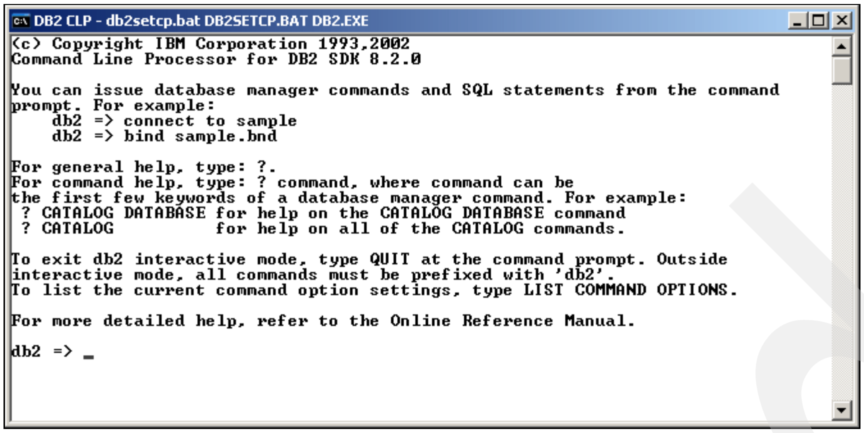

```
(c) Copyright IBM Corporation 1993,2002
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _
```

*Figure 3-27 DB2 command line processor*

To restore the databases from filesystem backups, issue the commands in Example 3-22 in the DB2 command line processor using the timestamps that you have selected.

*Example 3-22 Restore command from filesystem backups*

```
restore database TPCDB from “E:\TPC_database_backups” taken at 20061002153029
restore database IBMCDB from “E:\TPC_database_backups” taken at 20061002152837
```

If you restore from Tivoli Storage Manager, use different commands as in Example 3-23.

*Example 3-23 Restore command from Tivoli Storage Manager backups*

```
restore database TPCDB use TSM taken at 20060925105323
restore database IBMCDB use TSM taken at 20060925105100
```

Figure 3-28 on page 98 shows an example of the restore process dialog for the TPCDB database restore from a filesystem. Repeat this process for the IBMCDB database after the TPCDB database restore process completes.

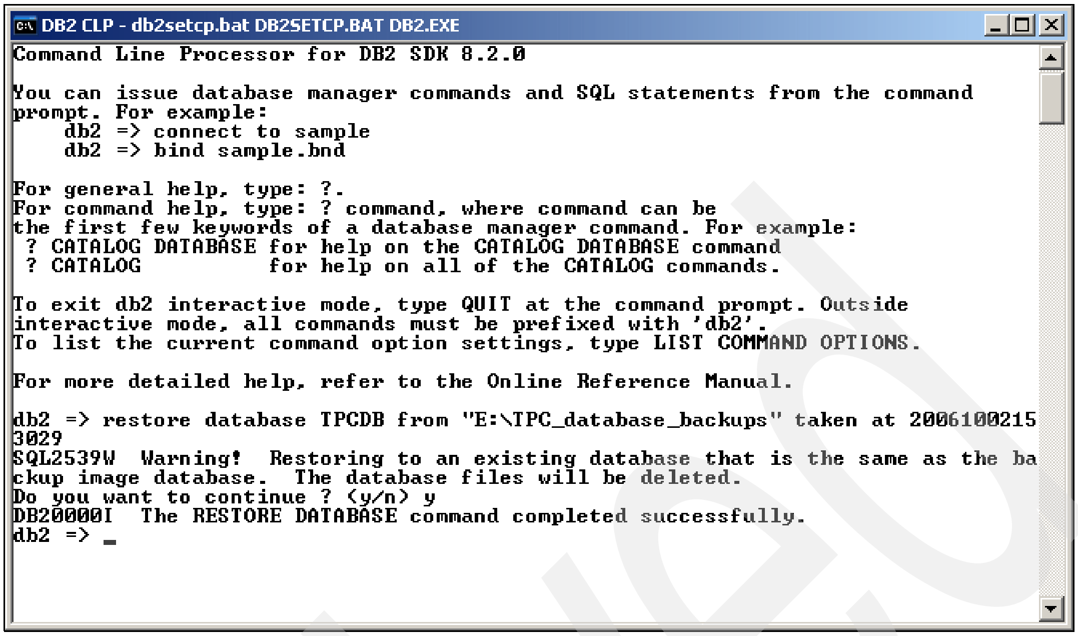

```
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => restore database TPCDB from "E:\TPC_database_backups" taken at 2006100215
3029
SQL2539W  Warning!  Restoring to an existing database that is the same as the ba
ckup image database.  The database files will be deleted.
Do you want to continue ? (y/n) y
DB20000I  The RESTORE DATABASE command completed successfully.
db2 => _
```

*Figure 3-28   Example of offline restore of TPCDB from a filesystem*

Perform the restore operation for the IBMCDB using the same method.

### Roll forward the databases

After the database restore processes complete, you can start the roll forward. You cannot start TPC at this point, because the databases will not open until you perform some type of roll forward.

Roll forward options in DB2 can be complex. We do not intend to provide a complete guide to DB2 roll forward recovery.

We describe how to roll forward in two ways:

- Roll forward to the end of the logs

  This rolls forward from the restore point through all available log files to the most recent consistent point-in-time. If you are using an old backup and there are many logs through which to roll, this method can take some time.

- Roll forward to a point-in-time

  With a point-in-time roll forward, you can specify a specific point-in-time for the roll forward process to stop, complete, and allow the database to open.

### Roll databases forward to the end of the logs

To roll the database forward to the end of all of the logs after a restore, type the following commands in the DB2 command line processor as seen in Figure 3-29 on page 99. When each command completes, it returns an audit of the process.

> **Note:** The *last committed transaction time* is displayed in a UTC-0 time format even if your local time zone is, for example, PDT (UTC-8). Perform this command for both databases.

Commands:

- rollforward database TPC to end of logs and complete
- rollforward database IBMCDB to end of logs and complete

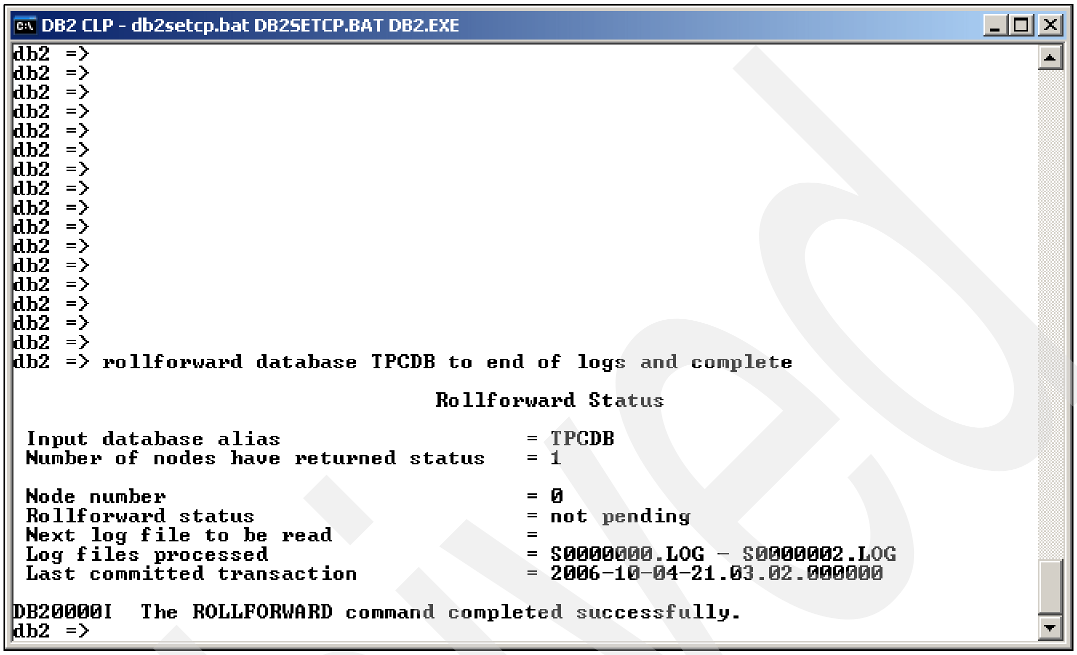

```
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 => rollforward database TPCDB to end of logs and complete

                                 Rollforward Status

 Input database alias                   = TPCDB
 Number of nodes have returned status   = 1

 Node number                            = 0
 Rollforward status                     = not pending
 Next log file to be read               =
 Log files processed                    = S0000000.LOG - S0000002.LOG
 Last committed transaction             = 2006-10-04-21.03.02.000000

DB20000I  The ROLLFORWARD command completed successfully.
db2 =>
```

*Figure 3-29   Roll forward TPCDB to the end of the logs and complete*

When complete, proceed to “Restart the TPC services” on page 100.

### *Roll forward databases to a point-in-time*

Here are the commands to roll the databases forward to a given point-in-time after the restore.

> **Note:** By default, DB2 uses UTC-0 time for the point-in-time roll forward. Add the **use local time** flag to the command if you want to specify a time in your local time zone.

Follow these steps:

1. Use the DB2 command line processor as seen in Figure 3-30 on page 100 to enter the `rollforward` command. In this example, we rolled the TPCDB database forward to a few minutes after the restore time. We entered the time using the **use local time** option.
2. Enter the point-in-time as *YYYY-MM-DD-HH.MM.SS.*

   The command for the TPCDB database is:

   ```
   rollforward database TPCDB to 2006-10-04-13.36 using local time and complete
   ```

   The command for the IBMCDB:

   ```
   rollforward database IBMCDB to 2006-10-04-13.36 using local time and complete
   ```

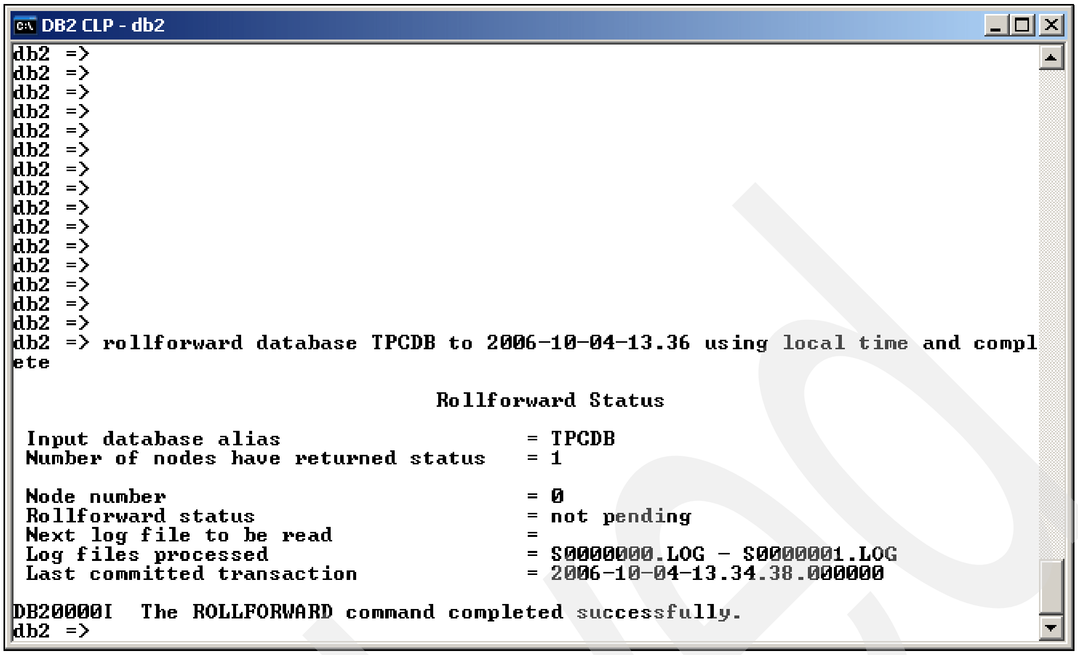

```
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 =>
db2 => rollforward database TPCDB to 2006-10-04-13.36 using local time and compl
ete

                                 Rollforward Status

 Input database alias                   = TPCDB
 Number of nodes have returned status   = 1

 Node number                            = 0
 Rollforward status                     = not pending
 Next log file to be read               =
 Log files processed                    = S0000000.LOG - S0000001.LOG
 Last committed transaction             = 2006-10-04-13.34.38.000000

DB20000I  The ROLLFORWARD command completed successfully.
db2 =>
```

*Figure 3-30 Roll forward the TPCDB to point-in-time and complete*

Notice that the actual *last committed transaction* time is slightly different than the time that is requested in the roll forward. This is the closest that DB2 can get to the requested time and still keep the database in a consistent state.

### *Restart the TPC services*

After you complete the restore operation and the roll forward for the TPC and IBMCDB databases, restart the TPC services to bring the server back online. To do this on a Windows TPC server, issue the commands shown in Example 3-24 in a command window.

*Example 3-24 Windows commands to start TPC*

```
net start “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
net start “IBM TotalStorage Productivity Center - Data Server”
net start “IBM WebSphere Application Server V5 - Device Server”
```

## 3.13.3 Potential agent issues after the restore

Following the restore of the TPC databases to a previous state, there is always a possibility that new agents were deployed to machines after the time of the restore. If this happens, there are agents, which are not registered in the TPC server and which are now running with a older version of the database, out in the environment. To correct this problem, you need to instruct the orphaned agents to re-register themselves with the TPC server.

The steps for this process are:

1. Create a PROBE_ME file in the TPC for Data top level directory of the agent.
2. Restart the agent.
3. If successful, the agent deletes the PROBE_ME file to acknowledge the action.

### Create a PROBE_ME file

Locate the root level directory for the TPC for Data agent on the affected machines.

For example:

- C:\Profile Files\IBM\TPC\ca\subagents\TPC\Data - On Windows
- /usr/tivoli/ep/subagents/TPC/Data - On AIX
- /opt/tivoli/ep/subagents/TPC/Data - Linux® or UNIX

The actual location depends on your choice of installation directory, but the final part of the path will end in \subagents\TPC\Data. Locate this directory on your machine and create an empty file named `PROBE_ME` that is upper case with no other file extension.

### Restart the agent

With the PROBE_ME file in place, stop the agent that is on the agent machine.

- For Windows, restart the service by locating the "IBM Tivoli Common Agent" service and restarting it.
- For AIX and UNIX:
  - **`cd`** ***`<agent_install_dir>`*** /usr/tivoli/ep, for example
  - **`./endpoint.sh restart`**

### Check PROBE_ME file

If the restart is successful, the agent removes the PROBE_ME file when the agent restarts. This process can take up to 30 seconds to complete.

The agent is now registered with TPC and you can add the agent to scan and probe jobs.

## 3.14 TPC repository database sizing

This section is a guide to estimating the storage requirements for the TPC repository database. This section is not designed to be an exact tool; instead, it focuses on the primary data types that consume the majority of space within the TPC repository database.

You can break down the sizing information into three sections:

- Storage subsystem performance collection data:
  - IBM DS*xxxx*, ESS, and non-IBM subsystems
  - IBM SAN Volume Controller (SVC) systems
- SAN Fabric performance collection data
- TPC for Data analysis collection records

TPC collects performance data from each device at timed intervals. For instance, you might decide to collect performance information from your production DS8000™ at 5-minute intervals. This consumes much more repository storage than if you choose to sample your production DS8000 at 30-minute intervals.

Not all devices support the same level of performance monitoring granularity. For instance, the SVC has a minimum sample rate of 15-minute intervals.

**Important:** If you plan to collect performance data from a large number of devices at small time intervals, the infrastructure that you put in place needs to be able to support it. The CIMOMs that collect the data will need to work harder, and the TPC repository database will need to support a higher number of transactions to keep pace.

You might be required to perform additional tuning of the DB2 database as monitored devices increase or if you plan to use a high sample rate.

As performance data arrives at TPC, it is stored in the TPC repository database. Over time, TPC creates hourly and daily averages of this data. The averaged data requires less storage space in the repository over a longer period of time. It also makes reporting over a longer time period more meaningful and easier to display.

For this scenario, you plot the performance of a subsystem for the last 30 days. If you did so using 15-minute samples, there are 2,800 plot points, making the graph both jagged and difficult to plot.

Because TPC averages the 15-minute samples into both hourly and daily samples, known as *aggregates,* you can choose to plot the same 30-day period with only 30 or 720 points on the graph, making it much less jagged and more legible.

The process of collecting the individual samples into hourly and daily samples is called *History Aggregation*, and TPC has a configuration panel that controls how much history you keep over time.

**Important:** The history aggregation process is a global setting, which means that the values set for history retention are applied to *all* performance data from *all* devices. You cannot set history retention on an individual device basis.

Figure 3-31 on page 103 shows the TPC panel for setting the history retention for performance monitors as well as other types of collected statistics.

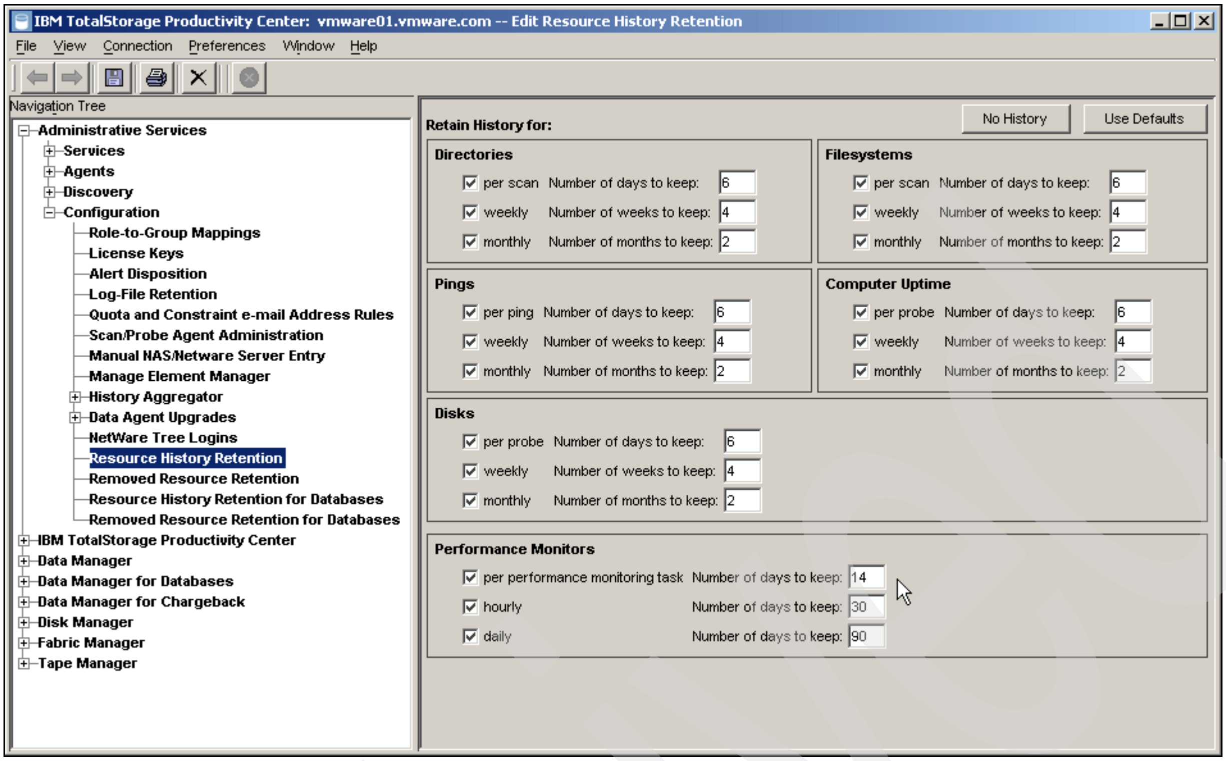
IBM TotalStorage Productivity Center: vmware01.vmware.com -- Edit Resource History Retention
File View Connection Preferences Window Help
Navigation Tree
Administrative Services
Services
Agents
Discovery
Configuration
Role-to-Group Mappings
License Keys
Alert Disposition
Log-File Retention
Quota and Constraint e-mail Address Rules
Scan/Probe Agent Administration
Manual NAS/Netware Server Entry
Manage Element Manager
History Aggregator
Data Agent Upgrades
NetWare Tree Logins
Resource History Retention
Removed Resource Retention
Resource History Retention for Databases
Removed Resource Retention for Databases
IBM TotalStorage Productivity Center
Data Manager
Data Manager for Databases
Data Manager for Chargeback
Disk Manager
Fabric Manager
Tape Manager
No History
Use Defaults

Retain History for:

Directories
- ☑ per scan Number of days to keep: 6
- ☑ weekly Number of weeks to keep: 4
- ☑ monthly Number of months to keep: 2

Filesystems
- ☑ per scan Number of days to keep: 6
- ☑ weekly Number of weeks to keep: 4
- ☑ monthly Number of months to keep: 2

Pings
- ☑ per ping Number of days to keep: 6
- ☑ weekly Number of weeks to keep: 4
- ☑ monthly Number of months to keep: 2

Computer Uptime
- ☑ per probe Number of days to keep: 6
- ☑ weekly Number of weeks to keep: 4
- ☑ monthly Number of months to keep: 2

Disks
- ☑ per probe Number of days to keep: 6
- ☑ weekly Number of weeks to keep: 4
- ☑ monthly Number of months to keep: 2

Performance Monitors
- ☑ per performance monitoring task Number of days to keep: 14
- ☑ hourly Number of days to keep: 30
- ☑ daily Number of days to keep: 90

*Figure 3-31 Setting the history retention for performance monitors*

The **Performance Monitors** values in Figure 3-31are:

- **Per performance monitoring task**

  The value that you set here defines the number of days that TPC keeps individual data samples for all of the devices that send performance data. The example shows 14 days. When per sample data reaches this age, TPC permanently deletes it from the database.

  Increasing this value allows you to look back at device performance at the most granular level at the expense of consuming more storage space in the TPC repository database.

  Data held at this level is good for plotting performance over a small time period but not for plotting data over many days or weeks because of the number of data points. Consider keeping more data in the hourly and daily sections for longer time period reports.

  Checking this field determines whether history retention is `on` or `off`. If you remove the check, TPC does not keep any history for “per sample data.”

- **Hourly**

  This value defines the number of days that TPC holds performance data that has been grouped into hourly averages. Hourly average data potentially consumes less space in the database. For example, if you collect performance data from an SVC at 15-minute intervals, retaining the hourly averages requires four times less space in the database.

  The check box determines whether history retention is `on` or `off`. If you remove the check, TPC does not keep any history for hourly data.

- **Daily**

  This value defines the number of days that TPC holds performance data that has been grouped into daily averages. After the defined number of days, TPC permanently deletes records of the daily history from the repository.

Daily averaged data requires 24 times less space in the data for storage compared to hourly data. This is at the expense of granularity; however, plotting performance over a longer period (perhaps weeks or months) becomes more meaningful.

The check box determines whether history retention is `on` or `off`. If you remove the check, TPC does not keep any history for daily data.

## 3.14.1 Storage subsystem performance data sizing

There is a significant difference in the sizing calculation between the SVC and other subsystems, both IBM and non-IBM. For this reason, the sizing tables are separated.

### Sizing the repository for ESS, DS*xxxx,* and non-IBM subsystems

You can use the example worksheet in Table 3-6 to get an understanding of the likely storage requirements that are needed for the repository to hold a given amount of data.

Table 3-6 shows working examples for four storage subsystems in an environment and the amount of storage space that performance collection uses for each example. The total figure represents the amount of storage needed for the "per sample" data. Continue through this section to calculate the complete amount of storage needed for hourly and daily history types.

Calculation method example for **ESS_Production**:

60/5 x 24 = **288** samples per day x **1,500** volumes x **200** bytes per sample = **86,400,000** bytes

*Table 3-6 Per sample repository database sizing for ESS, DSxxxx, and non-IBM subsystems*

| (a) Subsystem name | (b) Number of volumes (LUNs) sampled | (c) Performance collection interval (minutes) | (d) Performance data record size | (e) Daily amount of data collected<br>(60/(c) x 24) x (b) x (d) = (e) |
|---|---|---|---|---|
| ESS_Production | 1,500 | 5 | 200 bytes | 86,400,000 |
| DS8100_Live | 1,500 | 15 | 200 bytes | 28,800,000 |
| DS4700_Jason | 500 | 30 | 200 bytes | 4,800,000 |
| EMC_remote | 1,000 | 30 | 200 bytes | 9,600,000 |
| | | | | |
| | | | | |
| | | | | |
| | | (f) Total required per day | | 129,600,000 |
| | | (g) Number of days to retain per sample = 14 days<br>(f) x (g)/1,024,000 + 50% | | **2,658 MB** |

**Note:** Notice that the final figure includes an additional 50%. This amount provides for DB2 table indexes and other database overhead.

You can see that the amount of space that is required increases dramatically as the sample rate increases. Remember this when you plan the appropriate sample rate for your environment.

Next, use Table 3-7 to calculate the amount of storage that is needed to hold the performance data for the hourly and daily history averages. When complete, add together the totals from Table 3-6 on page 104 and Table 3-7 to give you the total repository requirement for these types of storage subsystems as seen in Table 3-8.

Calculation method example for **ESS_Production**:

**1,500** volumes x **200** bytes per sample x 24 = **7,200,000** bytes for hourly history average

*Table 3-7 Hourly and daily repository database sizing for ESS, DSxxxx, and non-IBM storage*

| (a) Subsystem name | (b) Number of volumes sampled (LUNs) | (c) Performance data record size (bytes) | (d) Hourly requirement<br>(b) x (c) x 24 | (e) Daily requirement<br>(b) x (c) |
|---|---|---|---|---|
| ESS_Production | 1,500 | 200 | 7,200,000 | 300,000 |
| DS8100_Live | 1,500 | 200 | 7,200,000 | 300,000 |
| DS4700_Jason | 500 | 200 | 2,400,000 | 100,000 |
| EMC_remote | 1,000 | 200 | 4,800,000 | 200,000 |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | **Daily totals** | 21,600,000 | 900,000 |
| | | Hourly days x 30 | (f) 648,000,000 | |
| | | Daily days = 90 | | (g) 81,000,000 |
| | | | **Total MB**<br>(f) + (g)/<br>1,024,000 + 50% | **1,068 MB** |

Table 3-8 shows the total TPC repository space required for ESS, DS*xxxx*, and non-IBM storage subsystems. The total TPC repository space is the sum of the totals of both Table 3-6 on page 104 and Table 3-7.

*Table 3-8 Total TPC repository space required for ESS, DSxxxx, and non-IBM subsystems*

| Total space required MB |
|---|
| 2,658 |
| 1,068 |
| **3,726 MB** |

## Sizing the repository for SVC (TPC Version 3.1.3 and above)

Use this section to size the repository when you use TPC Version 3.1.3 and above. Starting with this version, TPC collects a larger number of performance metrics from SVC devices in order to take advantage of SVC V4.1 software, which can now report on a larger number of performance metrics.

The repository database tables have increased in size significantly to enable TPC to store these additional performance metrics. If you run at TPC Version 3.1.2 or earlier, the next section gives you the details for sizing at that version.

> **Important:** Upgrading from TPC 3.1.2 to TPC 3.1.3 or later causes the database schema to change for SVC performance data. The tables in this section help you plan for this increase in size.

Complete Table 3-9 for the each SVC that you will monitor in the environment. The table assumes a fixed sample rate is chosen for all SVCs. If you plan to monitor some SVCs at 15-minute intervals and other SVCs at 30-minute intervals, you need to fill out this table twice: one table for each chosen sample rate, and then, add the two tables together to give you an overall total.

*Table 3-9 Repository sizing for SVC and TPC Version 3.1.3 and above*

| **Subsystem** | **Number of VDisks** | **Number of MDisks** | **I/O groups** | **MDisk groups** | **Cluster pairs** |
|---|---|---|---|---|---|
| TEC_SVC | 900 | 500 | 1 | 4 | 1 |
| SVC_Jason | 3,000 | 1,500 | 2 | 6 | 2 |
| | | | | | |
| | | | | | |
| | | | | | |
| **Totals** | 3,900 | 2,000 | 3 | 10 | 3 |
| Record size (bytes) | 198 | 78 | 500 | 128 | 492 |
| Byte totals | 772,200 | 156,000 | 1,500 | 1,280 | 1,476 |
| | | | (a) Sample rate (bytes) | | 932,456 |
| | | | (b) Hourly amount @ 15-minute sample rate<br>(60/15) x (a) | | 3,729,824 |
| | | | (c) Daily amount<br>(b) x 24 | | 89,515,776 |
| | | | (d) 14-day retention of samples<br>(b) x 24 x 14 | | 1,253,220,864 |
| | | | (e) 30-day retention of hourly<br>24 x (a) x 30 | | 671,368,320 |
| | | | (f) 90-day retention of daily<br>(a) x 90 | | 83,921,040 |
| | | | (g) Overall total required (MB)<br>(d) + (e) + (f)/1,024,000 + 50% | | **2,942 MB** |

**Important:** Notice that the overall figure in (g) adds 50% to the amounts calculated through the table. The majority of this overhead takes the DB2 table indexes for this data plus database page overhead into account.

## Sizing the repository for SVC - TPC Version 3.1.2 and below

TPC Version 3.1.2 and below were enabled to work with SVC V3 software. SVC V3 reported a smaller number of performance metrics through the Storage Management Initiative Specification (SMI-S) interface. For this reason, the performance tables in the TPC repository database are much smaller than those of TPC V3.1.3 and above. TPC V3.1.3 and above have been engineered to work with SVC V4 software that reports much more performance information through the SMI-S interface.

Complete Table 3-10 for each SVC that you want to monitor in the environment. The table assumes a fixed sample rate is chosen for all SVCs. If you plan to monitor some SVCs at 15-minute intervals and other SVCs at 30-minute intervals, you need to fill out this table twice: one time for each chosen sample rate, and then, add the two tables together to give you an overall total.

*Table 3-10 Repository sizing for SVC and TPC Version 3.1.2 and below*

| **Subsystem** | **Number of VDisks** | **Number of MDisks** | **I/O groups** | **MDisk groups** |
|---|---|---|---|---|
| TEC_SVC | 900 | 500 | 1 | 4 |
| SVC_Jason | 3,000 | 1,500 | 2 | 6 |
| | | | | |
| | | | | |
| | | | | |
| **Totals** | 3,900 | 2,000 | 3 | 10 |
| Record size (bytes) | 44 | 76 | 46 | 112 |
| Byte totals | 171,600 | 152,000 | 138 | 1,120 |
| | | (a) Sample rate (bytes) | | 324,858 |
| | | (b) 15-minute sample rate<br>(60/15) x (a) | | 1,299,432 |
| | | (c) Daily amount<br>(b) x 24 | | 31,186,368 |
| | | (d) 14-day retention of samples<br>(b) x 24 x 14 | | 436,609,152 |
| | | (e) 30-day retention of hourly<br>24 x (a) x 30 | | 233,897,760 |
| | | (f) 90-day retention of daily<br>(a) x 90 | | 29,237,220 |
| | | (g) Overall total required (MB)<br>(d) + (e) + (f)/1,024,000 + 50% | | **1,025 MB** |

**Important:** Notice that the overall figure in (g) adds 50% to the amounts calculated through the table. The majority of this overhead takes the DB2 table indexes for this data plus database page overhead into account. TPC 3.1.2 and below used approximately a third of the repository space of TPC 3.1.3 and above.

### Sizing the repository for SAN fabric performance data

This section describes sizing for SAN fabric performance collection. Fabric port record sizes per data sample are relatively large. We also observed that the indexing requirements for this data also tend to be high.

We based the byte sizing that we provide here on real world observations of actual database utilization over a period of time.

**Note:** Table 3-11 shows all of the switches sampled at 5-minute intervals. If you plan to monitor some switches at one rate and other switches at another rate, create a separate table for each rate.

The final figure includes a 50% uplift for indexing and DB2 storage overhead.

*Table 3-11 SAN switch performance repository data sizing*

| Switch name | (a) Number of ports | (b) Size (bytes) | (c) Sample rate (minutes) | (d) Hourly amount (bytes) (60/(c)) x (a) x (b) |
|---|---|---|---|---|
| TEC_Switch_1 | 32 | 400 | 5 | 153,600 |
| TEC_Switch_2 | 32 | 400 | 5 | 153,600 |
| Remote_Switch_1 | 64 | 400 | 5 | 307,200 |
| Remote_switch_2 | 64 | 400 | 5 | 307,200 |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| **Totals** | (e) 192 | (f) Total sample rate per hour | | 921,600 |
| | | (g) 14 days retain sample rate<br>(f) x 24 x 14 | | 309,657,600 |
| | | (h) 30 days retain hourly rate<br>(e) x (b) x 24 x 30 | | 55,296,000 |
| | | (i) 90 days retain daily rate<br>(e) x (b) x 90 | | 6,912,000 |
| | | Overall Total MB<br>(g) + (h) + (i)/1,024,000 + 50% | | **544.72 MB** |

## Sizing the repository for TPC for Data requirements

Repository sizing for TPC for Data is more difficult to accurately model due to the dynamic nature of the collected data. Performance data collection sizing is simple in that it collects a set amount of data at regular intervals.

However with TPC for Data, a policy or profile collects a variable amount of data from each monitored server based on what and how much data of a matching type is found on each machine.

Key factors in sizing TPC for Data:

- Total number of operating system registered users storing files
- Total number of filesystems monitored
- Total number of different file types (that is, *.txt, *.exe, *.doc, *.mp3, and so forth)
- Number of machines with data agents deployed and collecting data
- Total number of file names collected and stored for reporting

Key largest repository tables:

- T_STAT_USER_HIST - User history file
- T_STAT_FTYPE_HIST - File type history
- T_STAT_FILE - Stored file names

Figure 3-12, Table 3-13 on page 110, and Table 3-14 on page 111 help to estimate the worst case sizing for these key tables.

*Table 3-12 Estimating the user history repository requirement*

| **Statistic name** | **Number of filesystems covered** | **Number of users covered** | **Days to keep scan history** | **Number of weeks of scan history** | **Number of months of scan history** |
|---|---|---|---|---|---|
| Custom_stat | 300 | 800 | 30 | 52 | 24 |
| UNIX_stat | 250 | 1500 | 30 | 52 | 24 |
| Windows_stat | 500 | 2000 | 30 | 52 | 24 |
| | | | | | |
| Totals | (a) 1050 | (b) 4300 | (c) 90 | (d) 156 | (e) 72 |
| | | | Worst case total requirement (bytes)<br>a x b x (c + d + e) x 45 bytes | | 64,609,650,000 |
| | | | Realistic expectation reduce to 10% of worst | | 6,460,965,000 |
| | | | Divide 1,024,000 = MB | | **6,310 MB** |

**Note:** Unlike the performance tables, we must estimate much more here. For example, there might be 500 filesystems covered by the Windows_stat and 2,000 users with data across the 500 filesystems, but not all of the 500 filesystems have files owned by the 2,000 users. There is likely only a subset of filesystems with data for all 2,000 users. This is why the realistic figure is reduced to only 10% of the worst case figure. You might want to change the 10% factor to your specific requirements.

Use Table 3-13 to calculate the repository space required to store file type history information.

Estimating the file type history buildup is more accurate than estimating the user table history, because the data entering this table is more constant for a given profile.

*Table 3-13 Estimating the file type repository requirement*

| **Statistic profile name** | **Number of file types** | **Number of TPC agents covered** | **Days to keep scan history** | **Number of weeks of scan history** | **Number of months of scan history** |
|---|---|---|---|---|---|
| Win_types | 50 | 200 | 30 | 52 | 24 |
| UNIX_servers | 50 | 150 | 60 | 52 | 24 |
| Media files | 30 | 50 | 60 | 52 | 24 |
| | | | | | |
| Totals | (a) 130 | (b) 400 | (c) 150 | (d) 156 | (e) 72 |
| | | | Total - bytes<br>a x b x (c + d + e) x 55 bytes | | 1,081,080,000 |
| | | | Total MB - total/1,024,000 | | **1,056 MB** |

The third TPC for Data repository table of significant size is the T_STAT_FILE table. This table holds a record of the file names, which have been collected by profiles for largest, most obsolete, orphan files, and so forth.

**Note:** If you plan to use TPC for Data for duplicate file spotting or to archive specific files for you, it is likely that you will increase the number of file names that each agent will collect. See 2.2.1, “How does TPC identify a duplicate file” on page 24 to learn how to configure TPC for improved duplicate file reporting, which will impact the size of this table.

When completing Table 3-14 on page 111, the “Total file names per agent” will be the total of all types as seen in Figure 3-32 on page 111. In this example, it will be `1,800` file names per agent.

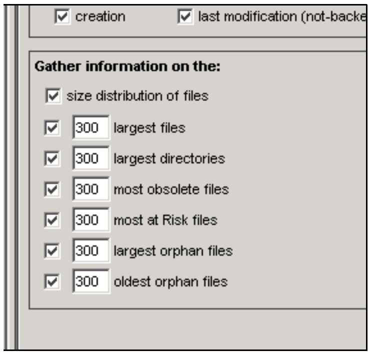

☑ creation ☑ last modification (not-backe

Gather information on the:

☑ size distribution of files

☑ 300 largest files

☑ 300 largest directories

☑ 300 most obsolete files

☑ 300 most at Risk files

☑ 300 largest orphan files

☑ 300 oldest orphan files

*Figure 3-32 Adding up all file names*

*Table 3-14 Estimating the file name repository requirement*

| Statistic profile name | (a) Total file names collected per agent | (b) Number of agents to which this profile applies | Total files per statistic<br>a x b |
|---|---|---|---|
| Duplicate file spot | 2,000 | 500 | 1,000,000 |
| Control audio files | 200 | 150 | 30,000 |
| Archive old data | 200 | 50 | 10,000 |
| | | | |
| | | Total files in table | 1,040,000 |
| | | Size (bytes) =<br>Total x 250 bytes | 420,000,000 |
| | | Size/1,024,000 = MB | **410 MB** |

The final step for sizing the TPC for Data repository is to total the three tables and add an overhead for the default statistics. The average overhead for the default statistic types is provided at 1.5 MB per TPC agent. Therefore:

Default TPC for Data overhead = Total agents x 1.5 MB

Example:

1,000 x 1.5 = 1,500 MB

Enter this figure in Table 3-15.

*Table 3-15 TPC for Data repository total*

| Source | Amount in MB |
|---|---|
| User history | 6,310 |
| File type history | 1,056 |
| File names | 410 |
| Default statistics overhead | 1,500 |
| **Total requirement (MB)** | **9,276** |

## 3.15 Simple repository database tuning

By default, the TPC databases and their DB2 active logs are stored on the same filesystem. You can achieve performance improvements by placing the logs on a separate filesystem or a separate disk drive to balance the I/O requirements of both tasks.

Of the two databases, the TPCDB is the most heavily used and, therefore, the prime candidate for you to move its logs. The IBMCDB database is used by the Agent Manager and is a much less demanding database from an I/O perspective.

To move the logs for the TPCDB database to a new location, use the following steps:

1. Choose a new log path location. For this example, `E:\DB2_active_logs\TPCDB`
2. Start a DB2 command line processor (Figure 3-33).

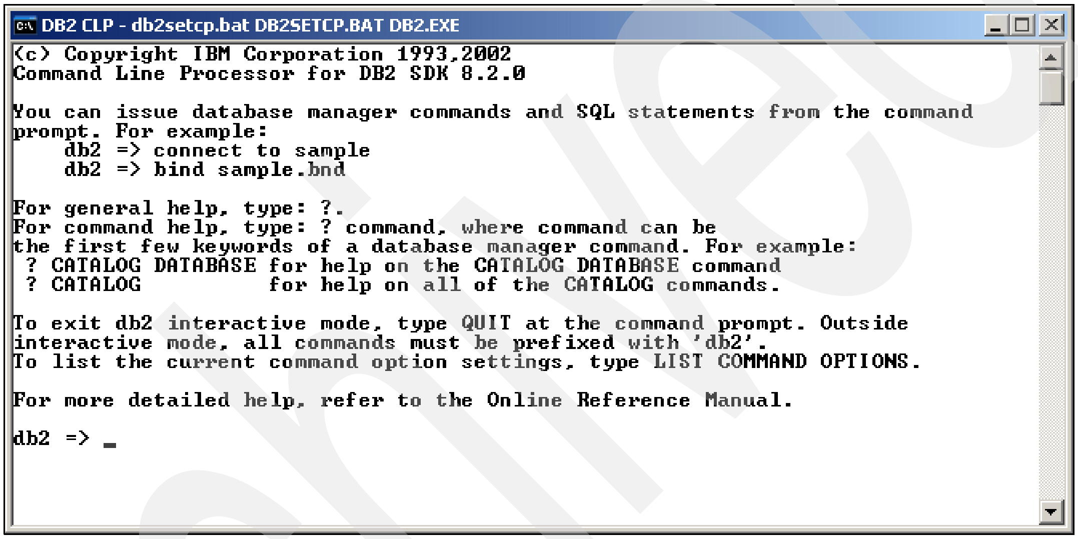

```
(c) Copyright IBM Corporation 1993,2002
Command Line Processor for DB2 SDK 8.2.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _
```

*Figure 3-33 DB2 command line processor*

3. Issue the following commands in the window:

   ```
   update db cfg for TPCDB using newlogpath E:\DB2_active_logs\TPCDB
   quit
   exit
   ```

The new log path goes into effect the next time that the database closes and opens. Stop TPC and restart it to use the new log path. Another method to use the new log path is to reboot the TPC server.

### Stop TPC services

Stop the TPC server by issuing the commands that are shown in Example 3-25 in a command window on Windows.

*Example 3-25 Windows commands to start TPC*

```
net stop "IBM WebSphere Application Server V5 - Tivoli Agent Manager"
net stop "IBM TotalStorage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V5 - Device Server"
```

### Start the TPC Services

Start the TPC server by issuing the commands that are shown in Example 3-26 in a command window on Windows.

*Example 3-26 Windows commands to start TPC*

```
net start “IBM WebSphere Application Server V5 - Tivoli Agent Manager”
net start “IBM TotalStorage Productivity Center - Data Server”
net start “IBM WebSphere Application Server V5 - Device Server”
```

## 3.16 Repository calculation templates

This section has blank versions of the worksheets used in 3.14, “TPC repository database sizing” on page 101 to calculate the DB2 repository space requirements for a given environment. Use these worksheets to help you size your individual requirements.

## Worksheet - Sizing SVC performance collection for TPC 3.1.3 and above

Refer to “Worksheet - Sizing SVC performance collection for TPC 3.1.3 and above” on page 114 to see a working example.

*Table 3-16 Repository sizing for SVC and TPC Version 3.1.3 and above*

| Subsystem | Number of VDisks | Number of MDisks | I/O groups | MDisk groups | Cluster pairs |
|---|---|---|---|---|---|
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| **Totals** | | | | | |
| Record size (bytes) | 198 | 78 | 500 | 128 | 492 |
| Byte totals | | | | | |
| | | | (a) Sample rate (bytes) | | |
| | | | (b) 15-minute sample rate<br>(60/15) x (a) | | |
| | | | (c) Daily amount<br>(b) x 24 | | |
| | | | (d) 14-day retention of samples<br>(b) x 24 x 14 | | |
| | | | (e) 30-day retention of hourly<br>24 x (a) x 30 | | |
| | | | (f) 90 days of daily retention<br>(a) x 90 | | |
| | | | (g) Overall total required (MB)<br>(d) + (e) + (f)/1,024,000 + 50% | | |

## Worksheet - Sizing SVC performance collection for TPC 3.1.2 and below

Refer to “Sizing the repository for SVC - TPC Version 3.1.2 and below” on page 107 for a working example.

*Table 3-17 Repository sizing for SVC and TPC Version 3.1.2 and below*

| Subsystem | Number of VDisks | Number of MDisks | I/O groups | MDisk groups |
|---|---|---|---|---|
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| **Totals** | | | | |
| Record size (bytes) | 44 | 76 | 46 | 112 |
| Byte totals | | | | |
| | | (a) Sample rate (bytes) | | |
| | | (b) 15-minute sample rate<br>(60/15) x (a) | | |
| | | (c) Daily amount<br>(b) x 24 | | |
| | | (d) 14-day retention of samples<br>(b) x 24 x 14 | | |
| | | (e) 30-day retention of hourly<br>24 x (a) x 30 | | |
| | | (f) 90 days of daily retention<br>(a) x 90 | | |
| | | (g) Overall total required (MB)<br>(d) + (e) + (f)/1,024,000 + 50% | | |

### Worksheet - Sizing performance collection for ESS, DSxxxx, and non-IBM

Refer to “Sizing the repository for ESS, DSxxxx, and non-IBM subsystems” on page 104 for a working example of this table.

Table 3-18 is the first of two worksheets necessary to calculate the repository space that is required for these types of subsystems.

*Table 3-18 Per sample repository database sizing for ESS, DSxxxx, and non-IBM subsystems*

| **(a) Subsystem name** | **(b) Number of volumes (LUNs) sampled** | **(c) Performance collection interval (minutes)** | **(d) Performance data record size** | **(e) Daily amount of data collected**<br>**(60/(c) x 24) x (b) x (d) = (e)** |
|---|---|---|---|---|
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | | 200 | |
| | | (f) Total required per day | | |
| | | (g) Number of days to retain per sample = 14 days<br>(f) x (g)/1,024,000 + 50% | | |

Table 3-19 is the second table needed to calculate repository space required for these types of subsystems.

*Table 3-19 Hourly and daily repository database sizing for ESS, DSxxxx, and non-IBM storage*

| (a) Subsystem name | (b) Number of volumes sampled (LUNs) | (c) Performance data record size (bytes) | (d) Hourly requirement<br>(b) x (c) x 24 | (e) Daily requirement<br>(b) x (c) |
|---|---|---|---|---|
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | 200 | | |
| | | **Daily totals** | | |
| | | Hourly days x 30 | (f) | |
| | | Daily days = 90 | | (g) |
| | | | **Total MB**<br>(f) + (g)/1,024,000 + 50% | |

## Worksheet - Sizing SAN switch performance collection

Refer to “Sizing the repository for SAN fabric performance data” on page 108 for a working example of how to calculate the amount for storage that is required to hold SAN switch performance data.

If you will monitor SAN switches at different time intervals, use a separate worksheet for each sample rate chosen.

*Table 3-20 SAN switch performance repository data sizing*

| **Switch name** | **(a) Number of ports** | **(b) Size (bytes)** | **(c) Sample rate (minutes)** | **(d) Hourly amount (bytes) (60/(c)) x (a) x (b)** |
|---|---|---|---|---|
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| | | 400 | | |
| **Totals** | (e) | (f) Total sample rate per hour | | |
| | | (g) 14 days retain sample rate<br>(f) x 24 x 14 | | |
| | | (h) 30 days retain hourly rate<br>(e) x (b) x 24 x 30 | | |
| | | (i) 90 days retain daily rate<br>(e) x (b) x 90 | | |
| | | Overall Total MB<br>(g) + (h) + (i)/1,024,000 + 50% | | |

## Worksheet - Sizing TPC for Data repository

Refer to 3.14, “TPC repository database sizing” on page 101 to see working examples of these tables.

*Table 3-21 Estimating the user history repository requirement*

| Statistic name | Number of filesystems covered | Number of users covered | Days to keep scan history | Number of weeks of scan history | Number of months of scan history |
|---|---|---|---|---|---|
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| Totals | (a) | (b) | (c) | (d) | (e) |
| | | | Worst case total requirement (bytes)<br>a x b x (c + d + e) x 45 bytes | | |
| | | | Realistic expectation - reduce to 10% of worst | | |
| | | | Divide 1,024,000 = MB | | |

*Table 3-22 Estimating the file type repository requirement*

| Statistic profile name | Number of file types | Number of TPC agents covered | Days to keep scan history | Number of weeks of scan history | Number of months of scan history |
|---|---|---|---|---|---|
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| Totals | (a) | (b) | (c) | (d) | (e) |
| | | | Total - bytes<br>a x b x (c + d + e) x 55 bytes | | |
| | | | Total MB - total/1,024,000 | | |

*Table 3-23 Estimating the file name repository requirement*

| **Statistic profile name** | **(a) Total file names collected per agent** | **(b) Number of agents to which this profile applies** | **Total files per statistic**<br>**a x b** |
|---|---|---|---|
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | Total files in table | |
| | | Size (bytes) =<br>Total x 250 bytes | |
| | | Size/1,024,000 = MB | |

The default overhead for the default statistic types is 1.5 MB per TPC for Data agent. Therefore:

Default TPC for Data overhead = Total agents x 1.5 MB

Example:

1,000 x 1.5 = 1,500 MB

Enter this figure in Table 3-15 on page 111.

*Table 3-24 TPC for Data repository total*

| **Source** | **Amount in MB** |
|---|---|
| User history | |
| File type history | |
| File names | |
| Default statistics overhead | |
| **Total requirement MB** | |

# 4

# Reporting and monitoring basics

In an existing environment, you have monitoring, alerting, and reporting mechanisms that are already implemented with several tools. Storage devices have Call Home features and can send SNMP traps. You can manage SAN switches with tools that are provided by vendors, such as EFC Manager by McData or Fabric Manager by Brocade, and also can send SNMP traps. Servers might have Tivoli endpoint monitors, and databases have their own monitoring tools.

IBM TotalStorage Productivity Center (TPC) provides central managing, monitoring, and reporting functions for storage-related tasks. TPC uses the Storage Management Initiative Specification (SMI-S) standard. Therefore, TPC can handle vendor comprehensive products.

TPC does not replace the need for other monitoring and alerting tools. It can only replace storage-related monitoring, managing, and reporting tasks, or run as an additional enhancement to the existing environment.

In this chapter, we highlight the basic monitoring, alerting, and reporting steps with TPC.

# 4.1 Monitoring and alerting

The chapter is a brief overview about how to set up the monitoring and alerting for IBM TPC after its installation with its agents.

We first want to clarify the differences between alerting and monitoring within TPC. *Monitoring* runs the data collection and stores the samples in the TPC database repository. The *Alerting mechanism* in TPC is related to alerting management.

TPC provides a set of default configurations that provides a quick start after the installation without much configuration effort.

## 4.1.1 How alerting works

The user-defined monitoring (data collecting) tasks within TPC are: Discovery, Probes, Scans, Pings, and Performance monitoring.

You can classify the source of the alerts into the following structure:

- Discovery
- Ping, scan, quota, and constraints (TPC for Data)
- Probes
- Performance data (TPC for Disk and Fabric)
- Monitoring job triggering condition (monitor failed)
- Jobs
- SNMP traps from the SAN switches
- Indications from the fabric (common fabric agent)
- Indications from the storage subsystem (CIMOM)

Figure 4-1 on page 123 shows an overview of the alerting flow within TPC.

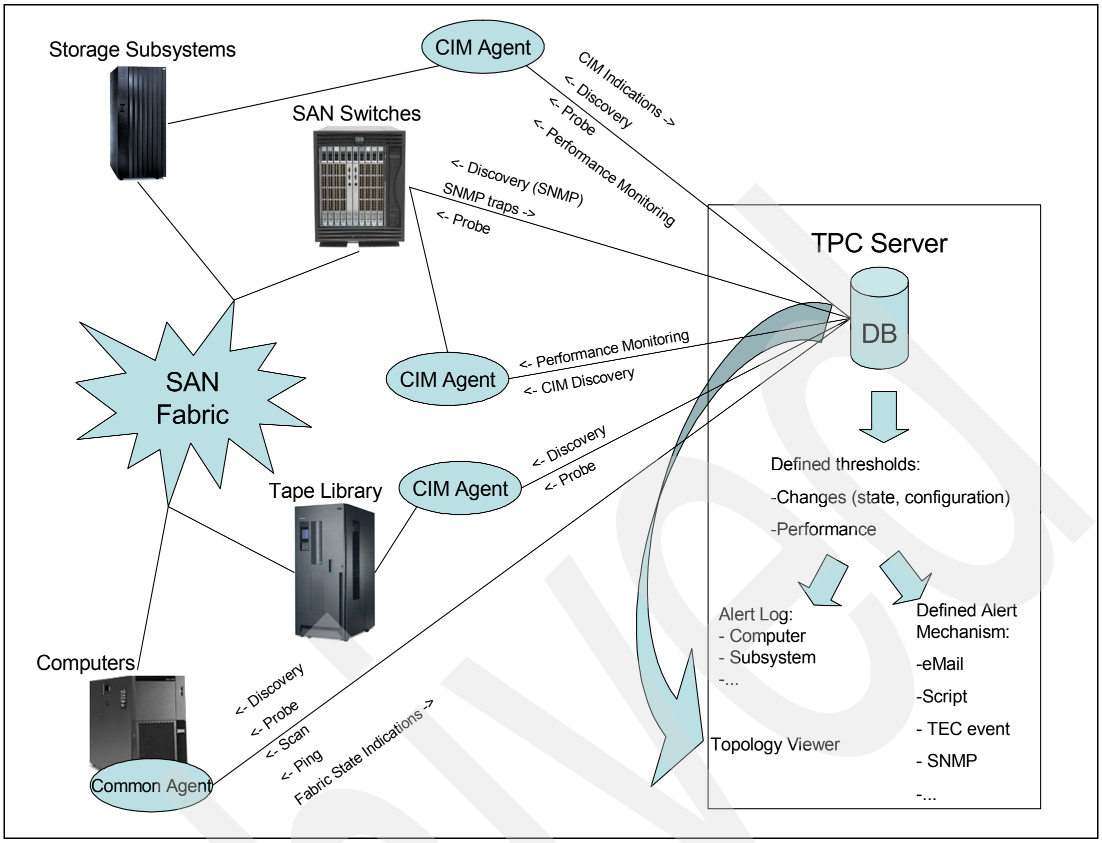

*Figure 4-1 Alerting flow overview*

You use the alerting mechanism in TPC to notify predefined receivers about specific storage condition violations in the environment. Many conditions produce alerts. The following summary shows the types of alerts that can be produced during different types of monitoring processes.

Probe identifies the following alerts:

- Computer alerts
- Storage subsystem alerts
- Fabric alerts and SAN switch alerts
- Endpoint device alerts
- Instance alerts
- Database-tablespace alerts

Scans produce the following alerts:

- Filesystem alerts
- Directory alerts
- Table alerts
- Constraint alerts

Quota job run produces the following alert:

- User and user group consumption alerts

Pings produce the following alert:

- Computer unreachable

Performance monitoring produces the following :

- Performance alerts disk subsystem
- Performance alerts fabric switches

The monitoring process defined in a TPC environment is normally run on a defined interval. Alerts are only generated at the process run time. Often that time interval is too long. For example, a Subsystem Probe runs every 24 hours. To avoid this delay in generating alerts, TPC supports SNMP traps from SAN switches and CIM indications from the CIM Agent.

You configure the alerts in each TPC section (Data, Disk, Fabric), and the alerts show in the Alerting node in the **IBM TotalStorage Productivity Center** → **Alerting** → **Alert Log** path.

Discovery is the primary process to gather information about the connections between TPC and the monitored systems. It includes finding new locations of monitored resources that were moved.

In a typical environment today, the number of CIM Agents is clearly defined. Therefore, in most cases, it is acceptable to disable SLP discovery when the available CIM Agents are properly configured in TPC. This configuration reduces the duration of a CIMOM discovery, because TPC does not have to wait for the SLP broadcast responses. To disable SLP discovery, in the TPC tree structure, go to **Administrative Services** → **Discovery** → **CIMOM** and remove the check in **Scan local subnet.** See Figure 4-2.

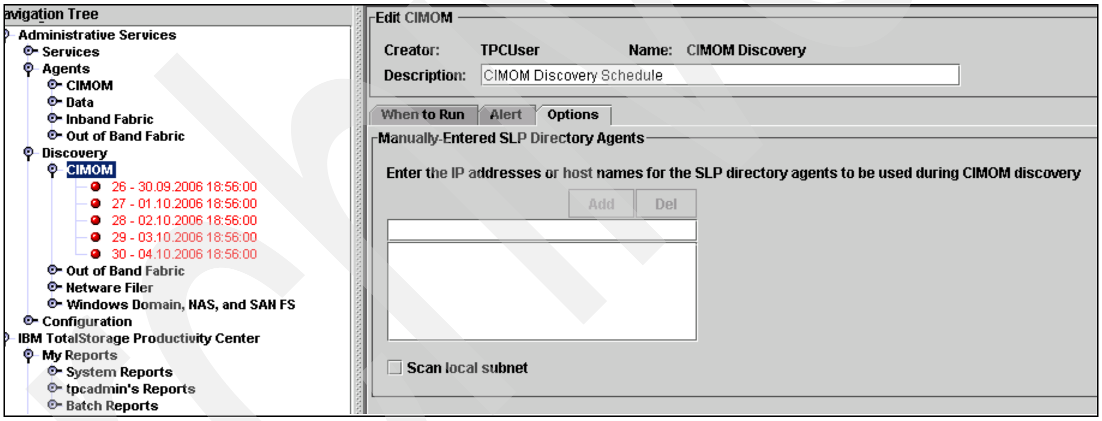

avigation Tree

- Administrative Services
  - Services
  - Agents
    - CIMOM
    - Data
    - Inband Fabric
    - Out of Band Fabric
  - Discovery
    - CIMOM
      - 26 - 30.09.2006 18:56:00
      - 27 - 01.10.2006 18:56:00
      - 28 - 02.10.2006 18:56:00
      - 29 - 03.10.2006 18:56:00
      - 30 - 04.10.2006 18:56:00
    - Out of Band Fabric
    - Netware Filer
    - Windows Domain, NAS, and SAN FS
  - Configuration
- IBM TotalStorage Productivity Center
  - My Reports
    - System Reports
    - tpcadmin's Reports
    - Batch Reports

Edit CIMOM

Creator: TPCUser Name: CIMOM Discovery

Description: CIMOM Discovery Schedule

Manually-Entered SLP Directory Agents

Enter the IP addresses or host names for the SLP directory agents to be used during CIMOM discovery

☐ Scan local subnet

*Figure 4-2 CIMOM discovery options*

## 4.1.2 Setting up monitoring and alerting for TPC for Data

IBM TotalStorage Productivity Center for Data (TPC for Data) uses the Common Data Agent on each server (Computer) to collect data. TPC runs different jobs, such as discovery, pings, scans, and probes periodically.

You use the information that is gathered by these jobs for:

- Alerts, if thresholds are violated
- Reports

Table 4-1 on page 125 lists the predefined default groups for TPC for Data, a description of the group, and where to find the information in the TPC tree structure.

*Table 4-1 Predefined default groups for TPC for Data*

| Where to find | Group name | Description |
|---|---|---|
| **Monitoring** → **Groups** → **Computer** | TPCUser.Default Computer Group | Every newly discovered computer is assigned to this group. |
| **Monitoring** → **Groups** → **Filesystem** | TPCUser.Default FS Group | Every newly discovered filesystem is assigned to this group. |
| **Monitoring** → **Groups** → **User** | TPCUser.Default User Group | Every newly discovered computer user is assigned to this group. |
| **Monitoring** → **Groups** → **OS User Group** | TPCUser.Default OSGroup Group | OS user groups are the operating system groups that TPC finds that own files and directories in the environment. |

It is helpful to use named groups for computers, storage subsystems, filesystems, directories, users, and OS user groups for use within monitoring jobs, such as pings, probes, scans, and alerting. Afterwards, you can easily assign groups to a monitoring job and if a new device needs to be monitored by TPC, include the device in the existing groups. Define groups by selecting **Data Manager** → **Monitoring** → **Groups**.

## Agents

Monitoring the agents that report to the TPC for Data server is important for several reasons. You want to identify failing or troubled agents. A large number of failed agents negatively impact the startup time of TPC, because TPC attempts to contact each agent during the startup phase. Failed agents also impact the information available for reports and alerts.

As a part of agent maintenance, we recommend that you check agent status periodically. You can check agent status in various ways:

- Check agent status by using the path **Administrative Services** → **Agents** → **Data** in the tree structure. Agent status needs to show as green. The repeated red status of an agent might signal that there is a problem with the agent, that the agent has been removed (from the environment) improperly, or that there is a connectivity issue.
- Generate the following report: **Data Manager** → **Reporting** → **Asset** → **System-wide** → **Computers** → **By Probe Time**. Data in this report can indicate which agents are not being probed, which might signal that there is a problem with the agent, that the agent has been removed (from the environment) improperly, or that there is a connectivity issue.

## Pings

TPC uses a simple network ping to check the network connectivity to each computer. Pings collect the information about the availability of the computers in your enterprise. See pings in the availability reporting by selecting **Data Manager** → **Reporting** → **Availability** → **Ping** in the tree structure.

After the installation of TPC for Data, the ping job in Table 4-2 on page 126 is predefined. The default ping monitor pings your computers every 30 minutes. Disable this if your computers are behind a firewall, and you do not want use computer availability for alerting or reporting.

If you intend to use computer alerts, such as “`Computer unreachable`,” you might decrease the ping frequency for those computers.

*Table 4-2 Predefined default values for TPC for Data*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring → Pings** | TPCUser.Default Ping | The group TPCUser.Default Computer Group is assigned to this monitor. Thus, by default, all computers are pinged in a 30-minute frequency. |

## Probe

Probes collect the structure information about the computer in your enterprise, such as hard disks, clusters, filesystems, directories, and information, such as OS level, serial number, RAM, swap space, and CPU.

After the installation of TPC for Data, the probe job listed in Table 4-3 is predefined. It probes your computers once a day. You can see the collected information in the asset, capacity, and storage subsystem reporting.

*Table 4-3 Predefined default values for probes*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring → Probes** | TPCUser.Default Probe | Probes TPCUser.Default Computer Group once a day. |

If you intend to use computer alerts, you might decrease the probe rate for those computers.

## Scan

Unlike the TPC components for Disk and Fabric, TPC for Data has a scan function. Use scans to collect statistics about the usage and trending of storage consumption. Scan data is stored in the TPC repository, and you can use scan data for several types of reports and alert conditions.

The TPC server schedules and manages the scan job. The Common Agent on the Server performs the scan job. Agent scans can consume significant resources, such as CPU and memory, during execution and might run for a long time, depending on the number of files and filesystems. Therefore, run a scan at non-peak time.

You can see the collected information in the usage, usage violation, and backup reporting.

TPC for Data runs a scan once a day by default using the TPCUser.Default FS Group. This group includes all filesystems by default.

You use profiles to specify what information is gathered during a scan. Profiles are associated with scan jobs. TPC for Data has several predefined profiles, which meet most requirements. Table 4-4 on page 127 lists the default profiles.

*Table 4-4 Default profiles within TPC for Data*

| Where to find | Condition | Description | Assigned to default scan |
|---|---|---|---|
| **Monitoring → Profiles** | Most at risk | This default profile gathers statistics about the *n* files that have been modified the longest per drive and have not been backed up since they were modified. Default amount is 20 files and the files in ?:\WINNT\system*\% are excluded. (Windows only) | Yes |
| **Monitoring → Profiles** | By access | This default profile gathers statistics by the length of time since the last access of the files. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | Yes |
| **Monitoring → Profiles** | By creation | This default profile gathers storage statistics by the length of time since the creation of the files. This profile applies to scans run against Windows machines only. Scan jobs with this profile do not gather any creation time data from non-Windows machines. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | Yes |
| **Monitoring → Profiles** | By modification not backed up | This default profile gathers statistics by the length of time since the last modification (only for files not backed up since modification). By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. (Windows only) | Yes |

| Where to find | Condition | Description | Assigned to default scan |
|---|---|---|---|
| **Monitoring → Profiles** | By modification | This default profile gathers statistics by the length of time since the last modification of files. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | Yes |
| **Monitoring → Profiles** | File size distribution | This default profile gathers information about the size distribution of files. | Yes |
| **Monitoring → Profiles** | Largest directories | This default profile gathers statistics about the *n* largest directories. (The default amount is 20.) | Yes |
| **Monitoring → Profiles** | Largest files | This default profile gathers statistics about the *n* largest files. (The default amount is 20.) | Yes |
| **Monitoring → Profiles** | Largest orphans | This default profile gathers statistics about the *n* largest orphaned files. (The default amount is 20.) | Yes |
| **Monitoring → Profiles** | Most obsolete files | This default profile gathers statistics about the *n* "most obsolete" files (files that have not been accessed or modified for the longest period of time). (The default amount is 20.) Files in ?:\WINNT\system*\%, /usr/bin/%, /sbin/%, /usr/sbin/%,/etc/% are excluded. | Yes |
| **Monitoring → Profiles** | Oldest orphans | This default profile gathers statistics about the *n* oldest orphaned files. (The default amount is 20.) | Yes |
| **Monitoring → Profiles** | Summary by file type | This default profile summarizes space usage by file types (for example, .exe, .dll, .doc, .mp3, and so forth). The default is the 100 file types using the most space. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | Yes |

| Where to find | Condition | Description | Assigned to default scan |
|---|---|---|---|
| **Monitoring → Profiles** | Summary by filesystem/directory | This default profile summarizes space usage by filesystem or directory. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | Yes |
| **Monitoring → Profiles** | Summary by group | This default profile summarizes space usage by OS group. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | No |
| **Monitoring → Profiles** | Summary by owner | This default profile summarizes space usage by owner. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | Yes |
| **Monitoring → Profiles** | Temporary files | This default profile gathers statistics about the non-OS files not accessed in the last year and orphan files. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. By default, files with extensions *.tmp, *.temp, and *.bak are included. | No |
| **Monitoring → Profiles** | Wasted space | This default profile gathers statistics about the non-OS files not accessed in the last year and orphaned files. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. | Yes |

## Alerting

Within TPC for Data, you can use an alert for computers, filesystems, and directories. TPC for Data detects alert condition violations during ping, discovery, probe, scan, and constraint jobs. If a condition exists, Data Manager sends a notification as defined in the alert that contained the condition.

After an TPC for Data installation, there are no alert conditions that are defined by default. Table 4-5 shows an extract of the computer alerts. We classify the alerts as:

- *Configuration change alerts,* such as `RAM Increased` or `New Filesystem Detected`. You might use configuration change alerts for historical purposes.
- *Operational alerts,* such as `New Disk Defect Detected`, `Disk failure Predicted`, or `Computer unreachable`. These alerts relate to operations and can negatively affect the service.

Refer to 4.9, “Alerting mechanism (event)” on page 179 if you want use an alert mechanism to notify receivers.

The requirements for every enterprise are extremely different. If you want use alerts, a good starting point is to define at least the operational alerts and adjust the alerts continuously. The alerts that are listed in Table 4-5 through Table 4-7 on page 131 are examples for an alerting configuration for a computer, a filesystem, and a directory.

*Table 4-5 Extract of computer alerts*

| Where to find | Condition | Explanation |
|---|---|---|
| **Alerting → Computer Alerts** | New disk detected | A new disk is discovered on a managed computer. |
| **Alerting → Computer Alerts** | Disk not found | A disk is removed from a managed computer. |
| **Alerting → Computer Alerts** | New disk defect found | A new disk defect is detected on a managed disk. |
| **Alerting → Computer Alerts** | Grown disk defects exceed *n* value | The grown disk defect threshold is exceeded. |
| **Alerting → Computer Alerts** | Disk failure predicted | A managed disk predicts that a disk failure is imminent. |
| **Alerting → Computer Alerts** | Virtual server removed | A virtual server is removed. |
| **Alerting → Computer Alerts** | Virtual server moved | A new virtual server is moved. |
| **Alerting → Computer Alerts** | Computer unreachable | A managed computer cannot be reached. |
| **Alerting → Computer Alerts** | Computer discovered | A new non-managed computer is discovered. |
| **Alerting → Computer Alerts** | Computer status change offline | A managed computer goes offline. |
| **Alerting → Computer Alerts** | Computer status change online | A computer comes online. |

*Table 4-6 Extract of filesystem alerts*

| Where to find | Condition | Explanation |
|---|---|---|
| **Alerting → Filesystem Alerts** | Filesystem not found | A filesystem is removed or unmounted from a managed computer. |
| **Alerting → Filesystem Alerts** | Filesystem free space less than *x* | A filesystem has less than a value of Percent or MB of free space. |

*Table 4-7 Extract of directory alerts*

| Where to find | Condition | Explanation |
|---|---|---|
| **Alerting → Directory Alerts** | Directory not found | A directory cannot be found. |
| **Alerting → Directory Alerts** | Directory consumes more than *x* | A user or directory storage quota consumes more than a value of Percent or MB. |

### 4.1.3 Policy management

You can use policy management for the fine adjustment of alerts. Policy management identifies quotas, constraints, and filesystem extensions. You can use policy management to schedule script or archive and backup actions. You can find it at **Data Manager → Policy Management**. Table 4-8 on page 132 lists the various policy definitions.

*Table 4-8 Policies*

| Policy | Description |
|---|---|
| Quotas | Define limits on the amount of storage that a user or a group of users can consume. Quotas allow you to specify limits at three different levels for a user or group of users: at the filesystem level, at the computer level, and at the level of the whole network. You can work with Network Appliance™ quotas, specifically, import quotas from NetApp® filers, view the definition for each imported quota, and determine how you will be alerted when the hard limit defined with a NetApp quota is close to being violated. Define the run time of the quota check job and define alerts for each violated quota. |
| Constraints | Define the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment. Constraints enable you to restrict users from putting certain files (such as MP3 files) on monitored computers. Request a Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage. |
| Filesystem extensions | Create additional space in the filesystems of managed hosts. The filesystem extension policy enables you to extend filesystems manually or automatically. For managed hosts that have access to an Enterprise Storage Server®, you can also allocate additional LUNs. |
| Scheduled actions | Schedule actions to execute scripts against your storage-related resources. |
| Archive and backup | View and edit archive and backup jobs that were created based on files that were selected from a **Reporting** → **Usage** → **Files** report (select and right-click). |

## Quotas

There are no default quotas defined. Quotas can meet the condition *User/User Group Consumes More Than ...*. A quota run uses the TPC repository. Therefore, schedule the quota run after the scan job to ensure that you have the latest information. For every alert, you can define the alert mechanism. For example, you can automatically inform the owner of files or a directory about a threshold violation, or you can trigger a script to run. Refer to "Setting up quota and constraint e-mail" on page 134 for e-mail notification details. Define quotas at **Data Manager** → **Policy Management** → **Quotas**.

## Constraints

Constraints are a powerful way to analyze your filesystems. You can trigger alerts and use constraints for reporting. You assign constraint definitions to a filesystem. After each scan

(**Data Manager** → **Monitoring** → **Scans**) of a desired filesystem, the results are compared with the associated constraints. Figure 4-3 illustrates the composition of a constraint.

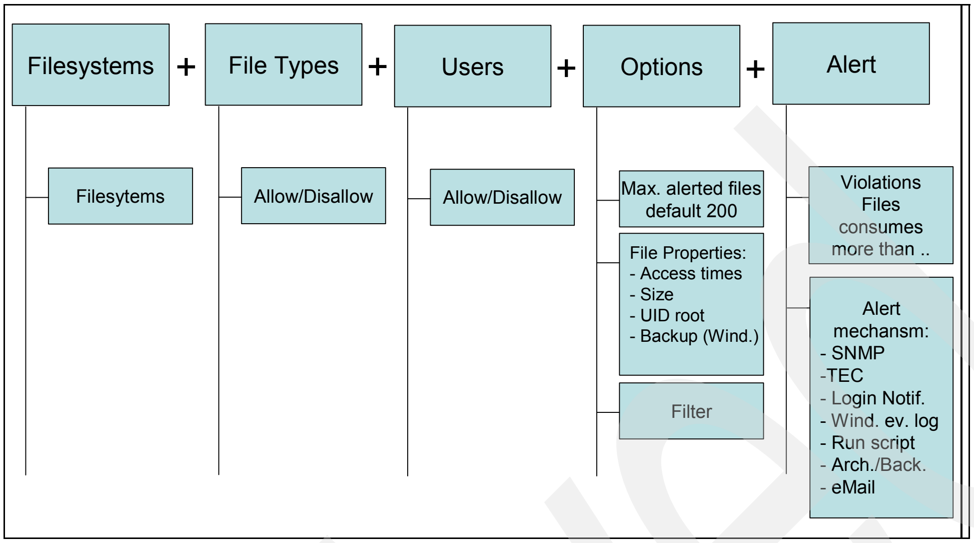

*Figure 4-3 Constraint composition*

Define constraints at **Data Manager** → **Policy Management** → **Constraints**.

During the installation of TPC for Data, the following default constraints are defined but are not associated with any filesystems.

*Table 4-9 Default constraints*

| Constraints | Description |
|---|---|
| TPCUser.At Risk File Constraint | This constraint finds files (directories and some system file names are excluded) with `LAST MODIFIED < 7 days 00:00` and `Not Backed up` (Windows).<br>The result has to be > `1 KB` to trigger an alert. The result count of file names is limited to `200` entries. See the violation alert in the ALERT LOG and see the result at **IBM TotalStorage Productivity Center** → **Data Manager** → **Reporting** → **Usage Violations** → **Constraint Violations.** |
| TPCUser.Obsolete File Constraint | This constraint finds all files (some system file names are excluded) with `LAST ACCESSED < 365 days 06:00.`<br>The result has to be > `1 KB` to trigger an alert. The result count of file names is limited to `200` entries. See the violation alert in the Alert Log and see the result at **Data Manager** → **Reporting** → **Usage Violations** → **Constraint Violations.** |
| TPCUser.Orphaned File Constraint | This constraint finds all `ORPHANED`. Orphaned files are files for which the owner does not exist any longer. The result has to be > `1 KB` to trigger an alert. The result count of file names is limited to `200` entries. See the violation alert in the Alert Log and see the result at **Data Manager** → **Reporting** → **Usage Violations** → **Constraint Violations**. |

Quotas and constraints offer an advanced e-mail notification mechanism. Therefore, you are able to send violation e-mails directly to the owner or the owner group. To use that alerting mechanism, the receiver's e-mail address must have similarities with the login account, such as *user name, first name,* or *last name*.

## Setting up quota and constraint e-mail

The steps are:

1. Define your mail server settings and especially the *Default Domain*. Set the e-mail configuration at **Administrative Services** → **Configuration** → **Alert Disposition**. The destination e-mail address always has the extension @<*your defined domain*>. TPC uses only this e-mail domain.

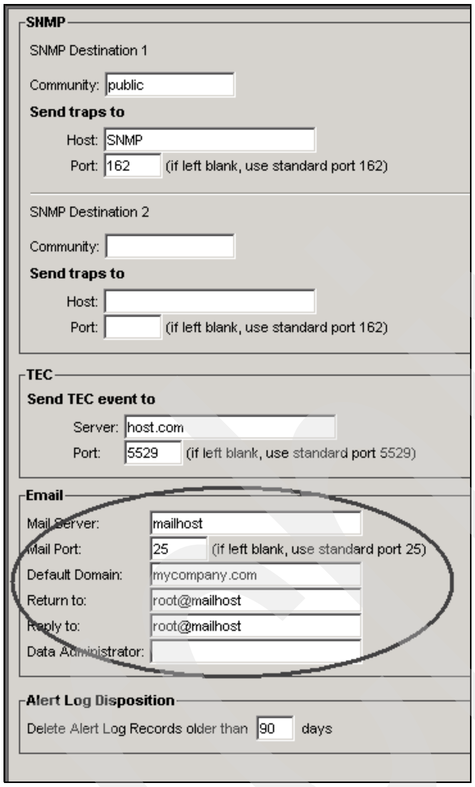

SNMP

SNMP Destination 1

Community: public

Send traps to

Host: SNMP

Port: 162 (if left blank, use standard port 162)

SNMP Destination 2

Community:

Send traps to

Host:

Port: (if left blank, use standard port 162)

TEC

Send TEC event to

Server: host.com

Port: 5529 (if left blank, use standard port 5529)

Email

Mail Server: mailhost

Mail Port: 25 (if left blank, use standard port 25)

Default Domain: mycompany.com

Return to: root@mailhost

Reply to: root@mailhost

Data Administrator:

Alert Log Disposition

Delete Alert Log Records older than 90 days

*Figure 4-4 Administrative Services → Configuration → Alert Disposition*

2. Set the quota and constraint e-mail address rules at **Administrative Services** → **Configuration** → **Quota and Constraint e-mail Address Rules**. Here you can define the e-mail receiver address composition. Only one rule is possible. You can use the *username*, *firstname*, *lastname*, just a substring of them, or a text phrase. By default, TPC uses the `LASTNAME` and adds the first letter of the `FIRSTNAME`. The user with Lastname=`Mueller` and Firstname=`Peter` has the e-mail address `muellerp@mycompany.com`.

   Use the functions **Add After**, **Add Before**, **Edit**, and **Delete** to order and edit the variables. To edit a variable, mark it by selecting it and click the desired function. Figure 4-6 on page 135 shows the panel that you use to define a substring.

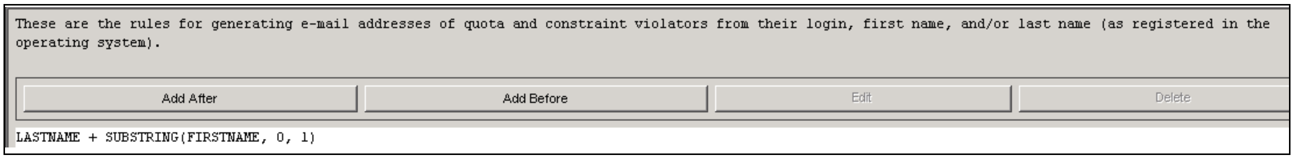

These are the rules for generating e-mail addresses of quota and constraint violators from their login, first name, and/or last name (as registered in the operating system).

```
LASTNAME + SUBSTRING(FIRSTNAME, 0, 1)
```

*Figure 4-5 Administrative Services → Configuration → Quota and Constraint e-mail Address Rules*

Use the variable type from the drop-down menu and choose the character positions that you want to use as shown in Figure 4-6.

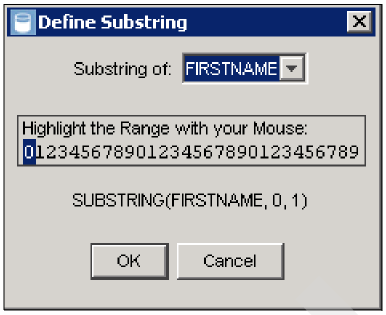

Substring of: FIRSTNAME

Highlight the Range with your Mouse:
0123456789012345678901234567890

SUBSTRING(FIRSTNAME, 0, 1)

*Figure 4-6 Define substring*

3. Select **Email to Quota Violator** for each triggering condition where you want to notify the violator. You can select it on the **Alert** tab of the quota or constraint definitions. With **Edit e-mail**, you can define the structure and content for the e-mail.

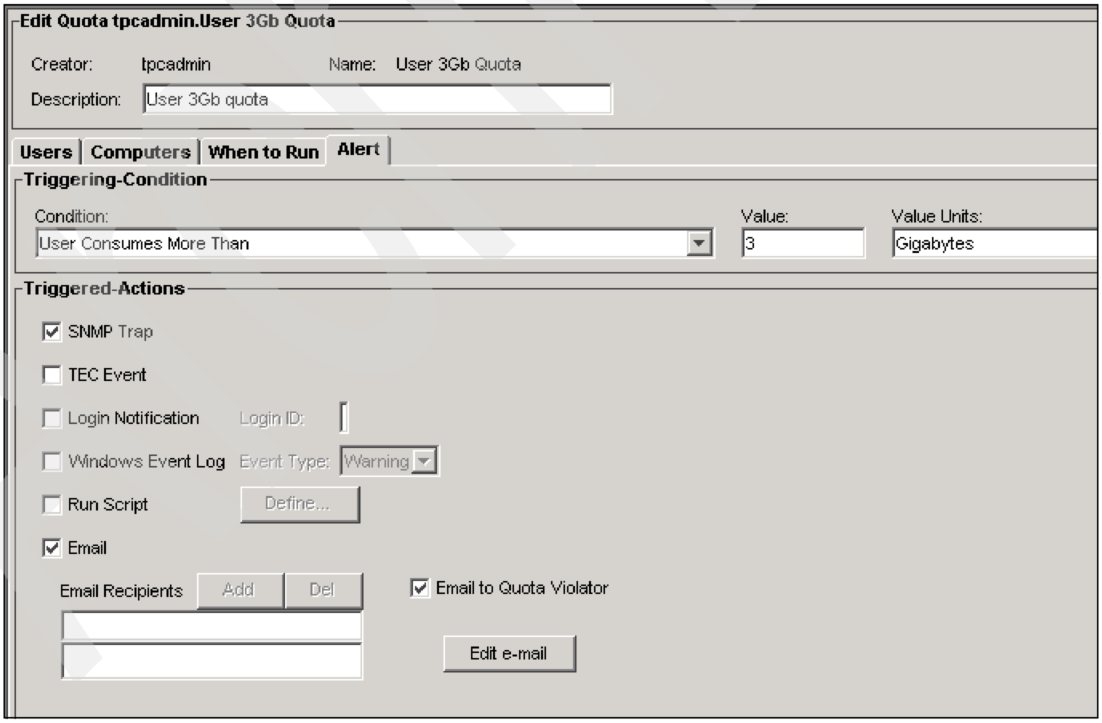

Edit Quota tpcadmin.User 3Gb Quota

Creator: tpcadmin Name: User 3Gb Quota

Description: User 3Gb quota

Triggering-Condition

Condition: User Consumes More Than

Value: 3

Value Units: Gigabytes

Triggered-Actions

☑ SNMP Trap

☐ TEC Event

☐ Login Notification Login ID:

☐ Windows Event Log Event Type: Warning

☐ Run Script

☑ Email

Email Recipients

☑ Email to Quota Violator

*Figure 4-7 IBM TotalStorage Productivity Center → Data Manager → Policy Management → Quotas*

4. Now, the settings are ready for sending an e-mail notification to the violator. TPC sends an e-mail to each violator at the next quota job run.

## 4.1.4 Filesystem extension

Filesystem extension (using **Data Manager** → **Policy Management** → **Filesystem Extension**) allows you to create additional space in the local filesystems of managed hosts. You can extend filesystems manually or set up a policy to extend filesystems automatically. You can configure a policy to extend filesystems at a specified time, or when utilization reaches a specified threshold. For managed hosts that have access to an ESS, you can also allocate additional LUNs when there is insufficient space to extend filesystems in the local volume group.

TPC supports filesystem extension for JFS filesystems running on AIX 5.1 and VxFS filesystems running on Sun Solaris™ 2.8. It provides automatic policy-based extension of a filesystem. For the latest support matrix, refer to the TPC Web site.

You can define the triggering condition as:

- Extend filesystem regardless of remaining free space
- Extend filesystem when free space is less than <*number*> <*GB/MB/%*>

You can also choose to have no extension performed when the policy is applied. Instead, any policy actions that have been performed are written to a log file. You can use this feature, called *LOG ONLY* mode, to preview the results of a policy before extending any filesystems (see Figure 4-8).

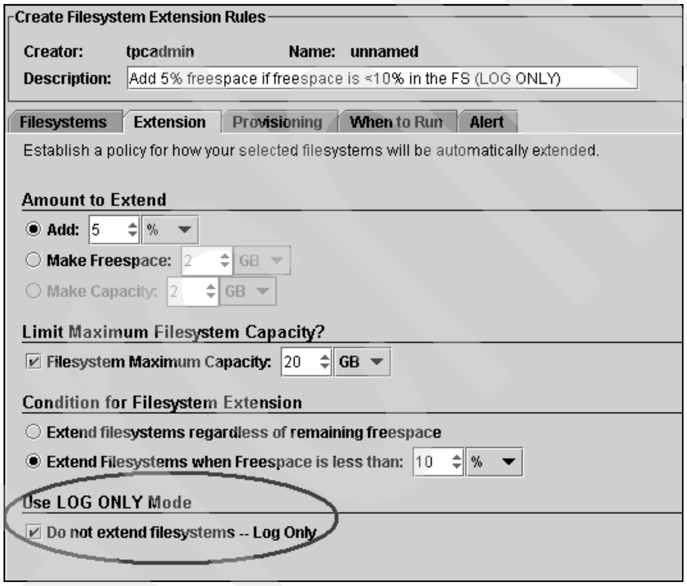

Create Filesystem Extension Rules

Creator: tpcadmin    Name: unnamed

Description: Add 5% freespace if freespace is <10% in the FS (LOG ONLY)

Establish a policy for how your selected filesystems will be automatically extended.

Amount to Extend

● Add: 5 %

○ Make Freespace: 2 GB

○ Make Capacity: 2 GB

Limit Maximum Filesystem Capacity?

☑ Filesystem Maximum Capacity: 20 GB

Condition for Filesystem Extension

○ Extend filesystems regardless of remaining freespace

● Extend Filesystems when Freespace is less than: 10 %

Use LOG ONLY Mode

☑ Do not extend filesystems -- Log Only

*Figure 4-8 Data Manager → Policy Management → Filesystem Extension*

**Note:** For simplification, we recommend that you use Tivoli Provision Manager (TPM) for profile-based filesystem extensions instead of TPC.

### Scheduled actions

Use scheduled actions (**Data Manager** → **Policy Management** → **Scheduled Actions**) for the scheduler to run scripts on clients. You can only run scripts on clients if the Common Agent configuration allows it. You set this option during Common Agent installation.

A scheduled action is a composition of:

- Computer or Computer Group
- Script name (a local file in \<*TPC_installation_directory*>\data\scripts)
- Scheduler
- Alert

The only triggering condition for using an alert is *Script Failed.*

> **Note:** The scripts that appear in the Script Name drop-down list are stored on the machine where the Data server component exists. If you want to run a script that is stored on a data agent, you must type the name of that script in the script name field. If TPC finds the script name on the agent and locally, it uses the agent script. Therefore, you can adjust the script on the agents to satisfy special requirements if needed.

The operating system under which an agent runs determines the scripts that can be run by that agent:

- UNIX and Linux - An agent running under a UNIX or Linux operating system does not run or receive (from the server) scripts that have an extension. Agents running under UNIX or Linux accept and run script files that do not contain an extension within their name. The first line in a script file determines what interpreter is needed to run the script.
- Windows - An agent running under a Windows operating system runs or receives (from the server) scripts that have an extension. The extension of a script file determines what interpreter is needed to run the script.

Figure 4-9 shows an example from a simple `hello_world.bat` script located on the Common Data Agent server. TPC only controls the submission of the scheduled scripts and batch processes. It does not handle OS-related return codes.

The output log contains the path to the script showing that the script was run at the Common Agent server.

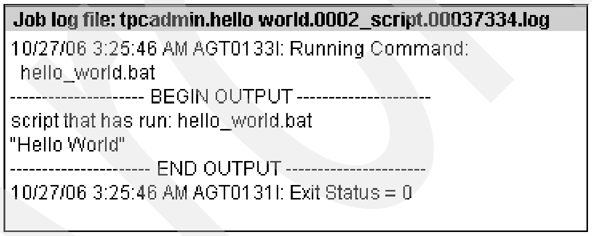

```
Job log file: tpcadmin.hello world.0002_script.00037334.log
10/27/06 3:25:46 AM AGT0133I: Running Command:
 hello_world.bat
-------------------- BEGIN OUTPUT --------------------
script that has run: hello_world.bat
"Hello World"
--------------------- END OUTPUT ---------------------
10/27/06 3:25:46 AM AGT0131I: Exit Status = 0
```

*Figure 4-9 Script without error*

Figure 4-10 shows a script with an unsuccessful exit code.

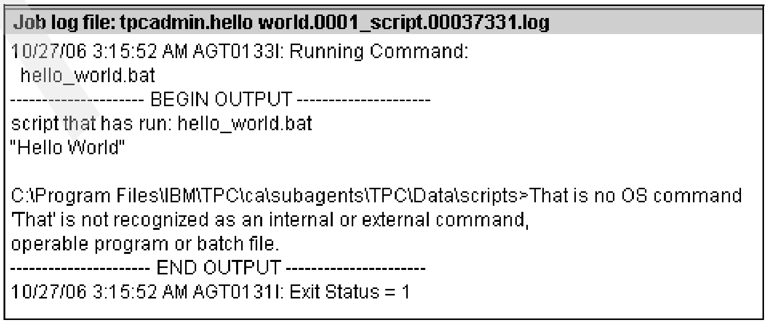

```
Job log file: tpcadmin.hello world.0001_script.00037331.log
10/27/06 3:15:52 AM AGT0133I: Running Command:
 hello_world.bat
-------------------- BEGIN OUTPUT --------------------
script that has run: hello_world.bat
"Hello World"

C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\scripts>That is no OS command
'That' is not recognized as an internal or external command,
operable program or batch file.
--------------------- END OUTPUT ---------------------
10/27/06 3:15:52 AM AGT0131I: Exit Status = 1
```

*Figure 4-10 Script with error*

### Archive and backup

TPC can use the Tivoli Storage Manager Client on the Common Agent Computer to archive or back up files. You can create file lists with reports. The Archive/Backup log (**Data Manager** → **Policy Management** → **Archive/Backup**) shows the details about the job with all its warnings and errors. Remember that only failed jobs trigger an alert; therefore, you need to check your job logs regularly. Figure 4-11 shows an example of a successful selective backup of a file. Figure 4-12 shows the same job; however in this case, it failed because the Tivoli Storage Manager node was locked on the Tivoli Storage Manager Server. Refer to 2.5, "Tivoli Storage Manager automatic archiving and integration" on page 48 for detailed information about integration of Tivoli Storage Manager in TPC.

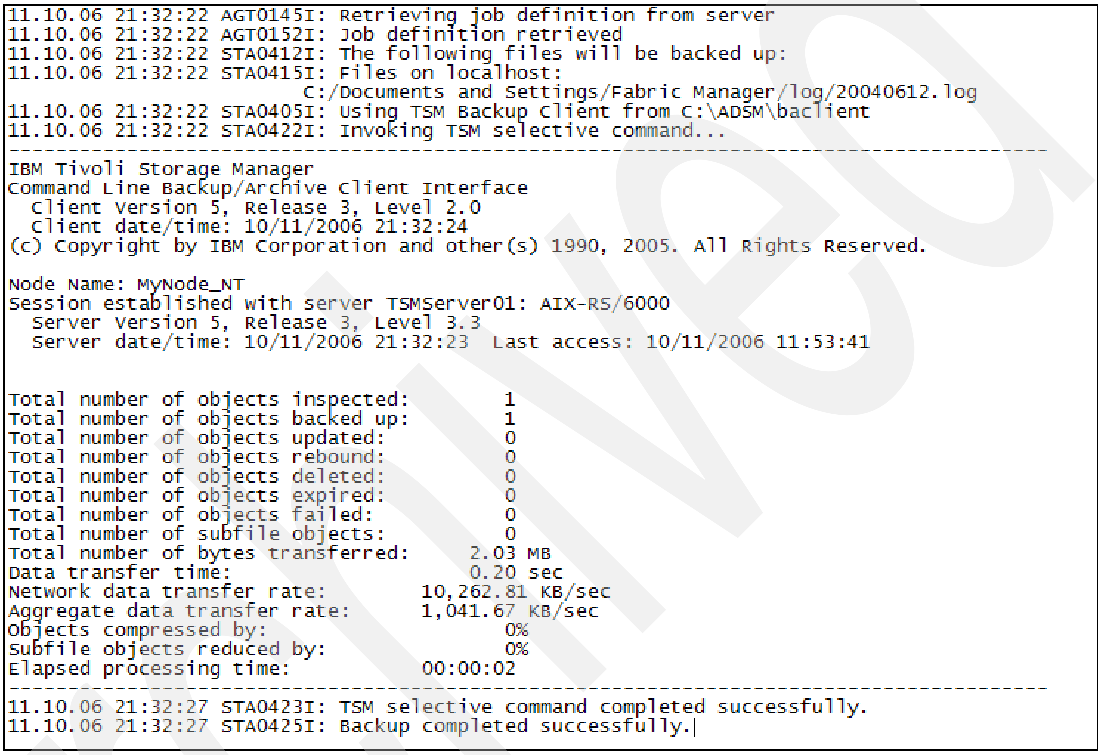

```
11.10.06 21:32:22 AGT0145I: Retrieving job definition from server
11.10.06 21:32:22 AGT0152I: Job definition retrieved
11.10.06 21:32:22 STA0412I: The following files will be backed up:
11.10.06 21:32:22 STA0415I: Files on localhost:
                         C:/Documents and Settings/Fabric Manager/log/20040612.log
11.10.06 21:32:22 STA0405I: Using TSM Backup Client from C:\ADSM\baclient
11.10.06 21:32:22 STA0422I: Invoking TSM selective command...
-----------------------------------------------------------------------------------------
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 3, Level 2.0
  Client date/time: 10/11/2006 21:32:24
(c) Copyright by IBM Corporation and other(s) 1990, 2005. All Rights Reserved.

Node Name: MyNode_NT
Session established with server TSMServer01: AIX-RS/6000
  Server Version 5, Release 3, Level 3.3
  Server date/time: 10/11/2006 21:32:23  Last access: 10/11/2006 11:53:41


Total number of objects inspected:        1
Total number of objects backed up:        1
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects expired:          0
Total number of objects failed:           0
Total number of subfile objects:          0
Total number of bytes transferred:     2.03 MB
Data transfer time:                    0.20 sec
Network data transfer rate:        10,262.81 KB/sec
Aggregate data transfer rate:      1,041.67 KB/sec
Objects compressed by:                    0%
Subfile objects reduced by:               0%
Elapsed processing time:           00:00:02
-----------------------------------------------------------------------------------------
11.10.06 21:32:27 STA0423I: TSM selective command completed successfully.
11.10.06 21:32:27 STA0425I: Backup completed successfully.|
```

*Figure 4-11 Successful selective backup of a file*

If the Tivoli Storage Manager node is locked, it generates an alert because the backup job failed.

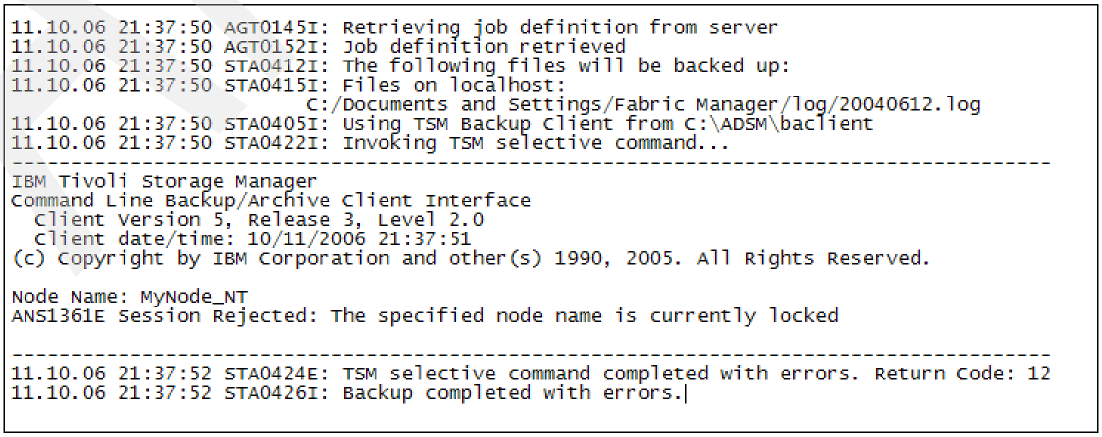

```
11.10.06 21:37:50 AGT0145I: Retrieving job definition from server
11.10.06 21:37:50 AGT0152I: Job definition retrieved
11.10.06 21:37:50 STA0412I: The following files will be backed up:
11.10.06 21:37:50 STA0415I: Files on localhost:
                         C:/Documents and Settings/Fabric Manager/log/20040612.log
11.10.06 21:37:50 STA0405I: Using TSM Backup Client from C:\ADSM\baclient
11.10.06 21:37:50 STA0422I: Invoking TSM selective command...
-----------------------------------------------------------------------------------------
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 3, Level 2.0
  Client date/time: 10/11/2006 21:37:51
(c) Copyright by IBM Corporation and other(s) 1990, 2005. All Rights Reserved.

Node Name: MyNode_NT
ANS1361E Session Rejected: The specified node name is currently locked

-----------------------------------------------------------------------------------------
11.10.06 21:37:52 STA0424E: TSM selective command completed with errors. Return Code: 12
11.10.06 21:37:52 STA0426I: Backup completed with errors.|
```

*Figure 4-12 Failed selective backup that triggers an alert*

The failed job appears in the alert log, which is shown in Figure 4-13. Select **IBM TotalStorage Productivity Center** → **Alerting** → **Alert Log** → **All**. You can define a triggering action on this event, such as running scripts, e-mails, or TEC, for example.

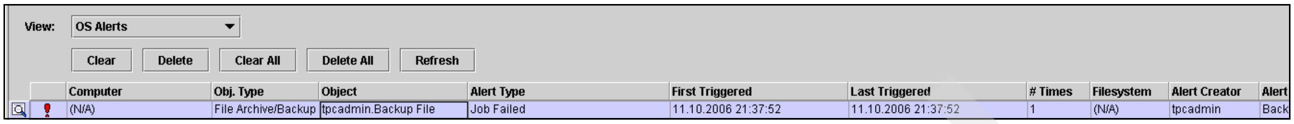

View: OS Alerts

| | | Computer | Obj. Type | Object | Alert Type | First Triggered | Last Triggered | # Times | Filesystem | Alert Creator | Alert |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | (N/A) | File Archive/Backup | tpcadmin.Backup File | Job Failed | 11.10.2006 21:37:52 | 11.10.2006 21:37:52 | 1 | (N/A) | tpcadmin | Back |

*Figure 4-13 Failed job is reported in the Alert Log*

# 4.2 Data Manager for databases

Database probes and scans are the primary mechanisms for collecting information about the RDBMS storage assets within an enterprise environment. This collected information helps identify, evaluate, control, and predict relational database needs.

Databases are usually critical for an enterprise. Identifying potential problem areas, such as running out of space, are important to identify and address in order to keep operations running smoothly. This section briefly discusses how to set up database probes and scans to collect this information.

## 4.2.1 RDBMS login

Before database monitoring can occur, you need to create a RDBMS login for the databases that you want to monitor. To create a RDBMS login for a database, perform the following steps:

1. Launch the IBM TotalStorage Productivity Center GUI.
2. Expand the **Administrative Services** node and navigate to the **License Keys** node under **Configuration**.
3. Select the **License Key** node.
4. On the right side pane, double-click **IBM TPC for Data - Databases**.
5. Select the **RDBMS Logins** tab.
6. Select **Add New**.
7. Complete the database information in the RDBMS Login Editor dialog. When completed, select **Save**. After selecting **Save**, the RDBMS Login Editor verifies the database before it creates a new RDBMS login entry.

The User ID that you enter in the **User** field must have database privileges as specified in Table 4-10.

*Table 4-10 RDBMS information*

| Database | User ID required privileges |
|---|---|
| Oracle® | CREATE SESSION<br>SELECT ANY DIRECTORY<br>ANALYZE ANY |
| Microsoft SQL Server | Must have *Permit* access |
| DS2 and Sybase | DBA privileges for the specified instance |

## 4.2.2 TPC for Databases default groups

After the installation of TPC for Data, you get the default groups in Table 4-11.

*Table 4-11 Predefined default groups for TPC for Databases*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring → Groups → Computer** | TPCUser.Default Computer Group | Every newly discovered computer is assigned to this group. The group is equal to **Data Manager → Monitoring → Groups → Computer → TPCUser.Default Group** (synchronous). |
| **Monitoring → Groups → Databases-Tablespaces** | TPCUser.Default Tablespace Group | Every newly discovered tablespace is assigned to this group. |
| **Monitoring → Groups → User** | TPCUser.Default User Group | Every newly discovered computer user is assigned to this group. The group is equal to **IBM TotalStorage Productivity Center → Data Manager → Monitoring → Groups → User → TPCUser.Default User Group** (synchronous). |

## 4.2.3 Probes

Run database probes to itemize and create inventory of the files, instances, logs, and objects that make up your enterprise's monitored RDBMSs. The results of the probe jobs are stored in the repository, and the results are used to supply the data that is necessary for viewing information about RDBMS storage.

Schedule database probe and scan jobs during times when user operations are at a minimum. Although it is possible to run many operations concurrently, the performance of any one operation decreases if there are other operations running at the same time. Schedule your database probe and scan jobs to avoid overlap.

*Table 4-12 Predefined default probes for TPC for Databases*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring → Probes** | TPCUser.Default Db Probe | Probes TPCUser.Default Computer Group once a day |

## 4.2.4 Scans

Run database scan jobs to collect statistics about the usage and trending of your actual storage consumption within an RDMBS. Scan jobs perform the majority of the work for the Data Manager for Databases by providing all of the data for usage reporting, as well as for quota analysis. The results of scans are stored in the repository, and the results used to supply the data that is necessary for capacity, usage, and usage violation reports.

*Table 4-13 Predefined default scans for TPC for Databases*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring → Scans** | TPCUser.Default Db Scan | The group TPCUser.Default Tablespace Group is assigned to it. Thus by default, all tablespaces are scanned once a day. All default profiles are assigned to this scan. |

Every scan has one or more profiles associated with it that define the data to filter and store in the TPC repository. You can view the collected information at **IBM TotalStorage Productivity Center → Data Manager for Databases → Reporting**. TPC for Databases has predefined profiles that meet most requirements. Table 4-14 lists the default profiles.

*Table 4-14 Predefined default profiles for TPC for Databases*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring → Profiles** | TPCUser.Db User Space | This default profile summarizes space usage by DB User (owner). By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. |
| **Monitoring → Profiles** | TPCUser.Largest Tables | This default profile gathers statistics about the *n* largest tables. (The default amount is `20`.) |
| **Monitoring → Profiles** | TPCUser.Segment Most Extents | This default profile gathers statistics about the *n* segments with the most extents. (The default amount is `20`, Oracle only.) |
| **Monitoring → Profiles** | TPCUser.Segment Most Unused Space | This default profile gathers statistics about the *n* segments with the most unused space. (The default amount is `20`, Oracle only.) |
| **Monitoring → Profiles** | TPCUser.Summary | This default profile summarizes space usage by database-tablespace. By default, every scan result is kept for seven days. The weekly scan result is kept for four weeks, and the monthly scan result is kept for three months. |

## 4.2.5 Alerts

TPC for Databases provides several alerts. The alerts differ from each database type. You can create alert conditions for one or all database types. Selecting **All Rdbms** provides the most limited choice of alert conditions. Each type of RDBMS has its own alert condition set. Figure 4-14 shows the Create Alert panel for instance alerts.

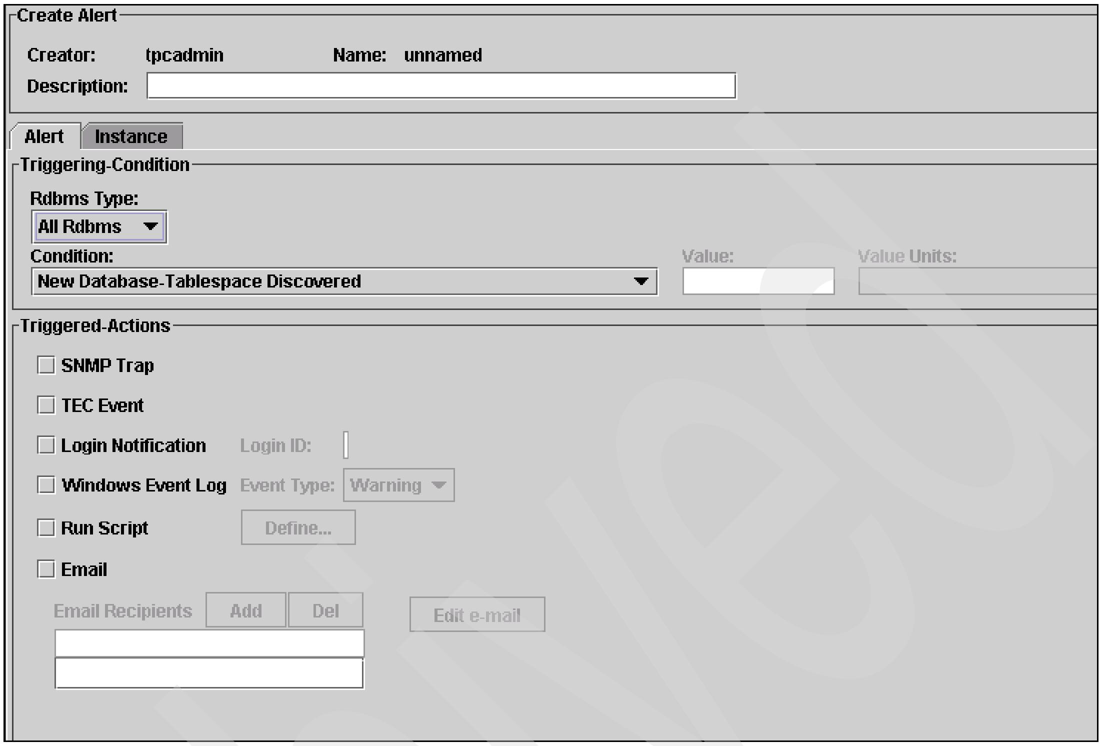

Creator: tpcadmin Name: unnamed

Rdbms Type: All Rdbms

Condition: New Database-Tablespace Discovered

☐ SNMP Trap

☐ TEC Event

☐ Login Notification

☐ Windows Event Log Event Type: Warning

☐ Run Script

☐ Email

*Figure 4-14 Create an Instance Alert*

The TPC Data Manager for Databases provides three level of alerts:

- Instance alerts
- Database-Tablespace alerts
- Table alerts

There are no default predefined alerts. You define alerts before or at the same time that you define probes and scans. You must define alerts before the probe and scan jobs run, because alerts are triggered based on the storage data that is collected by these monitoring jobs.

Unlike most TPC components, the Data Manager for Databases has its own Alert Log node as seen in Figure 4-15 on page 143.

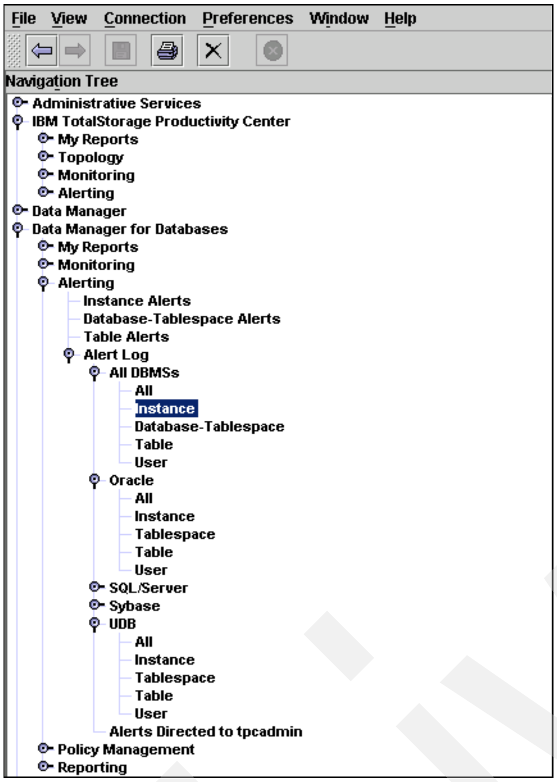

*Figure 4-15 Alert Log for Databases*

The selection of alerts differs from environment to environment. You can read a description of the alerts for each supported database in Appendix B of the manual, *IBM TotalStorage Productivity Center: User's Guide.*

## Policy management

Data Manager for Databases provides a quota management (**Data Manager for Databases** → **Policy Management**). It triggers the user and user group space usage within the databases. The User and User Groups consumption triggers or usage triggers can be defined at the following levels:

- Network (storage usage over several computers)
- Instance (storage usage within a database)
- Database - Tablespace (storage usage within a tablespace)

A quota is checked by aggregating the storage usage statistics that are stored in the repository. The quota checker job runs on the Data server and does not gather new statistics from any of the agent machines; the statistics that are used have already been gathered by separately scheduled scan jobs. For best results, schedule the quota checker to run after your scan jobs complete, which ensures that your quotas are checked against the most recent storage usage data.

# 4.3 Setting up monitoring and alerting for TPC for Disk

The first step in performance monitoring and alerting is to collect data in the TPC database repository.

## 4.3.1 Collecting data

TPC for Disk uses the CIM Agent as an interface to the disk subsystem. TPC uses the CIM Agent for discovery, probes, and performance monitoring. TPC also gets CIM indications from the CIM Agent. TCP creates an alert if the alert condition is violated.

TPC supports several storage subsystems, such as DS8000, ESS, and San Volume Controller (SVC). Table 4-15 shows naming relationships between TPC and storage devices. TPC uses synonyms, such as subsystem, storage pools, disk, or volume. For example, a storage pool is called an MDisk Group for an SVC; however, a storage pool is called a Rank for an ESS.

*Table 4-15 Naming relationships*

| Subsystem | Storage pool | Disk | Volume (LUN) |
|---|---|---|---|
| DS8000/DS6000™ | Extent pool | DDM | Volume |
| ESS | Rank | DDM | LUN |
| DS4000™ | Volume group | Disk | Volume (LUN) |
| SVC | MDisk group | MDisk | VDisk |

Use groups (**Disk Manager** → **Monitoring** → **Groups** → **Storage Subsystem**) to simplify adding new devices to the TPC monitors. Assign probes and especially alerts to groups. If you have groups and you need to add a new device, add the new device to the groups, and therefore, all alerts and probes are automatically assigned. TPC provides a default group (Table 4-16). You cannot use groups within performance monitors; within performance monitors, you have to assign each subsystem to one performance monitor.

In addition to discovering storage subsystems, the discovery job also tests the availability of the CIM Agents and queries the health information for the top-level devices. Because the discovery does not take very long, we recommend that you run the CIMOM discovery multiple times a day (for example, every four hours).

*Table 4-16 Default storage subsystem group*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring** → **Groups** → **Storage Subsystems** | TPCUser.Default Storage Subsystem Group | Every newly discovered subsystem is assigned to this group. |

**Note:** TPC supports indications to keep its repository up to date without any additional configuration effort. The currently available CIM Agents typically do not implement all needed indications, but they improve with every new CIM Agent. Therefore, not all events on devices are immediately reflected in TPC but might require a probe to complete. The Call Home functionality included in the subsystems might be sufficient for your environment. To avoid this situation, run the probe more frequently (perhaps every 12 hours) or configure the storage subsystem to directly send SNMP traps to a SNMP receiver, such as TEC.

## Probes

Probes examine the disk subsystem intensively. TPC is notified at the next probe of every change in the subsystem, for example, a new volume is created. Remember that a performance monitor can only report on probed entities. Therefore, run a probe frequently.

Probes collect a large amount of information. Therefore, do not run a probe more often than needed. If you use a concurrent alerting structure, which is immediately informing someone about problems, implement SNMP traps from the subsystem to a SNMP receiver, such as TEC. If the information delay is 12 hours, then set the probe interval to 12 hours and use the probe to catch the alerts that are missed by the indication from the subsystem. The default probe interval is 24 hours.

Because the probe process is intensive, shift the time that the probe is scheduled to run against each CIMOM, so that the probes do not run at the same time. A probe might take up to 20 minutes for each subsystem depending of the size of the subsystem, hardware configurations, and the current usage.

For small environments, use the following structure for an overview. For each subsystem, create a unique probe, similar to <*cimom_server_name*>_<*subsystem_name*>. In Table 4-16 on page 144, you can see that server `CIMOMServer01` manages Subsystem `DS8_01, DS8_02`, and `DS8_03`. The list is sorted alphabetically. You can run every probe on this CIM Agent with one hour between them.

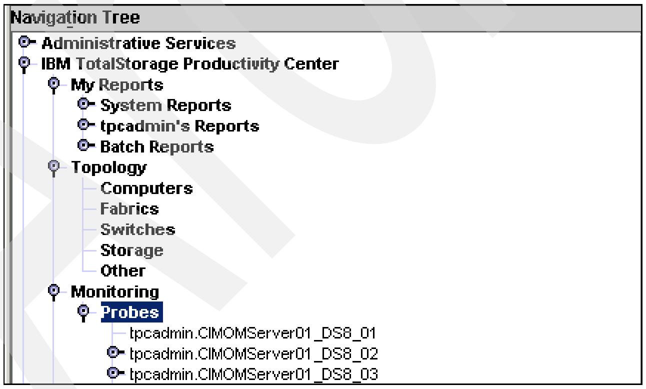

*Figure 4-16   Disk subsystem probes overview*

For large environments, it is more efficient to group the probes and subsystems. You might use the following convention to structure the group and probe names. The columns in Table 4-17 on page 146 show which DS8000 is registered to which CIMOM server. Two CIMOM servers are listed in the example. The last column shows the groupnames of the rows. Thus, the Group `DS8_GRP01` includes `DS8_01` and `DS8_05`. The group `DS8_GRP_02` includes `DS8_02` and `DS8_06`, and so on.

*Table 4-17 Subsystem group definitions for probes*

| CIMOMServer01 | CIMOMServer02 | CIMOMServer.. | Groupname |
|---|---|---|---|
| DS8_01 | DS8_05 | DS8_.. | DS8_GRP01 |
| DS8_02 | DS8_06 | DS8_.. | DS8_GRP02 |
| DS8_03 | DS8_07 | DS8_.. | DS8_GRP03 |
| DS8_04 | DS8_08 | DS8_.. | DS8_GRP04 |

Now, we can create the probe jobs. Every probe run is fixed to a unique time and in a fixed interval. Thus, you can use the following structure:

*TYPE_GROUPNAME_DAY_TIME_INTERVAL*

The probename gives information about the assigned storage subsystem group, which day the probe runs (`ANY`=Everyday), the start time, and the repeat time (12H=repeat probe every 12 hours).

In our example, the names of the probes are shown in Table 4-17.

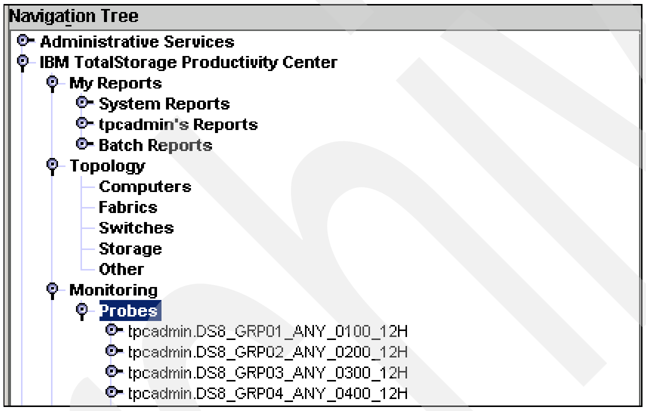

Navigation Tree

- Administrative Services
- IBM TotalStorage Productivity Center
  - My Reports
    - System Reports
    - tpcadmin's Reports
    - Batch Reports
  - Topology
    - Computers
    - Fabrics
    - Switches
    - Storage
    - Other
  - Monitoring
    - Probes
      - tpcadmin.DS8_GRP01_ANY_0100_12H
      - tpcadmin.DS8_GRP02_ANY_0200_12H
      - tpcadmin.DS8_GRP03_ANY_0300_12H
      - tpcadmin.DS8_GRP04_ANY_0400_12H

*Figure 4-17 Disk subsystem probes*

If you need to probe a new DS8000, add the subsystem to Table 4-17 and, therefore, into a subsystem group.

### Performance monitors

You can define performance monitors by selecting in the Navigation Tree **Disk Manager** → **Monitoring** → **Subsystem Performance Monitor**.

You can run performance monitors at intervals from five minutes to indefinitely. You can define the interval from five minutes to 60 minutes. The interval length depends on your requirements. To get the most information out of a performance monitor, use the lowest time interval. Otherwise, the subsystem performance monitoring loses peak loads. For DS8000, DS6000, DS4000, and ESS, you can use the interval cycle of five minutes; therefore, the performance data is based on a five minute average load. For SVC prior to Version 4.1, the minimum time interval is 15 minutes, because TPC uses the SVC internal performance monitor and this is the limit for this monitor. However, the minimum interval is five minutes if you use SVC Version 4.1 and TPC Version 3.1.3. You can assign each subsystem to exactly one performance monitor.

You can run a performance monitor indefinitely or for a defined period. A performance monitor that runs indefinitely displays less information in the job log during the run. Therefore, it is still a good practice to restart the performance monitor periodically. The following performance monitor example in Figure 4-18 starts the performance monitor every Sunday at `12:00` and runs until Sunday at `11:00`. In that example, we lose one hour of performance samples every Sunday from `11:00` until `12:00`.

You also can restart the performance monitor daily or monthly if required, but you always lose one hour of performance samples.

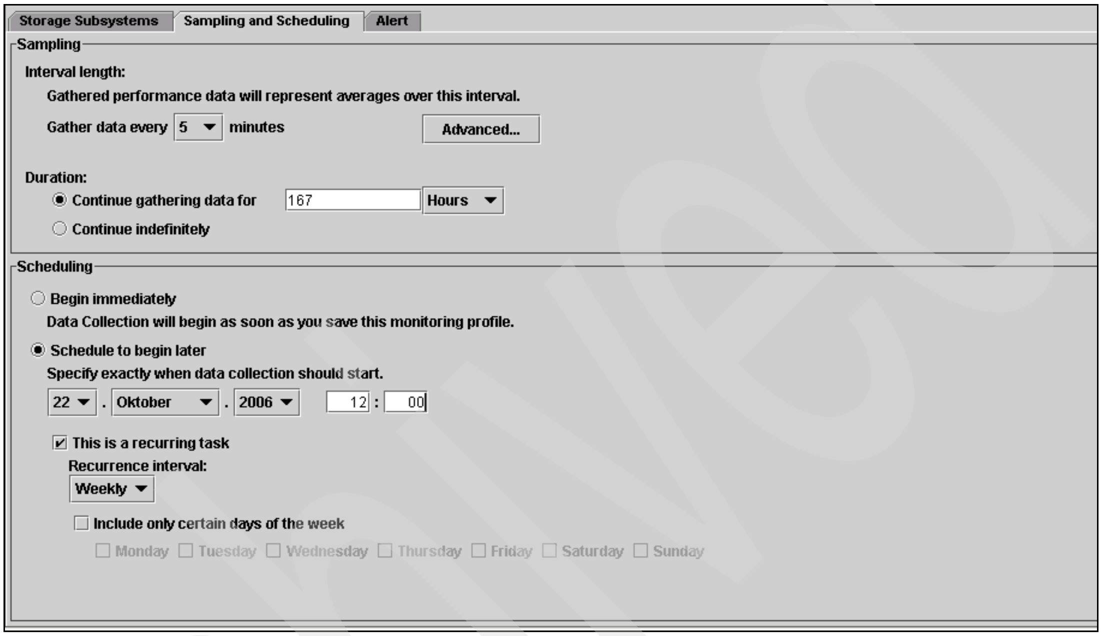

Sampling

Interval length:

Gathered performance data will represent averages over this interval.

Gather data every 5 minutes

Duration:

● Continue gathering data for 167 Hours

○ Continue indefinitely

Scheduling

○ Begin immediately

Data Collection will begin as soon as you save this monitoring profile.

● Schedule to begin later

Specify exactly when data collection should start.

22 . Oktober . 2006 12 : 00

☑ This is a recurring task

Recurrence interval:

Weekly

□ Include only certain days of the week

□ Monday □ Tuesday □ Wednesday □ Thursday □ Friday □ Saturday □ Sunday

*Figure 4-18 Performance Monitor weekly restart*

To be sure that the performance monitor collects performance data on all of your volumes, you can periodically check the time of your performance data collection. Use the report **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **By Volume** to verify the time of the last performance samples. If you have a lot of volumes, you can define a filter to show only the volumes with performance samples older than, for example, two days.The **Filter** function is on the same panel as the **Generate Report** function. Figure 4-19 on page 148 shows this filter. If you find volumes with old samples (older than the last probe), check your performance monitor log file. You might stop the performance monitor, run a probe, and restart the performance monitor.

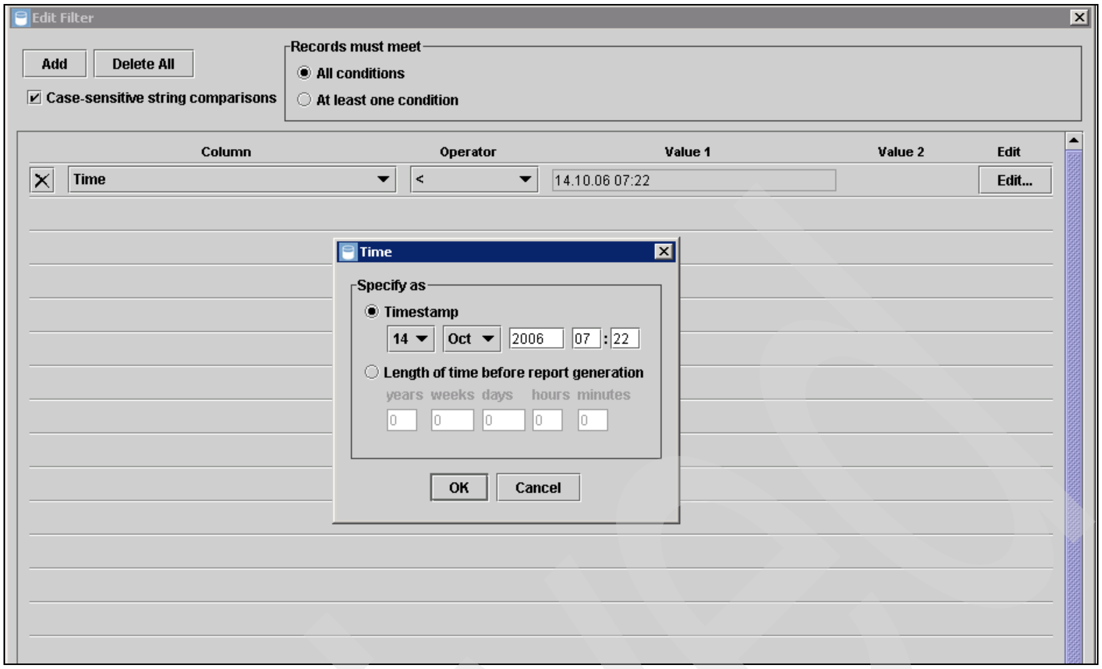

Edit Filter

Add | Delete All

☑ Case-sensitive string comparisons

Records must meet

● All conditions

○ At least one condition

| Column | Operator | Value 1 | Value 2 | Edit |
|---|---|---|---|---|
| Time | < | 14.10.06 07:22 | | Edit... |

Time

Specify as

● Timestamp

14 | Oct | 2006 | 07 : 22

○ Length of time before report generation

years weeks days hours minutes

0 0 0 0 0

OK | Cancel

Figure 4-19 Filter volumes with a time value

Within default settings, the performance monitor collects the performance data regardless of whether the samples are used for alerting or statistics. TPC provides the **Advanced** function (Figure 4-20) beside the interval definition (**Disk Manager** → **Monitoring** → **Subsystem Performance Monitor** → **Sampling and Scheduling**). This setting gives you the capability to store less performance sample data (for example, only every 15 minutes) than if you compare the performance data with the performance thresholds (for example, every five minutes).

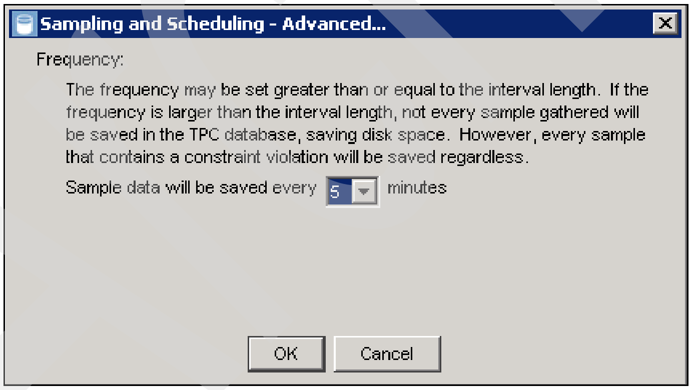

Sampling and Scheduling - Advanced...

Frequency:

The frequency may be set greater than or equal to the interval length. If the frequency is larger than the interval length, not every sample gathered will be saved in the TPC database, saving disk space. However, every sample that contains a constraint violation will be saved regardless.

Sample data will be saved every 5 minutes

OK | Cancel

Figure 4-20 Sampling and Scheduling - Advanced

You can only assign a storage subsystem to a performance monitor and not a group. Therefore, you might use following naming convention for a small environment:

*<CIMOM_SERVER_NAME>_<SUBSYSTEM_NAME>*

For a large environment, use a naming convention such as this one:

*<TYPE>_<DAY>_<TIME>_<INTERVAL>*

In our example, the performance monitor is named to `DS8_ANY_0112_05M`. `DS8` describes the Type DS8000, `ANY` means that the performance monitor runs every day during the week, `0112` is the start time of the performance monitor, and `05M` is the collection interval of five minutes.

## Alerting

You can define alerts by selecting **Disk Manager** → **Alerting** → **Storage Subsystem Alerts**.

TPC for Disk divides alerts in three categories:

- State changes alert for a device element, for example, a disk changes the status from `OK` to `ERROR`
- Configuration change alerts for a configuration of a subsystem that has changed, for example, `Cache Increased`
- Performance alerts for performance threshold violations

Table 4-18 and Table 4-19 on page 150 show the predefined alerts or monitoring groups for TPC for Disk.

*Table 4-18 Predefined default monitoring values for Disk Manager*

| Where to find | Condition | Description |
|---|---|---|
| **Monitoring** → **Groups** → **Storage Subsystem** | TPCUser.Default Storage Subsystem Group | Every newly discovered storage subsystem is assigned to this group. |
| **Alerting** → **Storage Subsystem Alerts** | TPCUser.Default Disk Array Discovery | If a new storage subsystem is discovered, it creates an alert. |

Table 4-19 on page 150 shows an extract of possible alert conditions. This is not a complete list and can differ between environments. This table gives you an idea or a starting point for defining alert conditions. You can get a complete list in Appendix B of the manual *IBM TotalStorage Productivity Center User's Guide*, GC32-1775.

*Table 4-19   Status change alerts TPC for Disk*

| Where to find | Condition | Description |
|---|---|---|
| **Alerting → Storage Subsystem Alerts** | New disk detected | A new disk is detected. For SVC, this is an MDisk. |
| **Alerting → Storage Subsystem Alerts** | Disk not found | A disk is not found. For SVC, this is an MDisk. |
| **Alerting → Storage Subsystem Alerts** | Storage subsystem not found | A previously detected storage subsystem is not found. |
| **Alerting → Storage Subsystem Alerts** | Subsystem status change offline | A subsystem goes offline. |
| **Alerting → Storage Subsystem Alerts** | Subsystem status change online | A subsystem comes online. |
| **Alerting → Storage Subsystem Alerts** | Subsystem version change | A subsystem version changes. |
| **Alerting → Storage Subsystem Alerts** | Back-end controller status change offline | A back-end controller goes offline (SVC). |
| **Alerting → Storage Subsystem Alerts** | Back-end controller status change online | A back-end controller comes online (SVC). |
| **Alerting → Storage Subsystem Alerts** | Volume status change offline | A volume goes offline. For SVC, this is a VDisk. |
| **Alerting → Storage Subsystem Alerts** | Pool status change | A pool is missing or rediscovered. For SVC, the pool is an MDisk Group. |
| **Alerting → Storage Subsystem Alerts** | Pool discovered | A pool is discovered. For SVC, the pool is an MDisk Group. |
| **Alerting → Storage Subsystem Alerts** | Pool status change offline | A pool goes offline. For SVC, the pool is an MDisk Group. |
| **Alerting → Storage Subsystem Alerts** | Pool status change online | A pool comes online. For SVC, the pool is an MDisk Group. |
| **Alerting → Storage Subsystem Alerts** | Node status change offline | A node goes offline (SVC). |
| **Alerting → Storage Subsystem Alerts** | Node status change online | A node comes online (SVC). |
| **Alerting → Storage Subsystem Alerts** | Node state change | An node is missing or rediscovered (SVC). |

We recommend that you use reporting for performance analyzing and capacity planning. You can use performance alerts for identifying a performance problem within your environment. In this case, we recommend that you use high alert thresholds (see Figure 4-21 on page 151); otherwise, you might be overwhelmed with alerts. A good strategy is to start with a high performance alerting threshold, solve performance bottlenecks, and lower the thresholds over time to an accurate value. Keep in mind that the storage subsystem-attached hosts and the associated applications determine a valid subsystem performance.

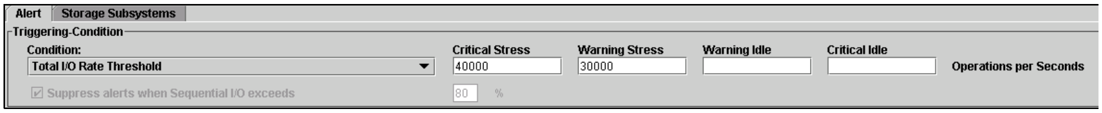

Alert | Storage Subsystems

Triggering-Condition

Condition: Total I/O Rate Threshold

Critical Stress: 40000

Warning Stress: 30000

Warning Idle:

Critical Idle:

Operations per Seconds

☑ Suppress alerts when Sequential I/O exceeds 80 %

*Figure 4-21 Performance threshold definitions*

We need to address a few points in order for you to understand threshold settings for a performance alert:

- There are two types of boundaries for each threshold: the upper boundary (stress) and the lower boundary (idle). When a metric's value exceeds the upper boundary or falls below the lower boundary, it triggers alerts.
- There are two levels of alerts: *warning* and *critical*. The combination of boundary type and level type generates four threshold settings: critical stress, warning stress, warning idle, and critical idle. Most threshold values are in descending order (critical stress has the highest value that indicates high stress on the device, and critical idle has the lowest value). Cache Holding Time is the only threshold in ascending order.
- If you are only interested in receiving alerts for certain boundaries, leave the other boundaries blank. The collection engine only checks boundary conditions with input values; therefore, no alerts are sent for conditions that are blank.

Figure 4-22 illustrates those boundaries with a performance graph. Each quadrant represents a performance sample and can create an alert if the sample violates a threshold.

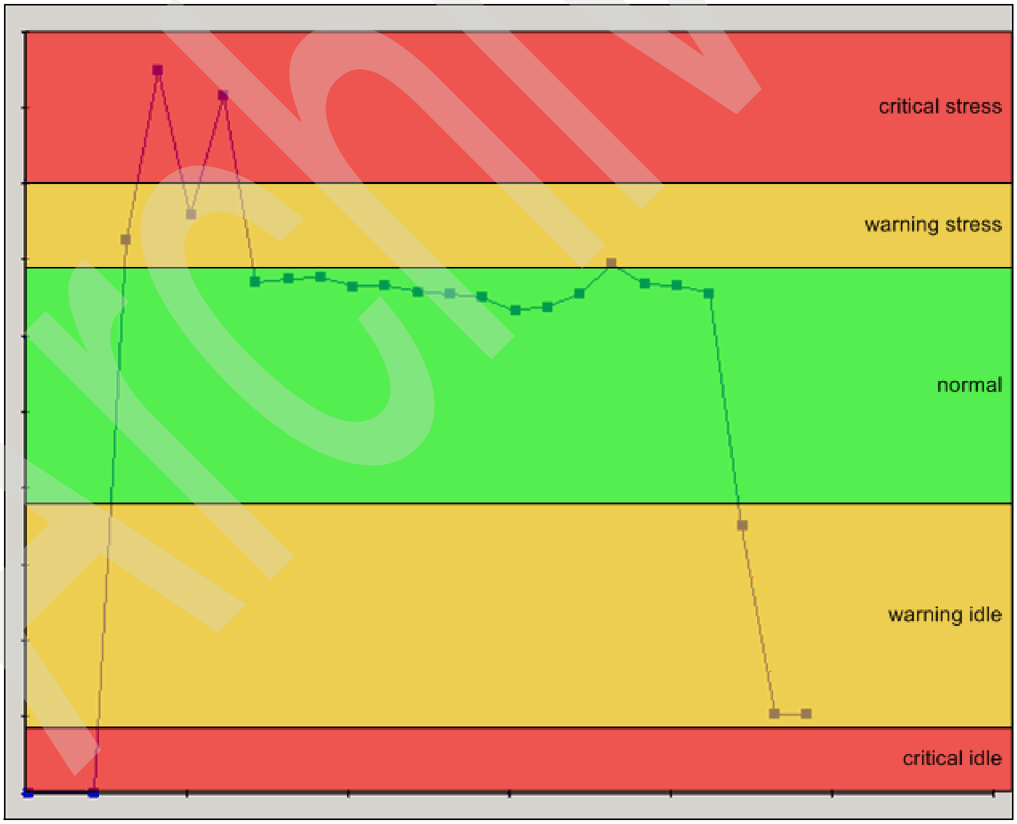

*Figure 4-22 Performance boundaries*

Every performance threshold violation is indicated as a constraint violation and is stored in the table T_PMM_EXCEPTION. Display these constraint violations in the report you display by selecting **IBM Total Storage Productivity Center → Disk Manager → Reporting → Constraint Violation**. The reports give you support to analyze the constraint violation. Refer to 4.10.5, “Disk Manager reports” on page 191 for an example.

The following constraints exist:

- Overall port response time threshold
- Total port data rate threshold
- Total port I/O rate threshold
- Cache holding time threshold
- NVS full percentage threshold
- Total data rate threshold
- Total I/O rate threshold
- Overall back-end response time threshold
- Total back-end data rate threshold
- Total back-end I/O rate threshold
- Disk utilization percentage threshold

Table 4-20 shows the available performance threshold alert conditions. You can expand the available alerts with future TPC versions. For some conditions, default thresholds exist.

*Table 4-20 Performance threshold alerts for storage subsystems*

| Where to find | Condition | Subsystem | Component | Explanation | Threshold |
|---|---|---|---|---|---|
| **Alerting → Storage Subsystem Alerts** | Disk utilization percentage | ESS, DS6000, and DS8000 | Array | Threshold on the arrays in a subsystem. The average percent of time that the disks associated with the array were busy. | Default thresholds are 80%, 50%, -1, and -1. For DS6000 and DS8000 subsystems, this threshold applies only to those ranks which are the only ranks in their associated extent pool. |
| **Alerting → Storage Subsystem Alerts** | NVS full | ESS, DS6000, and DS8000 | Controller | Threshold on the percentage of time that NVS space constraints caused I/O operations to be delayed for the subsystem controllers (clusters) | Default thresholds are 10%, 3%, -1, and -1. |

| Where to find | Condition | Subsystem | Component | Explanation | Threshold |
|---|---|---|---|---|---|
| **Alerting → Storage Subsystem Alerts** | Cache holding time | ESS, DS6000, and DS8000 | Controller | Threshold on the average cache holding time in seconds. Shorter time periods indicate adverse performance. The cache holding time metric for each controller is checked against the threshold boundaries of each collecting interval. | Default thresholds are 30 seconds, 60 seconds, -1,and -1. |
| **Alerting → Storage Subsystem Alerts** | Overall back-end response time | SVC, DS8000, DS6000, and ESS | MDisk (SVC), MDisk Group (SVC), subsystem, controller, and array | Average number of milliseconds that it took to service each I/O operation (read and write) for a particular component over a time interval. For SVC, this is the external response time of the MDisks. | No default |
| **Alerting → Storage Subsystem Alerts** | Total back-end data rate threshold | SVC, DS8000, DS6000, and ESS | MDisk (SVC), MDisk Group (SVC), subsystem, controller, and array | Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations. | No default |
| **Alerting → Storage Subsystem Alerts** | Total back-end I/O rate threshold | SVC, DS8000, DS6000, and ESS | MDisk (SVC), MDisk Group (SVC), subsystem, controller, and array | Average number of I/O operations per second for read and write operations. | No default |
| **Alerting → Storage Subsystem Alerts** | Total data rate threshold | SVC, DS8000, DS6000, and ESS | VDisk (SVC), I/O Group (SVC), MDisk Group (SVC), subsystem, controller, array, and volume | Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations. | No default |

| Where to find | Condition | Subsystem | Component | Explanation | Threshold |
|---|---|---|---|---|---|
| **Alerting → Storage Subsystem Alerts** | Total I/O rate threshold (overall) | SVC, DS8000, DS6000, and ESS | VDisk (SVC), I/O Group (SVC), MDisk Group (SVC), subsystem, controller, array, or volume | Average number of I/O operations per second for both sequential and nonsequential read and write operations. | No default |
| **Alerting → Storage Subsystem Alerts** | Overall port response time threshold | DS8000, DS6000, and ESS | Port | Average number of milliseconds that it took to service each operation (send and receive) for a particular port over a time interval. | No default |
| **Alerting → Storage Subsystem Alerts** | Total port data rate threshold | DS8000, DS6000, and ESS | Port | Average number of megabytes (2^20 bytes) per second that were transferred for send and receive operations for a particular port over a time interval. | No default |
| **Alerting → Storage Subsystem Alerts** | Total port I/O rate threshold | DS8000, DS6000, and ESS | Port | Average number of I/O operations per second for send and receive operations for a particular port over a time interval. | No default |

Many things influence how you set accurate performance thresholds. Analyzing the performance reports from your environment can help you. Also, check the vendor manuals and this Web site for useful information:

http://www.storageperformance.org

## 4.4 Setting up monitoring and alerting for TPC for Fabric

In this section, we discuss the basics for monitoring and setting up alerts for IBM TotalStorage Productivity Center for Fabric (TPC for Fabric).

## 4.4.1 Data collection

TPC for Fabric uses multiple methods to collect data from the SAN. The methods of data collection and the protocols used are:

- In-band Fabric agent - Fibre Channel Generic Services 3 standard (FC GS-3)
- Out-of-band Fabric agent - SNMP, except for Brocade switches, which use the Brocade API
- CIM Agent - Storage Management Initiative Specification (SMI-S)

Due to varying support provided by different protocols and vendors, you need to rely on a combination of various agents to fully support all of the functions of TPC for Fabric. Table 4-21 is a description of the information that each type of agent provides.

*Table 4-21 SAN agents*

| In-band Fabric agent | Out-of-band Fabric agent | CIM Agent |
|---|---|---|
| Topology and identification information (some switch attributes can only be gathered by out-of-band Fabric agents) | Information about the fabric by querying the switch for topology information | Switch and port diagnostic and performance information |
| Host level information for the local system | Zoning information and zone control (only Brocade) | N/A |
| Host Bus Adapter (HBA) information from the local system (vendor, model, and driver version) | Virtual SAN information (only Cisco) | N/A |
| Zoning information and zone control (except Brocade) | N/A | N/A |
| Events detected by the HBA on the local system | N/A | N/A |

## 4.4.2 SNMP trap notification

Another aspect of the SNMP configuration includes trap notification. The switch generates SNMP traps and directs them to the TPC Device server as an indication that something in the fabric has changed. The default configuration for handling switch traps is to send them from the switch to port 162 on the TPC Device server.

SNMP traps are synonymous to the events that are received by the in-band Fabric agent. For redundancy and if possible in your environment, we recommend as the best practice to use SNMP traps and in-band Fabric agents to receive events. After the event is received from the SAN, the device server issues queries to the agents pertaining to the event and updates the Fabric Manager data contained in the database. If any changes are detected and if alerts are configured for the change detected, an alert is issued.

## 4.4.3 Probes

You can define probes for Fabric at **IBM TotalStorage Productivity Center** → **Monitoring** → **Probes**.

TPC is configured to provide real-time monitoring of the SAN after you have successfully configured discovery and have in-band Fabric agents, or configured the switches to send

SNMP traps, or both. However, there are some changes within a SAN that might not generate an in-band event or SNMP trap, such as an attribute for a device changing. In order to collect those events and maintain updated fabric information, we recommend that you probe your fabric on a regular basis.

Depending on the size of your environment, the number of fabrics, and the number of devices within each fabric, you might want to have more than one probe defined to collect the fabric data. This helps you to load balance the amount of work that the Device server and the SAN perform at a given time. You can view the job status for a probe to determine how long the probe takes to complete by subtracting the Start Time from the Finish Time.

Probes examine the Fabric and switches intensively. Therefore, do not use probes more often than needed. A probe runs using in-band Fabric and out-of-band Fabric agents. Every change in the fabric, for example, a `new port blade`, is reported to TPC.

A probe interval might be 24 hours or less. You can define a probe run to one fabric or to a fabric group. We recommend that you schedule fabric probes at alternating intervals, not at the same time.

### Performance monitors

You can define the performance monitor for a SAN switch at **Fabric Manager** → **Monitoring** → **Switch Performance Monitors**.

For SAN switches, you can use an interval cycle from five minutes to 60 minutes. Use the lowest time interval to get the most precise performance monitor; otherwise, the switch performance monitoring loses peak loads. In addition to performance samples, the performance monitor collects error counters, such as *Link Failure Rate Threshold* and *Error Frame Rate Threshold*.

You can set up a performance monitor naming convention for a small environment in this way:

*CIMAgentServername_SWITCHNAME_INTERVAL*

In our example, the performance monitor is named `CIMAgentServer01_SWITCH01_05M`.

For large environments, you can set up a performance naming convention this way:

*<TYPE>_<DAY>_<TIME>_<INTERVAL>*

> **Note:** For Brocade SAN switches, the CIM Agent needs at least one proxy switch per Fabric. Over those proxy switches, the CIM Agent collects performance samples. Therefore, TPC is able to run performance monitors, although the Brocade switches in the Fabric are not registered at the Brocade CIM Agent.

## 4.4.4 Alerting

Define alerts for TPC for Fabric at **Fabric Manager** → **Alerting**.

TPC for Fabric divides alerts into three categories:

- Fabric alerts
- Switch alerts
- Endpoint device alerts

Table 4-22 shows the predefined alerts or monitoring groups for TPC for Fabric. Because normally, new fabrics and switches are not often discovered, you can leave these settings. For large fabrics, you can remove the *Endpoint Device Alert Discovery* to avoid getting too many alerts.

*Table 4-22 Predefined default monitoring values for Fabric Manager*

| Where to find | Variable | Description |
|---|---|---|
| **Monitoring → Groups → Fabric** | TPCUser.Default Fabric Group | Every newly discovered Fabric is assigned to this group. |
| **Alerting → Fabric Alerts** | TPCUser.SNMP Discovery | If a new fabric is discovered, it creates an alert. |
| **Alerting → Switch Alerts** | TPCUser.SNMP Discovery | If a new switch is discovered, it creates an alert. |
| **Alerting → Endpoint Device Alerts** | TPCUser.SNMP Discovery | If a new endpoint is discovered, it creates an alert. |

Table 4-23, Table 4-24 on page 158, and Table 4-25 on page 158 show an extract of possible state change alert conditions for fabric, switches, and endpoints. This is not a complete list and might differ between environments. This list gives you an idea or a starting point for defining alert conditions. For example, zoning alerts, such as `Zone to Zone Alias Change` or `Zone Set to Zone Change`, might be useless for your environment but might be useful if you want trace the changes in your SAN zonings.

*Table 4-23 State change alerts for Fabric*

| Where to find | Condition | Description |
|---|---|---|
| **Alerting → Fabric Alerts** | Fabric state change | A fabric is missing or rediscovered. |
| **Alerting → Fabric Alerts** | Fabric goes offline | A fabric goes offline. |
| **Alerting → Fabric Alerts** | Fabric goes online | A fabric comes online. |
| **Alerting → Fabric Alerts** | Fabric connection state change | A connection is missing or rediscovered. |
| **Alerting → Fabric Alerts** | Fabric to switch change | A switch to fabric association is discovered, rediscovered, or missing. |

*Table 4-24 State change alerts for switches*

| Where to find | Condition | Description |
|---|---|---|
| **Alerting → Switch Alerts** | Switch state change | A switch is missing or rediscovered. |
| **Alerting → Switch Alerts** | Switch property changes | A switch port is discovered, missing, or rediscovered. |
| **Alerting → Switch Alerts** | Switch status change offline | A switch goes offline. |
| **Alerting → Switch Alerts** | Switch status change online | A switch comes online. |
| **Alerting → Switch Alerts** | Switch blade change | A switch blade is discovered, missing, or rediscovered. |
| **Alerting → Switch Alerts** | Switch blade change offline | A switch blade goes offline. |
| **Alerting → Switch Alerts** | Switch blade change online | A switch blade comes online. |

*Table 4-25 State Change alerts for endpoints*

| Where to find | Condition | Description |
|---|---|---|
| **Alerting → Endpoint Alerts** | Endpoint state change | An endpoint is missing or is rediscovered. |

## Performance alert threshold settings

There are a few points that we need to address in order for you to understand threshold settings for a performance alert. See Figure 4-23.

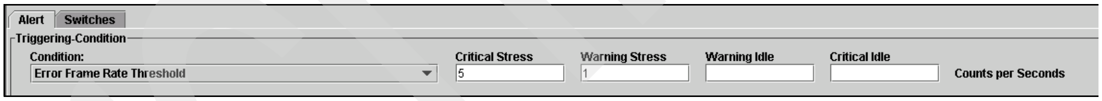

Triggering-Condition

| Condition: | Critical Stress | Warning Stress | Warning Idle | Critical Idle | |
|---|---|---|---|---|---|
| Error Frame Rate Threshold | 5 | 1 | | | Counts per Seconds |

*Figure 4-23 Threshold switch performance monitor*

Points to consider:

- There are two types of boundaries for each threshold: the upper boundary (stress) and lower boundary (idle). When a metric's value exceeds the upper boundary or falls below the lower boundary, it triggers alerts.
- There are two levels of alerts: warning and critical. The combination of boundary type and level type generates four threshold settings: critical stress, warning stress, warning idle, and critical idle. Most threshold values are in descending order (critical stress has the highest value that indicates high stress on the device, and critical idle has the lowest value). Cache Holding Time is the only threshold in ascending order.
- If you want to receive alerts for certain boundaries, leave the other boundaries blank. The collection engine only checks boundary conditions with input values; therefore, there are no alerts sent for blank conditions.

Figure 4-24 on page 159 illustrates those boundaries with a performance graph. Each gray quadrant represents a performance sample and can create an alert if the sample violates a threshold.

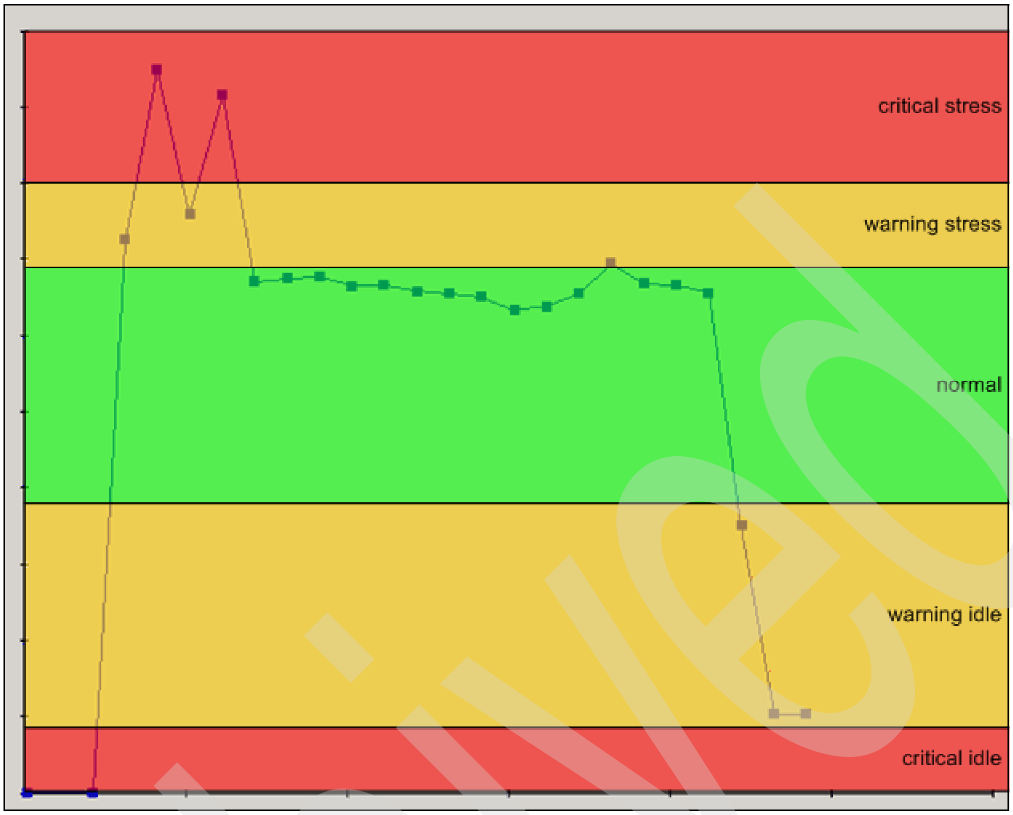

*Figure 4-24 Performance boundaries*

Table 4-26 on page 160 shows an extract of possible performance threshold alert conditions. This is not a complete list and might differ between environments. We recommend that you configure the *Link Failure Rate Threshold* and the *Error Frame Rate Threshold* alerts, because they provide predictive analysis for your SAN. These errors are indications of potential link failures in the SAN. An accurate threshold setting can vary from environment to environment. We do not have a good recommendation on the threshold settings. We recommend that you use trial and error with those thresholds in your environment to determine a proper setting.

Table 4-26 Performance thresholds for SAN devices

| Where to find | Condition | Explanation | Threshold |
|---|---|---|---|
| **Alerting → Switch Alerts** | Total port data rate threshold | A total port data rate threshold is reached. | You can only set total data rate (read and write) thresholds. This threshold can give you an idea which ports are most utilized. Therefore, a 2 GBps switch settings can be 300, 200, -1, or -1. |
| **Alerting → Switch Alerts** | Link failure rate threshold | A link failure rate threshold is reached. | We do not have an accurate threshold recommendation. Create a report of Link Failures. If no failure is recognized, start with these settings: 0, 0, -1,and -1. |
| **Alerting → Switch Alerts** | Error frame rate threshold | An error frame rate threshold is reached. | Use the same procedure as mentioned at Link failure rate threshold. |
| **Alerting → Switch Alerts** | Total port packet rate threshold | A total port packet rate threshold is reached. | You can only set total packet rate (read and write) threshold. |

## 4.5 Setting up monitoring and alerting for TPC for Tape

There are a few monitoring tasks available at the current version of TPC Version 3.1.2. TPC reports on the following tape library information as shown in Table 4-27:

Table 4-27 Tape reporting

| Component | What is reported |
|---|---|
| Library | Status, number of drives, number of changers, number of cartridges, number of slots, number of I/O ports, model, vendor, description, owner, contact, firmware version, element manager, lock present, locked state, security breach state, manually entered name, and three properties that you can define. |
| Drive | Status, needs cleaning, number of mounts, WWNN, firmware version, and location. |
| Changer | Status, media-flip supported, WWNN, and firmware version. Note for 3584, these really are the logical partitions, not the physical accessors. |
| I/O port | Extended state (whether accessible by an operator), description, and location. |
| Cartridge | Label, capacity, type, whether cleaner media, whether dual-sided, location, and media description. |

### 4.5.1 Alerting

We currently do not define TPC alerts for tape events. We only support alerts that are delivered to TPC through the SMI-S indication mechanism. Specifically, these are all SNMP tape alerts that are described in the *IBM TotalStorage UltraScalable Tape Library 3584 Planning and Operator Guide,* GA32-0408.

## 4.6 How long to retain monitoring and alerting data

During monitoring jobs, TPC collects a large amount of data. The performance monitor job especially collects a large amount of data. Consequently, the TPC data repository (database) increases in size with time. Therefore, TPC provides a simple Data Lifecycle Management configuration.

For Data Retention, TPC categorizes its data into:

- Job log files
- Alerts
- Statistical data (including performance data)

### 4.6.1 Job log files

Every task, such as performance monitors, probes, discovery, pings, scans, and history aggregator, creates a log file for each run. You configure the log file retention at **Administrative Services → Configuration → Log-File Retention**. Figure 4-25 shows the default configuration. The last five job run log-files are kept for a maximum of 90 days. You might only use the maximum of 90 days if you rerun a job, such as a continuously running performance monitor. Changes on the configuration go into effect at the next run of the history aggregator.

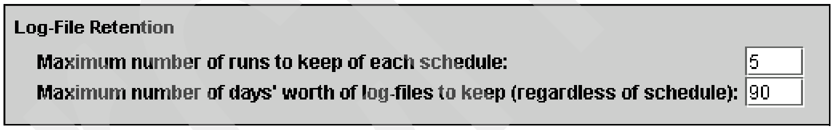

Log-File Retention

Maximum number of runs to keep of each schedule: 5

Maximum number of days' worth of log-files to keep (regardless of schedule): 90

*Figure 4-25 Log-File Retention*

### 4.6.2 Alerts

Every threshold violation, external indication, or job error saves an alert to the TPC repository.

In the **Alert Log Disposition** panel (**Administrative Services → Configuration → Alert Disposition**), you define how long to keep the alert logs as shown in Figure 4-26 on page 162. This value affects all alert logs, independently of whether a log has been *cleared* (acknowledged) or not. You can read the alert logs at **IBM TotalStorage Productivity Center → Alerting**. The default configuration is to keep alert logs for 90 days. Large environments with many alerts might decide to keep the alert log for fewer than 90 days, or just *delete* unwanted alerts manually.

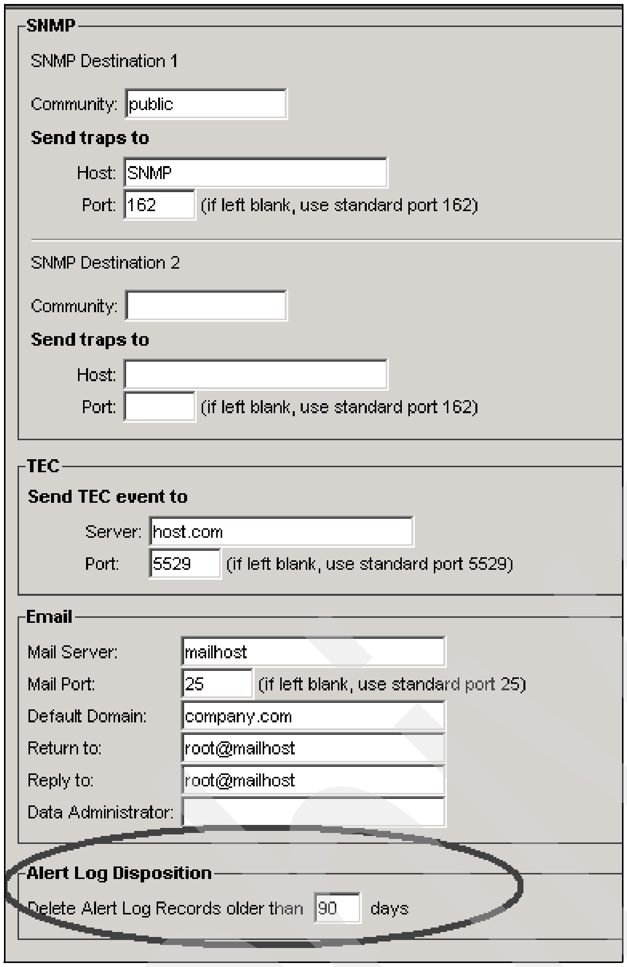

**SNMP**

SNMP Destination 1

Community: public

**Send traps to**

Host: SNMP

Port: 162 (if left blank, use standard port 162)

SNMP Destination 2

Community:

**Send traps to**

Host:

Port: (if left blank, use standard port 162)

**TEC**

**Send TEC event to**

Server: host.com

Port: 5529 (if left blank, use standard port 5529)

**Email**

Mail Server: mailhost

Mail Port: 25 (if left blank, use standard port 25)

Default Domain: company.com

Return to: root@mailhost

Reply to: root@mailhost

Data Administrator:

**Alert Log Disposition**

Delete Alert Log Records older than 90 days

*Figure 4-26 Alert Log Disposition*

## 4.6.3 Statistical data

Statistical data is most likely the largest and most varied amount of data in the TPC repository, and therefore, influences the size and performance of the repository database.

The resource history retention that you can see in Figure 4-27 on page 163 controls how long to keep the statistical information. By specifying a number for the days, the weeks, or the months for each element on this window, you can control the amount of data that is retained and available for historical analysis and charting. The longer you keep the data, the more informative your analysis is over time.

When these settings are too high, they can have a negative effect on the amount of time that it takes to generate reports. If report generation starts to slow down to an unacceptable level, this might be due to high resource history retention settings.

> **Note:** If you do not select a Performance Monitor check box, the data is never discarded and generates a great amount of data in the database. The same result occurs if you select a Performance Monitor check box and specify “0” for the number of days to keep.

Figure 4-27 shows the default configuration. Adjust those values according to your specific environment. Refer to Chapter 3, “IBM Total Productivity Center database considerations” on page 61 to calculate the effect to the TPC database.

*Sample data* is the data that is collected at the specified interval length of the monitor, for example, you can collect performance data every five minutes. The default retention period for sample data is 14 days, which means that by default, TPC keeps the individual five minute samples for 14 days before TPC purges the samples. TPC summarizes the Individual samples into hourly and daily data. For example, TPC saves the sum of 12 of the five minute samples as an hourly performance data record, and TPC saves the sum of 288 of these samples as a daily performance data record. The default retention periods for hourly and daily data are 30 days and 90 days, respectively.

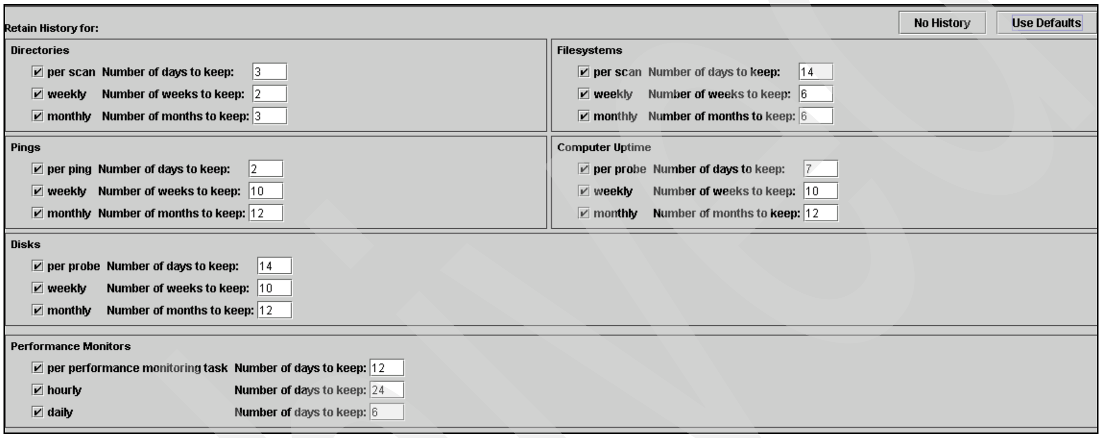

Retain History for:

Directories
☑ per scan Number of days to keep: 3
☑ weekly Number of weeks to keep: 2
☑ monthly Number of months to keep: 3

Filesystems
☑ per scan Number of days to keep: 14
☑ weekly Number of weeks to keep: 6
☑ monthly Number of months to keep: 6

Pings
☑ per ping Number of days to keep: 2
☑ weekly Number of weeks to keep: 10
☑ monthly Number of months to keep: 12

Computer Uptime
☑ per probe Number of days to keep: 7
☑ weekly Number of weeks to keep: 10
☑ monthly Number of months to keep: 12

Disks
☑ per probe Number of days to keep: 14
☑ weekly Number of weeks to keep: 10
☑ monthly Number of months to keep: 12

Performance Monitors
☑ per performance monitoring task Number of days to keep: 12
☑ hourly Number of days to keep: 24
☑ daily Number of days to keep: 6

*Figure 4-27 Resource history retention*

TPC for Databases has its own resource history retention configuration. You can see the configuration pane at **Administrative Services** → **Configuration** → **Resource History Retention for Databases**. Figure 4-28 shows the default settings.

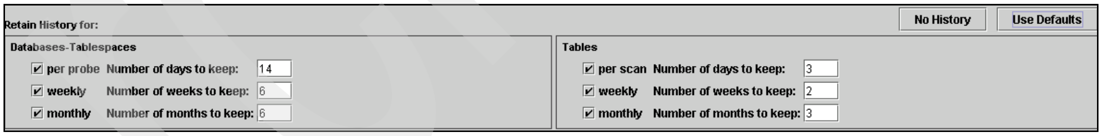

Retain History for:

Databases-Tablespaces
☑ per probe Number of days to keep: 14
☑ weekly Number of weeks to keep: 6
☑ monthly Number of months to keep: 6

Tables
☑ per scan Number of days to keep: 3
☑ weekly Number of weeks to keep: 2
☑ monthly Number of months to keep: 3

*Figure 4-28 Resource history retention for databases*

TPC for Data has an additional resource history retention configuration. You can see this setting in each profile (**Data Manager** → **Monitoring** → **Profiles**). Profiles are assigned to scans. Figure 4-29 on page 164 shows an example. You can set the retention period per sample, week, and month.

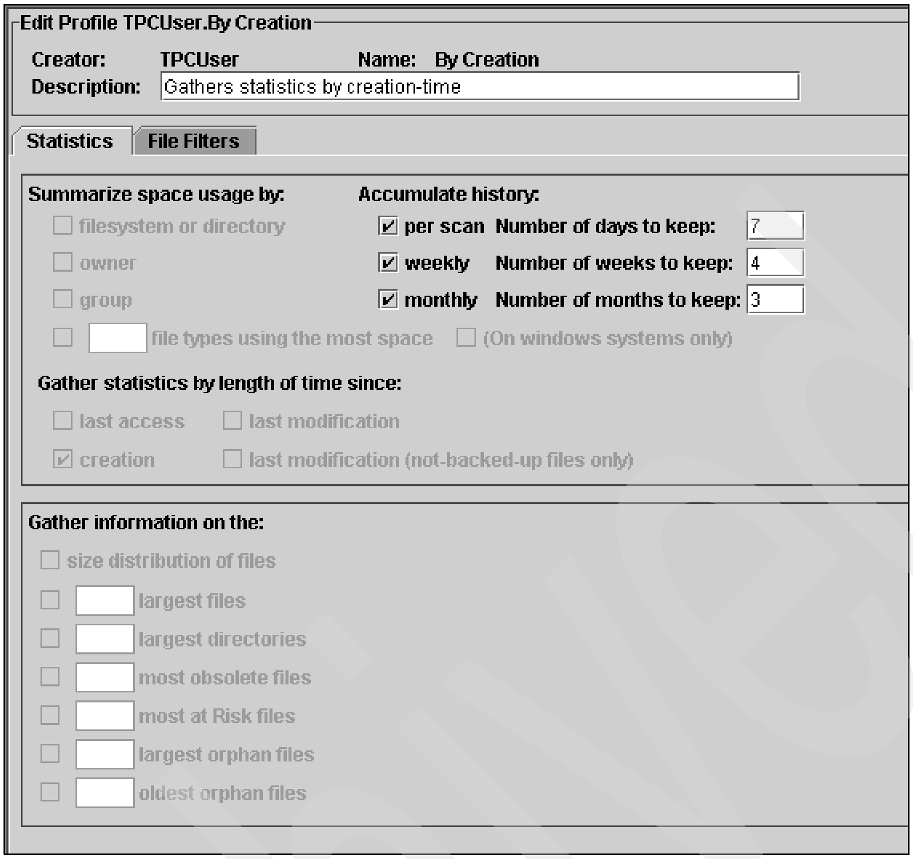

Edit Profile TPCUser.By Creation

Creator: TPCUser Name: By Creation
Description: Gathers statistics by creation-time

Summarize space usage by:

- ☐ filesystem or directory
- ☐ owner
- ☐ group
- ☐ file types using the most space ☐ (On windows systems only)

Accumulate history:

- ☑ per scan Number of days to keep: 7
- ☑ weekly Number of weeks to keep: 4
- ☑ monthly Number of months to keep: 3

Gather statistics by length of time since:

- ☐ last access ☐ last modification
- ☑ creation ☐ last modification (not-backed-up files only)

Gather information on the:

- ☐ size distribution of files
- ☐ largest files
- ☐ largest directories
- ☐ most obsolete files
- ☐ most at Risk files
- ☐ largest orphan files
- ☐ oldest orphan files

Figure 4-29 Profile defined for statistical data history

## 4.6.4 Removed resource retention

*Removed Resource Retention* controls how long an entity resides with its status missing in the TPC repository, and removed resource retention is displayed at the viewer.

For example, if a volume has been deleted, after the next probe, the storage subsystem health indicates a warning, because a volume is missing. This health status remains until you remove the missing entity manually or the Remove Resource Retention time has passed. You can only remove missing entities from the TPC repository, except that you can also remove the storage subsystem from the TPC repository. Select **Disk Manager** → **Storage Subsystems,** click the subsystem that you want to remove, and select **Remove**.

Select **Administrative Services** → **Configuration** → **Removed Resource Retention** to reach the control panel. Figure 4-30 displays the default settings.

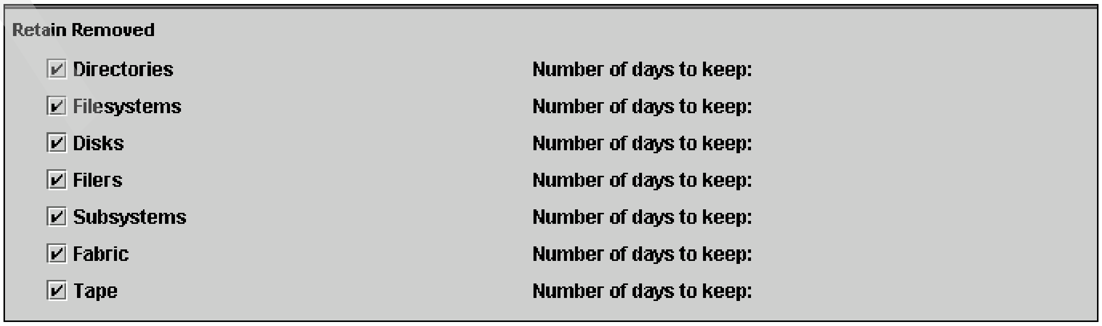

Retain Removed

- ☑ Directories Number of days to keep:
- ☑ Filesystems Number of days to keep:
- ☑ Disks Number of days to keep:
- ☑ Filers Number of days to keep:
- ☑ Subsystems Number of days to keep:
- ☑ Fabric Number of days to keep:
- ☑ Tape Number of days to keep:

Figure 4-30 Removed resource retention

TPC for Databases has its own resource history retention configuration. Select **Administrative Services → Configuration → Removed Resource Retention for Databases** to access the configuration panel. Figure 4-31 shows the default settings.

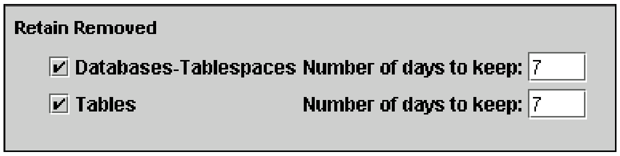

Retain Removed

☑ Databases-Tablespaces Number of days to keep: 7

☑ Tables Number of days to keep: 7

*Figure 4-31 Removed resource retention for Databases*

### 4.6.5 History Aggregator

The *History Aggregator* enables jobs to sum data in an enterprise environment for historical reporting purposes, such as summary statistics and user space usage statistics, and is extremely important. Check the state of the History Aggregator regularly. Improper maintenance in this area can negatively impact historical data trending. Run the History Aggregator daily, which is the default. Configure the History Aggregator at **Administrative Services → Configuration → History Aggregator**. Figure 4-32 shows the configuration pane. The History Aggregator is used within TPC for Data but also removes old Alert Log records. You do not use the History Aggregator to create hourly and daily performance samples.

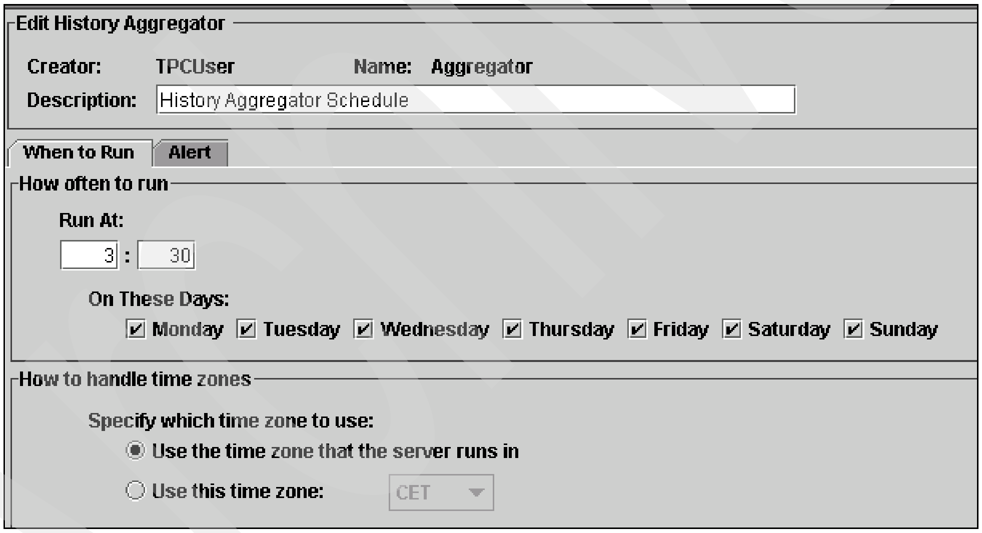

Edit History Aggregator

Creator: TPCUser Name: Aggregator

Description: History Aggregator Schedule

How often to run

Run At: 3 : 30

On These Days:

☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday ☑ Friday ☑ Saturday ☑ Sunday

How to handle time zones

Specify which time zone to use:

● Use the time zone that the server runs in

○ Use this time zone: CET

*Figure 4-32 History Aggregator*

## 4.7 What logs to review

Use the alert log to view and operate upon alerts (performance, state changes, and other alerts). The Alert Log node allows an administrator to view enterprise-wide alerts that are categorized by different criteria, which makes alert management easier. See the alert log at **IBM TotalStorage Productivity Center → Alerting → Alert Log**. Figure 4-33 on page 166 shows the structure of the Alerting node.

The Alert Log viewer is easy to use. If you have many alert logs, it might take some effort to handle them all.

The only way to export alerts (for example, you want to notify your colleagues and assign alerts to them, or for just reporting open alerts) from the Alert Viewer is the printing function with output in a *.csv file. You can only export all of the alerts in a section (node). Therefore, use a spreadsheet application, such as Excel® to modify the list (*.csv file) afterwards. The alert details are not exported (printed).

Refresh the Alerting overview at **File** → **Refresh Alerts**.

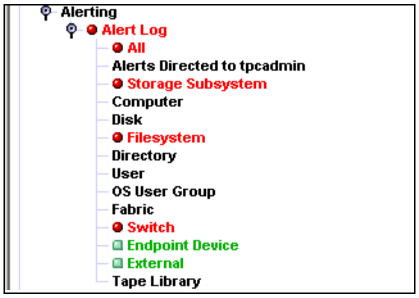

```
Alerting
  ● Alert Log
      ● All
      Alerts Directed to tpcadmin
      ● Storage Subsystem
      Computer
      Disk
      ● Filesystem
      Directory
      User
      OS User Group
      Fabric
      ● Switch
      🟩 Endpoint Device
      🟩 External
      Tape Library
```

*Figure 4-33 Alerting*

Every alert node is marked with either a red bullet, a green square, or no indication. A green square means that there are new log entries but no errors, such as "`new device discovered`" (neutral alert). A red bullet means that there are a minimum of one error alert (negative alert). Blank means, there are no new alerts. Table 4-28 describes the nodes in the Alert Log.

*Table 4-28 Alert Log nodes*

| **Section (node)** | **Description** |
|---|---|
| All | All alerts are listed here, including failed jobs such as discovery and probes, or failed monitor alerts such as performance monitor. You can only review detailed logs about failed jobs at the specific monitor node (for example, **Disk Manager** → **Monitoring** → **Subsystem Performance Monitoring**). |
| Alerts directed to <*username*> | All alerts with triggered actions, Login Notification, to a specific user. This section is only visible to the currently logged in user <*username*>. You cannot see alerts directed to another user. |
| Storage subsystems | All storage subsystem-related events, for example, disk utilization percentage threshold is violated or subsystem allocated capacity change. |
| Computer | All computer-related alerts, for example, disk not found. |
| Disk | All disk-related alerts, for example, a disk is not found. (filter view for OS-related or subsystem alerts) |
| Filesystem | All filesystem-related alerts, for example, filesystem low on free space, constraint violations, and more. |
| Directory | All directory-related alerts, for example, directory too big. |

| Section (node) | Description |
|---|---|
| User | All user-related alerts, for example, a user quota violation. |
| OS user group | All OS user group alerts, for example, a user group quota violation. |
| Fabric | All fabric-related alerts, for example, connection state change (missing/rediscovered). |
| Switch | All switch-related alerts, for example, total port data rate threshold violation or switch status change (offline/online). |
| Endpoint device | All endpoint device-related alerts, for example, endpoint device discovered. |
| External | An external event can be a CIM indication from a CIMOM or an SNMP trap from a SAN switch. The alert contains a timestamp and alert text. The source is the IP from the CIMOM or SNMP trap sender. Those events are only informational events and can be used to verify that TPC receives alerts from external sources. |
| Tape library | All tape library-related alerts |

By selecting a section in the Navigation Tree, the following structure shown in Figure 4-34 appears. Every alert uses one row.

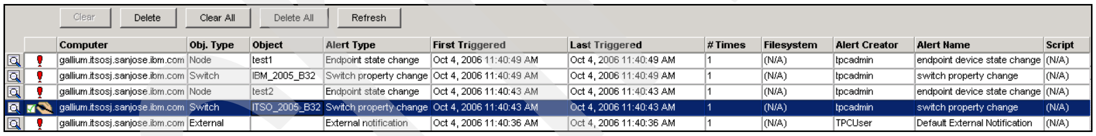

| | | Computer | Obj. Type | Object | Alert Type | First Triggered | Last Triggered | # Times | Filesystem | Alert Creator | Alert Name | Script |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 🔍 | ! | gallium.itsosj.sanjose.ibm.com | Node | test1 | Endpoint state change | Oct 4, 2006 11:40:49 AM | Oct 4, 2006 11:40:49 AM | 1 | (N/A) | tpcadmin | endpoint device state change | (N/A) |
| 🔍 | ! | gallium.itsosj.sanjose.ibm.com | Switch | IBM_2005_B32 | Switch property change | Oct 4, 2006 11:40:49 AM | Oct 4, 2006 11:40:49 AM | 1 | (N/A) | tpcadmin | switch property change | (N/A) |
| 🔍 | ! | gallium.itsosj.sanjose.ibm.com | Node | test2 | Endpoint state change | Oct 4, 2006 11:40:43 AM | Oct 4, 2006 11:40:43 AM | 1 | (N/A) | tpcadmin | endpoint device state change | (N/A) |
| 🔍 | ✓ | gallium.itsosj.sanjose.ibm.com | Switch | ITSO_2005_B32 | Switch property change | Oct 4, 2006 11:40:43 AM | Oct 4, 2006 11:40:43 AM | 1 | (N/A) | tpcadmin | switch property change | (N/A) |
| 🔍 | ! | gallium.itsosj.sanjose.ibm.com | External | | External notification | Oct 4, 2006 11:40:36 AM | Oct 4, 2006 11:40:36 AM | 1 | (N/A) | TPCUser | Default External Notification | (N/A) |

*Figure 4-34 Example Alerting → All*

By clicking the magnifying glass icon 🔍 (drill down) at the first column or double-clicking the row, details about the alert appear as shown in Figure 4-35.

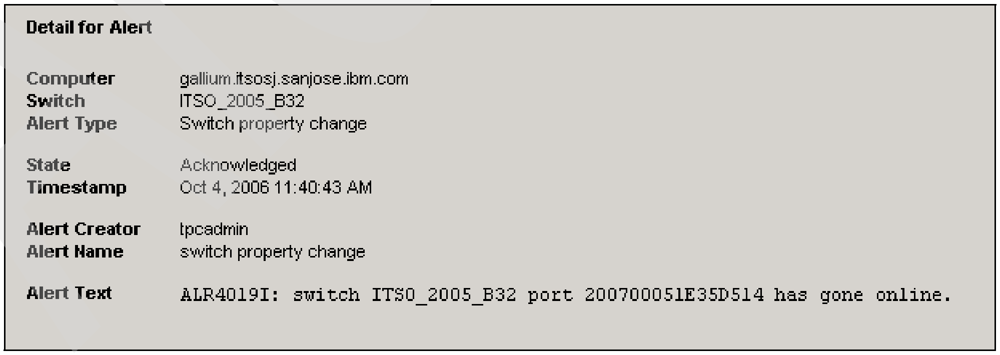

**Detail for Alert**

| | |
|---|---|
| **Computer** | gallium.itsosj.sanjose.ibm.com |
| **Switch** | ITSO_2005_B32 |
| **Alert Type** | Switch property change |
| **State** | Acknowledged |
| **Timestamp** | Oct 4, 2006 11:40:43 AM |
| **Alert Creator** | tpcadmin |
| **Alert Name** | switch property change |
| **Alert Text** | `ALR4019I: switch ITSO_2005_B32 port 200700051E35D514 has gone online.` |

*Figure 4-35 Detail for Alert*

After you investigate the alert (for example, use the element managers to analyze device alerts), you can clear or delete the alert by using “Cleared” or “Deleted” to acknowledge the alert. Just select the alert row and click Clear or Delete . Multiple selections are also possible. If you want keep the alert, use **Clear**. If you want to permanently remove the alert, use **Delete**. After clearing the alert, the second column shows a picture of a hand making a check mark ✓ instead of a red exclamation mark ❗. The Alert State in the detailed view changes from Active to Acknowledged. See Figure 4-34 on page 167. There is not an indication about who cleared (Acknowledged) the alert. Table 4-29 describes the columns available in the alert tables. Every section has its own set of information.

*Table 4-29 Alerting structure*

| Column label | Description | Available for node |
|---|---|---|
| Storage subsystem | Identifier for the subsystem | Storage Subsystem |
| Computer | The computer that reported the alert to TPC. This can be, for example, the TPC server, a CIM Agent, or something else. | All<br>Storage Subsystem<br>Computer<br>Filesystem<br>User<br>OS User Group<br>Fabric<br>Switch<br>Endpoint Device |
| External | Your TPC server host name. | External |
| User | Identifies the user, for example, on a quota violation. Username is presented in the format *username@hostname*. | User |
| Filesystem | Identifies the filesystem, for example, a constraint violation on the filesystem C:\. | Filesystem<br>User |
| Enclosure Type | Enclosure type of reporting alert, for example, computer or storage subsystem. | Disk |
| Enclosure | Enclosure label name of reporting alert, for example, a computer name. | Disk |
| Disk | Identifies the disk of the reporting alert, for example, `Western Digital WT400BB-23DEA0`." | Storage Subsystem<br>Disk |
| Object Type | Identifies the object type, such as disk, fabric, fabric zone, fabric alias, switch, node (endpoint), port, external, filesystem, directory, discovery, and more. | All<br>Computer<br>Fabric<br>Switch<br>Endpoint Device<br>External |
| Object | Identifies the object, such as disk serial numbers, filesystems, directories, fabric label, switch label, zone label, endpoint label, and more, or it also can be a job name. | All<br>Computer<br>Fabric<br>Switch<br>Endpoint Device<br>External |

| Column label | Description | Available for node |
|---|---|---|
| Alert Type | Defined alert condition of the alert, such as Computer Quota Violation, Total Port Data Rate Threshold, External notification, and more. | All<br>Storage Subsystem<br>Computer<br>Disk<br>Filesystem<br>User<br>Fabric<br>Switch<br>Endpoint Device<br>External |
| First Triggered | First triggered date and time of the alert. | All<br>Storage Subsystem<br>Computer<br>Filesystem<br>User<br>Fabric<br>Switch<br>Endpoint Device<br>External |
| Last Triggered | Last triggering date and time. This is used if the event happened more than one time, for example, the filesystem free space threshold was violated 10 times. | All<br>Storage Subsystem<br>Computer<br>Filesystem<br>User<br>Fabric<br>Switch<br>Endpoint Device<br>External |
| # Times | Number of times an alert occurred. This happens if a constraint violation occurred several times. | All<br>Storage Subsystem<br>Computer<br>Filesystem<br>User<br>Fabric<br>Switch<br>Endpoint Device<br>External |
| Alert Creator | What TPC user created the alert condition (owner), for example, TPCUser for default alerts. | All<br>Storage Subsystem<br>Computer<br>Filesystem<br>User<br>Fabric<br>Switch<br>Endpoint Device<br>External |

| Column label | Description | Available for node |
|---|---|---|
| Alert Name | User-defined label of the alert. | All<br>Storage Subsystem<br>Computer<br>Filesystem<br>User<br>Fabric<br>Switch<br>Endpoint Device<br>External |
| Script | The script name of the executed script. | All<br>Storage Subsystem<br>Computer<br>Filesystem<br>User<br>Fabric<br>Switch<br>Endpoint Device<br>External |

If there are dependent alerts in a section, such as for Disk (disk subsystem and OS alerts), you can filter the alerts by using the *View*.

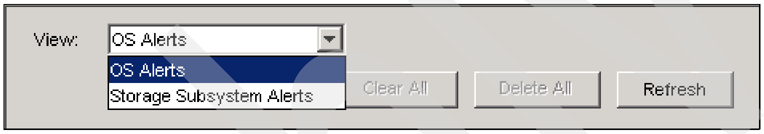

*Figure 4-36 View filter*

If you want to display alerts from a specific asset (only for computer and storage subsystems), open **Data Manager** → **Reporting** → **Asset**. Select an asset, for example, a Storage Subsystem with a right-click, and choose **View Alerts**. Only alerts related to this asset display in the right pane.

You can see monitoring job-related alerts (for example, `probe job failed`) within **IBM TotalStorage Productivity Center** → **Alerting** → **Alert Log** → **All** or at the specific job node.

You can see detailed logs about those jobs (probes and performance monitor jobs) on each job node for troubleshooting. Figure 4-37 shows job logs within the monitor node.

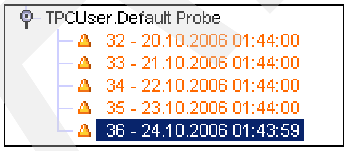
TPCUser.Default Probe
32 - 20.10.2006 01:44:00
33 - 21.10.2006 01:44:00
34 - 22.10.2006 01:44:00
35 - 23.10.2006 01:44:00
36 - 24.10.2006 01:43:59

*Figure 4-37 Job logs within the monitor node*

By selecting the failed probe, the job log displays as shown in Figure 4-38 on page 171.

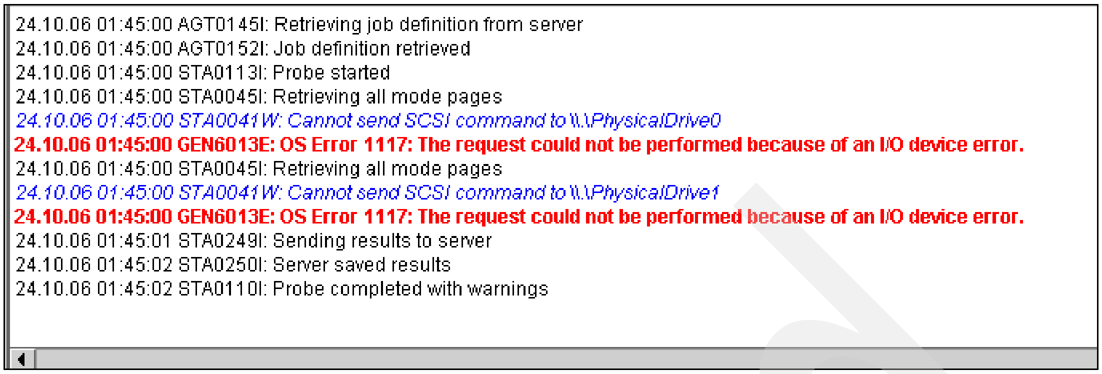

```
24.10.06 01:45:00 AGT0145I: Retrieving job definition from server
24.10.06 01:45:00 AGT0152I: Job definition retrieved
24.10.06 01:45:00 STA0113I: Probe started
24.10.06 01:45:00 STA0045I: Retrieving all mode pages
24.10.06 01:45:00 STA0041W: Cannot send SCSI command to \\.\PhysicalDrive0
24.10.06 01:45:00 GEN6013E: OS Error 1117: The request could not be performed because of an I/O device error.
24.10.06 01:45:00 STA0045I: Retrieving all mode pages
24.10.06 01:45:00 STA0041W: Cannot send SCSI command to \\.\PhysicalDrive1
24.10.06 01:45:00 GEN6013E: OS Error 1117: The request could not be performed because of an I/O device error.
24.10.06 01:45:01 STA0249I: Sending results to server
24.10.06 01:45:02 STA0250I: Server saved results
24.10.06 01:45:02 STA0110I: Probe completed with warnings
```

*Figure 4-38 Detailed job log*

For large log files, use the double arrows at the right side of the window to scroll through the file and use **Top** and **Bottom** (in the upper right corner) to navigate to the beginning or to the end of the log file as shown in Figure 4-39. Alternatively, you can search up and down the log file for strings.

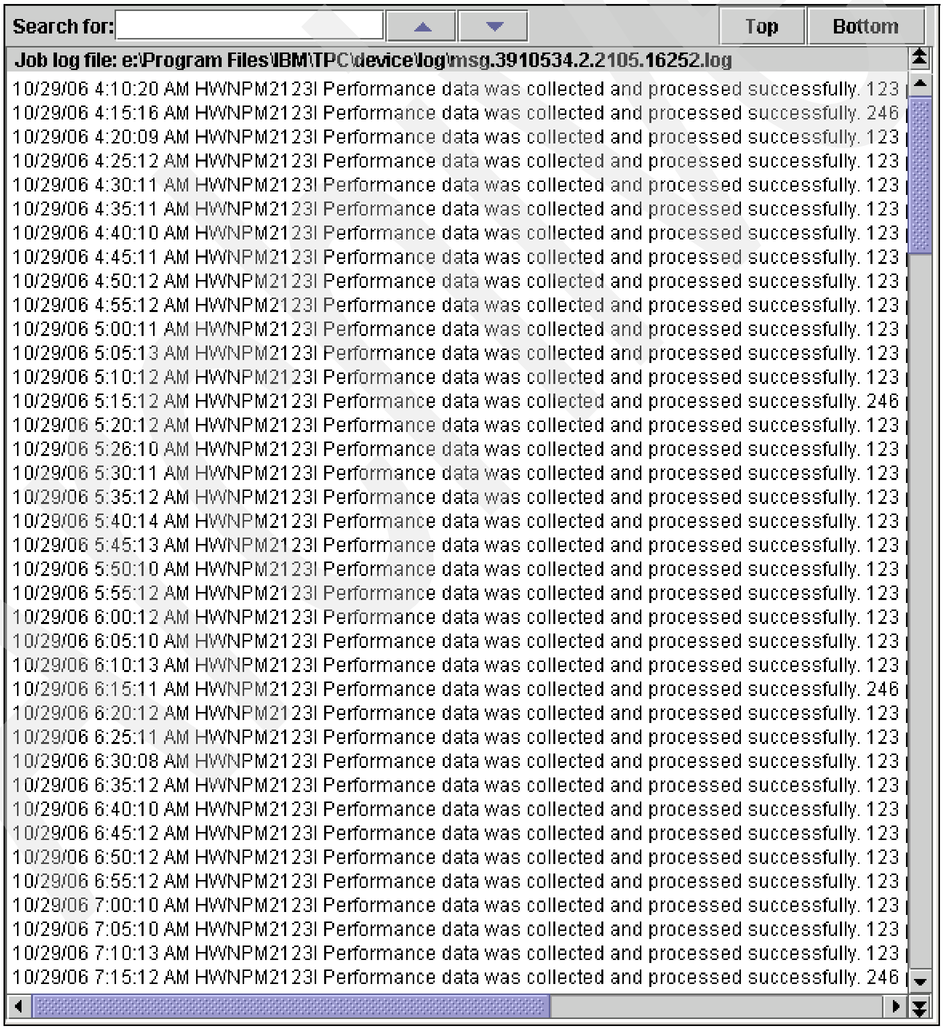

```
Job log file: e:\Program Files\IBM\TPC\device\log\msg.3910534.2.2105.16252.log
10/29/06 4:10:20 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:15:16 AM HWNPM2123I Performance data was collected and processed successfully. 246
10/29/06 4:20:09 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:25:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:30:11 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:35:11 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:40:10 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:45:11 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:50:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 4:55:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:00:11 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:05:13 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:10:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:15:12 AM HWNPM2123I Performance data was collected and processed successfully. 246
10/29/06 5:20:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:26:10 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:30:11 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:35:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:40:14 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:45:13 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:50:10 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 5:55:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:00:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:05:10 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:10:13 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:15:11 AM HWNPM2123I Performance data was collected and processed successfully. 246
10/29/06 6:20:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:25:11 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:30:08 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:35:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:40:10 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:45:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:50:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 6:55:12 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 7:00:10 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 7:05:10 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 7:10:13 AM HWNPM2123I Performance data was collected and processed successfully. 123
10/29/06 7:15:12 AM HWNPM2123I Performance data was collected and processed successfully. 246
```

*Figure 4-39 TPC log navigation*

Another easy way to look at the log files from a job (monitor) is to right-click the node and choose **History** (Figure 4-40).

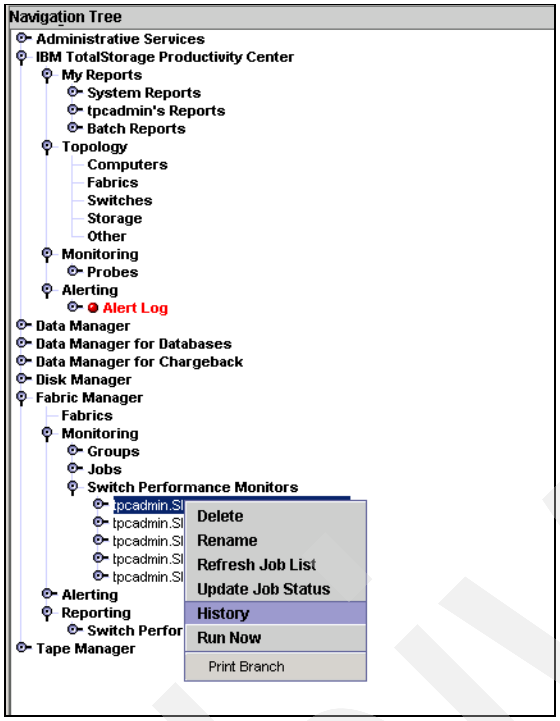

*Figure 4-40 Job history*

Figure 4-41 shows the job history.

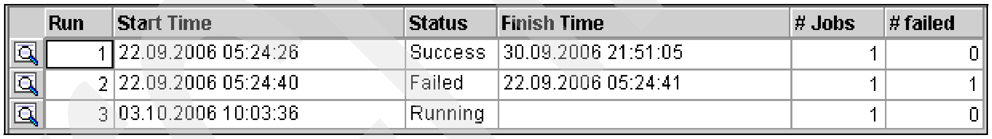

| | Run | Start Time | Status | Finish Time | # Jobs | # failed |
|---|---|---|---|---|---|---|
| | 1 | 22.09.2006 05:24:26 | Success | 30.09.2006 21:51:05 | 1 | 0 |
| | 2 | 22.09.2006 05:24:40 | Failed | 22.09.2006 05:24:41 | 1 | 1 |
| | 3 | 03.10.2006 10:03:36 | Running | | 1 | 0 |

*Figure 4-41 Job history detail*

There are different logs, such as alert logs and monitoring logs (for example, probe and performance). Check the status of those logs regularly. If a monitor job status is green, all is good. If the status is yellow, you can check out the warning messages. If the status is red, the monitor is stopped.

> **Note:** Warnings in job logs do not appear in the Alerting node and do not trigger an alert. A job only creates an alert if the job fails.

## 4.7.1 Messages

You can read the detailed explanations of the messages that are logged in log files in *IBM TotalStorage Productivity Center Messages,* GC32-7194. The messages are grouped in different sections. Use these sections to navigate the Messages manual. Table 4-30 on page 173 provides a brief overview of the TPC messages.

*Table 4-30 TPC message overview*

| Identifier | Number | Type |
|---|---|---|
| Messages start with a nonnumeric code. This code identifies which component of IBM TPC generated the message. When you look in the Messages manual, remember that descriptions are categorized by these designations first. | Messages contain a number following the code designation. Use this number to find the message description. | There are three types, and they are classified:<br>E = Error<br>W = Warning<br>I = Information |

## 4.7.2 TPC Web sites

This Web site is a helpful link to get a description of the error message. Insert the error code into the search text field:

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp

On the IBM Web site, you might find helpful information about APARs, software updates, and known issues (see Figure 4-42 on page 174):

http://www-03.ibm.com/servers/storage/support/software/tpc/

Often checking the compatibility matrix for your devices is useful.

Disk:

http://www-03.ibm.com/servers/storage/support/software/tpcdisk/installing.html

Fabric:

http://www-03.ibm.com/servers/storage/support/software/tpcfabric/installing.html

Data:

http://www-03.ibm.com/servers/storage/support/software/tpcdata/installing.html

And, you can use Google for an information repository:

http://www.google.com

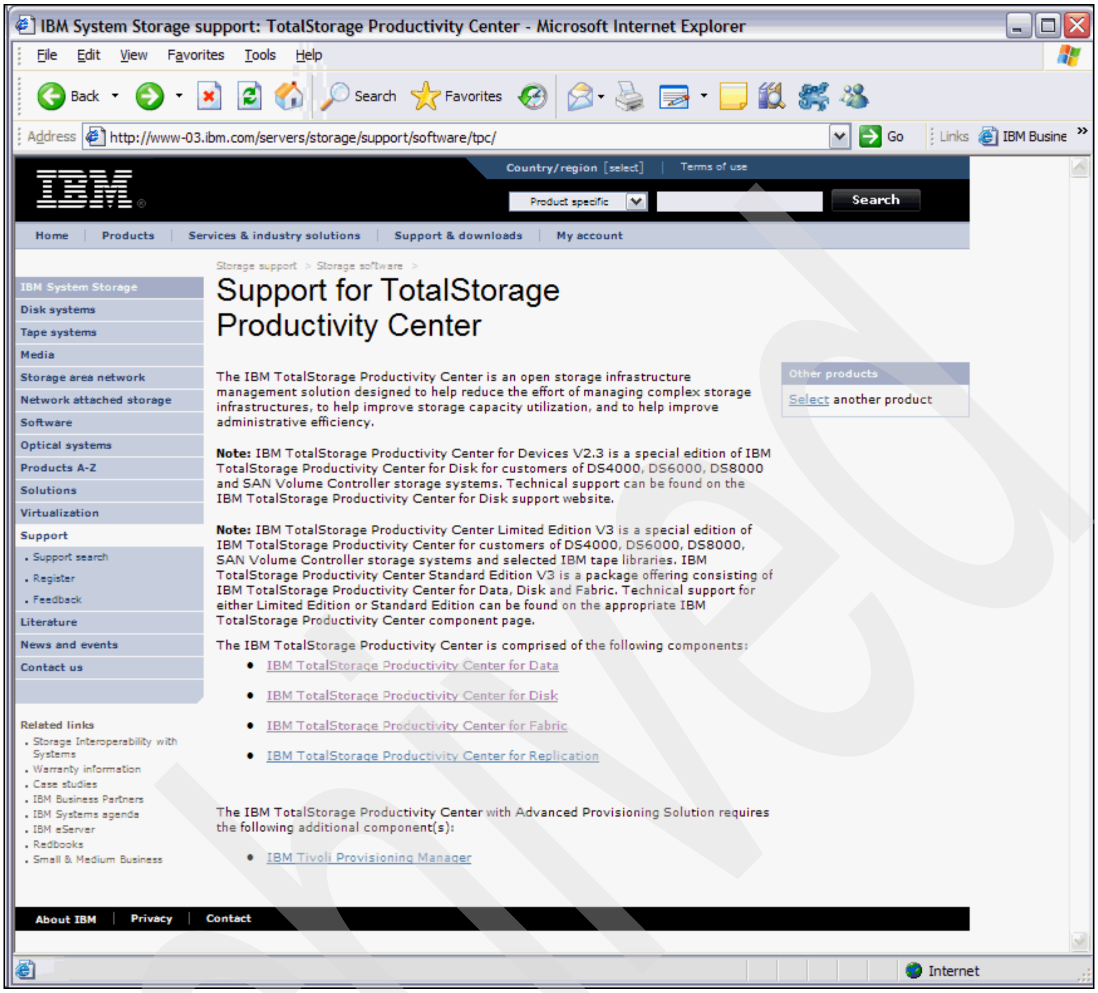

http://www-03.ibm.com/servers/storage/support/software/tpc/

Support for TotalStorage Productivity Center

The IBM TotalStorage Productivity Center is an open storage infrastructure management solution designed to help reduce the effort of managing complex storage infrastructures, to help improve storage capacity utilization, and to help improve administrative efficiency.

Select another product

**Note:** IBM TotalStorage Productivity Center for Devices V2.3 is a special edition of IBM TotalStorage Productivity Center for Disk for customers of DS4000, DS6000, DS8000 and SAN Volume Controller storage systems. Technical support can be found on the IBM TotalStorage Productivity Center for Disk support website.

**Note:** IBM TotalStorage Productivity Center Limited Edition V3 is a special edition of IBM TotalStorage Productivity Center for customers of DS4000, DS6000, DS8000, SAN Volume Controller storage systems and selected IBM tape libraries. IBM TotalStorage Productivity Center Standard Edition V3 is a package offering consisting of IBM TotalStorage Productivity Center for Data, Disk and Fabric. Technical support for either Limited Edition or Standard Edition can be found on the appropriate IBM TotalStorage Productivity Center component page.

The IBM TotalStorage Productivity Center is comprised of the following components:

- IBM TotalStorage Productivity Center for Data
- IBM TotalStorage Productivity Center for Disk
- IBM TotalStorage Productivity Center for Fabric
- IBM TotalStorage Productivity Center for Replication

The IBM TotalStorage Productivity Center with Advanced Provisioning Solution requires the following additional component(s):

- IBM Tivoli Provisioning Manager

*Figure 4-42 TPC Web site*

## 4.8 Alerts and Topology Viewer

The *Topology Viewer* uses the TPC database as the central repository for all of the data that it displays. The Topology Viewer actually reads the data in user-defined intervals from the database and updates, if necessary, the displayed information automatically.

The overall goal of the Topology Viewer is to provide a central location to view a storage environment, quickly monitor and troubleshoot problems, and gain access to additional tasks and functions within the TPC User Interface.

The Topology Viewer represents the current state of an entity or entity group. It presents the actual health state and performance state. By default, the performance presentation layer is disabled. To enable it, right-click the **Topology Viewer** panel → **Global Settings** → **Performance,** and check **Performance** (see Figure 4-43 on page 175).

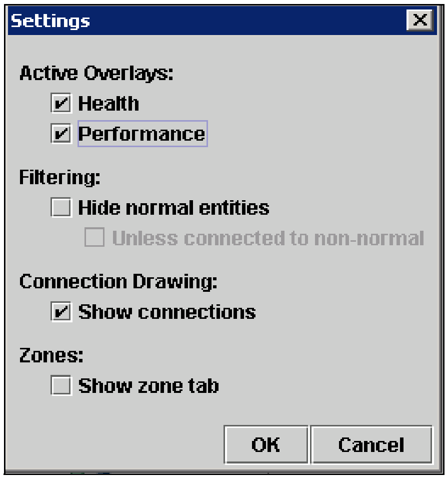

Settings

Active Overlays:

☑ Health

☑ Performance

Filtering:

□ Hide normal entities

□ Unless connected to non-normal

Connection Drawing:

☑ Show connections

Zones:

□ Show zone tab

*Figure 4-43 Enable the performance layer in the Topology Viewer*

To set the refresh interval, right-click the **Topology Viewer** panel → **Refresh Settings** → and select a New refresh rate. The default is 5 minutes, which works well for most requirements. The refresh can take a while and consume CPU cycles if you have multiple Topology Viewer views open at the same time. If you do not need a current reflection of the state of your environment, increase the refresh rate to 60 minutes in the Topology Viewer and refresh the Topology Viewer manually as needed. To refresh the Topology Viewer, right-click the **Topology Viewer** pane and select **Refresh View** or **Refresh All Views** as shown in Figure 4-44.

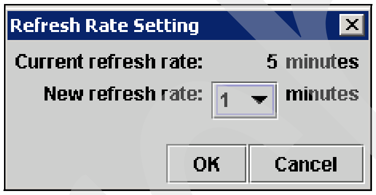

Refresh Rate Setting

Current refresh rate: 5 minutes

New refresh rate: 1 minutes

*Figure 4-44 Refresh Rate Setting*

The Topology Viewer provides colored and iconic overlays of information, such as:

- Health status of entities or group of entities
- Performance overlays
- Policy compliance and violations

An entity shows the health (upper square) and the performance overlay (lower bar chart under the health indicator) in the graphical view of the Topology Viewer as shown in Figure 4-45. The icons' status (Health Status and Performance Status) are also shown in the tabular view in the columns `Health` and `Performance`.

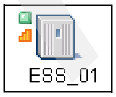

*Figure 4-45 Performance and health status*

You can hover the mouse over the Performance icon to get metrics from the last collected performance sample as shown in Figure 4-46 on page 176.

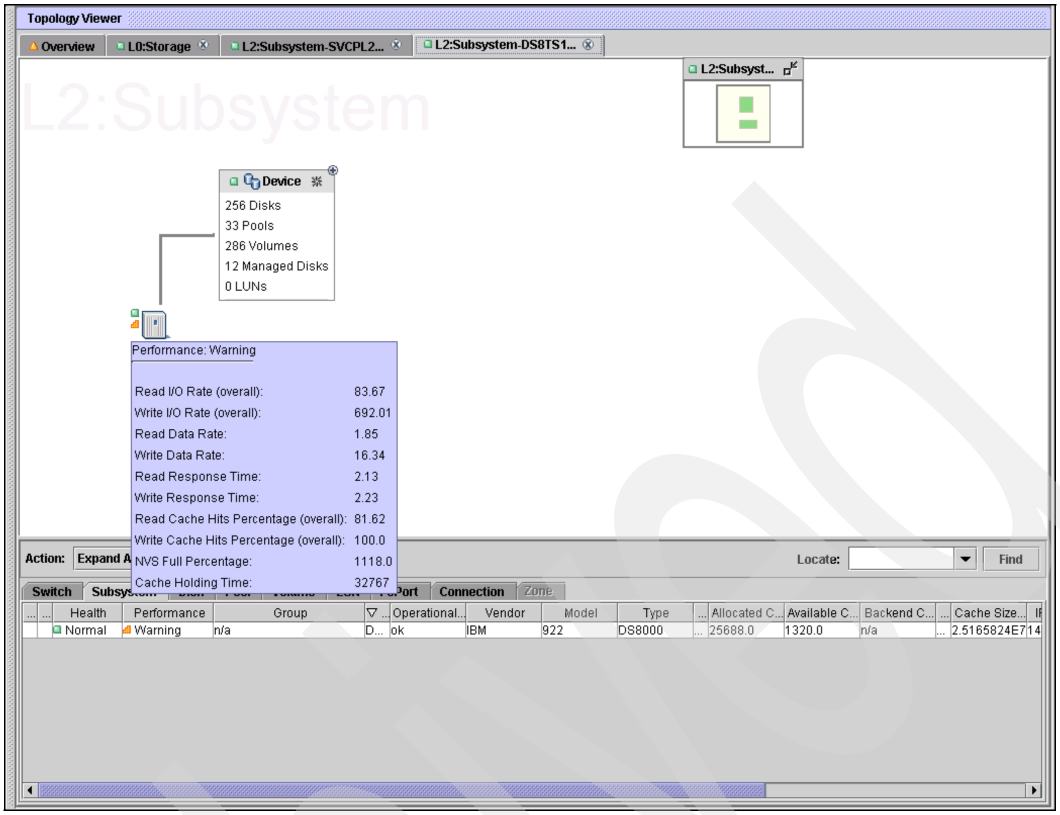

Performance: Warning

| | |
|---|---|
| Read I/O Rate (overall): | 83.67 |
| Write I/O Rate (overall): | 692.01 |
| Read Data Rate: | 1.85 |
| Write Data Rate: | 16.34 |
| Read Response Time: | 2.13 |
| Write Response Time: | 2.23 |
| Read Cache Hits Percentage (overall): | 81.62 |
| Write Cache Hits Percentage (overall): | 100.0 |
| NVS Full Percentage: | 1118.0 |
| Cache Holding Time: | 32767 |

| ... | ... | Health | Performance | Group | ▽ ... | Operational... | Vendor | Model | Type | ... | Allocated C... | Available C... | Backend C... | ... | Cache Size... | IF |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Normal | Warning | n/a | D... | ok | IBM | 922 | DS8000 | ... | 25688.0 | 1320.0 | n/a | ... | 2.5165824E7 | 14 |

Figure 4-46   Topology Viewer L2 showing performance hovering

The performance status displays predefined internal threshold violations. The Topology Viewer scans the PMM_Exception table for existing constraint violations. The performance status is aggregated up to the highest level entity, such as the switch or a subsystem. Generate the report **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **Constraint Violation** to get more information about your performance threshold violations.

Topology Viewer shows the performance status of an entity as either:

- *Normal*: Entity is operating at expected performance.
- *Warning*: There is at least one part of the entity that is not operating at expected performance.
- *Critical Entity*: The entity is operating below expected performance.
- *Unavailable*: This status is used if the entity is not available in certain situations.

The health status is also shown on connections between entities. Connections can show a normal state (green) or a critical state (red). If several connections are aggregated to show as one (thicker) line, this line might also show a warning state (yellow) according to the normal aggregation rules. Special connection types, such as trunks, are not explicitly displayed.

The Health Status icon in the Topology Viewer also represents the consolidated status of the device. The consolidated status of the device is not merely the status of the top-level device but also takes into consideration the status of the subordinate elements, such as pools, volumes, Fibre Channel (FC) ports, and more. You can use the drill-down feature of TPC to navigate to the next level of detail (that is, down to L1:Storage and L2:Storage). In these views, the entity groups have health indicators as well, so that you can easily navigate to the element in trouble, for example, a disk with an error.

If a group contains missing elements, then the overall group status is `warning` to indicate the missing element. You can configure the Removed Resource Retention to automatically delete those missing entities from the TPC repository after a specific time or delete them manually by right-clicking the entity and selecting **Remove from Database**.

The group status (aggregate) for health and performance state has these rules:

- RED (Critical): If at least one entity is critical and at least one entity is missing.
- YELLOW (Warning): If at least one entry is in a missing, warning, or critical state.
- GREEN (Normal): If all entities are in a normal or undefined state.

If all entities are in the same state, the aggregated state reflects this state.

The table in Figure 4-47 gives you an overview of the actual states and their graphical representation within the Topology Viewer.

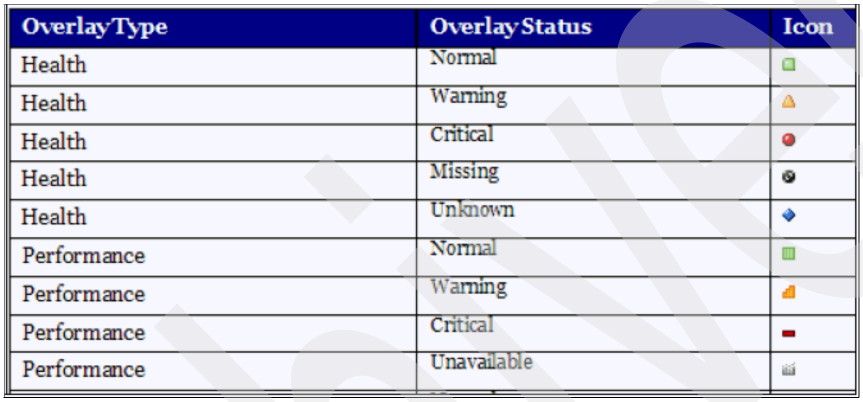

| OverlayType | OverlayStatus | Icon |
|---|---|---|
| Health | Normal | 🟩 |
| Health | Warning | ⚠ |
| Health | Critical | 🔴 |
| Health | Missing | ⚫ |
| Health | Unknown | 🔷 |
| Performance | Normal | ▦ |
| Performance | Warning | 📶 |
| Performance | Critical | ▬ |
| Performance | Unavailable | 📶 |

*Figure 4-47 Health state*

The health state of the entity, such as computers, storage subsystems, SAN switches, and so forth can have the following status shown in Table 4-31.

*Table 4-31 Health status*

| **Health state** | **Description** |
|---|---|
| Normal | Entity is operating normally. |
| Warning | There is at least some part of the entity that is not operating or has serious problems. |
| Critical | Entity is either not operating or has serious operational problems. |
| Missing | Entity was discovered and recognized by TPC but is no longer discovered as of the latest refresh cycle. |
| Unknown | Entity was discovered but not recognized by TPC. |

Note that the user-defined alerts, seen in the Alert Log node, are not propagated to the health status in the Topology Viewer within the current TPC version 3.1.3.

The tabular view of a device provides its health state and the operational state. The information in the Health column and in the Operational Status column is not the same. The Operational Status, which is only seen in the tabular view, shows more detailed information provided by the vendor device. In the example shown in Figure 4-48, the health of both SAN switches is marked Warning. The Operational Status gives more detailed information. In this example, the Operational Status was reported by the Brocade CIM Agent. As you can read in the Brocade document, *Brocade SMI Agent Developer's Guide,* it translates the status `SWITCH_STATUS_WARNING` to `Stressed` and `SWITCH_STATUS_FAULTY` to `Error`.

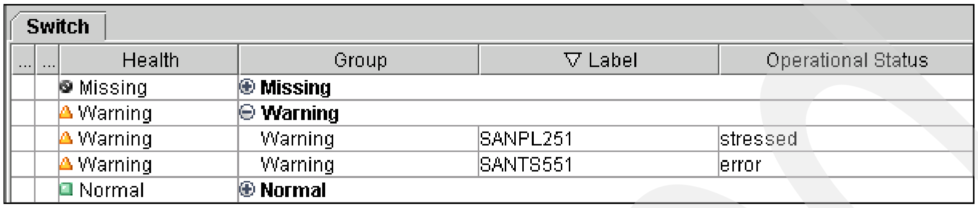

Switch

| ... | ... | Health | Group | ▽ Label | Operational Status |
|---|---|---|---|---|---|
| | | ● Missing | ⊕ **Missing** | | |
| | | ⚠ Warning | ⊖ **Warning** | | |
| | | ⚠ Warning | Warning | SANPL251 | stressed |
| | | ⚠ Warning | Warning | SANTS551 | error |
| | | 🟩 Normal | ⊕ **Normal** | | |

*Figure 4-48 Tabular view of two SAN switches with Operational Status stressed and error*

The Brocade SAN switch shows the following status by using the command **`switchStatusShow`**. Figure 4-49 shows the detailed information for the switch with operational status equal to `stressed`, and Figure 4-50 shows the detailed information for the switch with operational status equal to `error`. There are marginal ports that violate the Switch Status Policy (**`switchStatusPolicyShow`**).

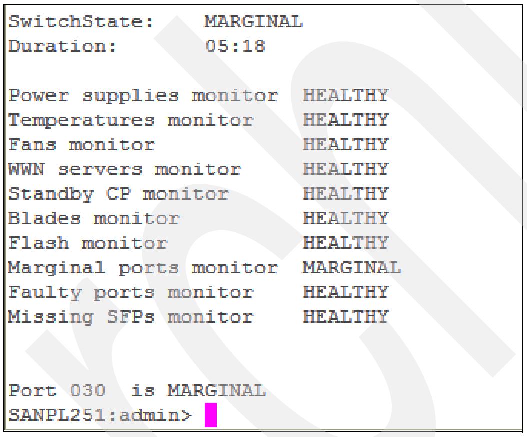

```
SwitchState:    MARGINAL
Duration:       05:18

Power supplies monitor  HEALTHY
Temperatures monitor    HEALTHY
Fans monitor            HEALTHY
WWN servers monitor     HEALTHY
Standby CP monitor      HEALTHY
Blades monitor          HEALTHY
Flash monitor           HEALTHY
Marginal ports monitor  MARGINAL
Faulty ports monitor    HEALTHY
Missing SFPs monitor    HEALTHY


Port 030  is MARGINAL
SANPL251:admin>
```

*Figure 4-49 Switch status marginal*

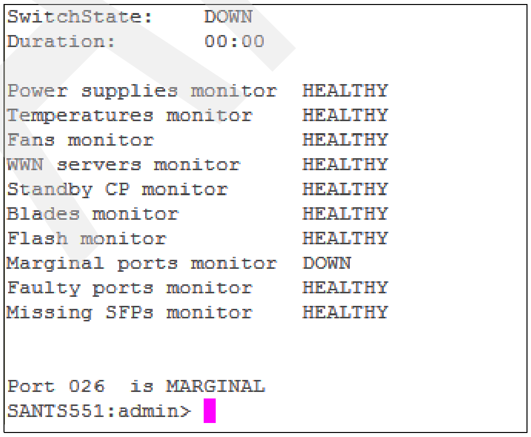

```
SwitchState:    DOWN
Duration:       00:00

Power supplies monitor  HEALTHY
Temperatures monitor    HEALTHY
Fans monitor            HEALTHY
WWN servers monitor     HEALTHY
Standby CP monitor      HEALTHY
Blades monitor          HEALTHY
Flash monitor           HEALTHY
Marginal ports monitor  DOWN
Faulty ports monitor    HEALTHY
Missing SFPs monitor    HEALTHY


Port 026  is MARGINAL
SANTS551:admin>
```

*Figure 4-50 Switch status down*

In this example, the solution is to fix the marginal ports on the SAN switch. Possibly, the attached device is not properly running (a new installation) or there has been a change to the Switch Status Policy.

We adjust the Switch Status Policy with the command **`SwitchStatusPolicySet`**, because the switch has many ports and there are many new server installations at the moment. After adjusting the policy, the Switch Status changes to `Healthy` (Figure 4-51).

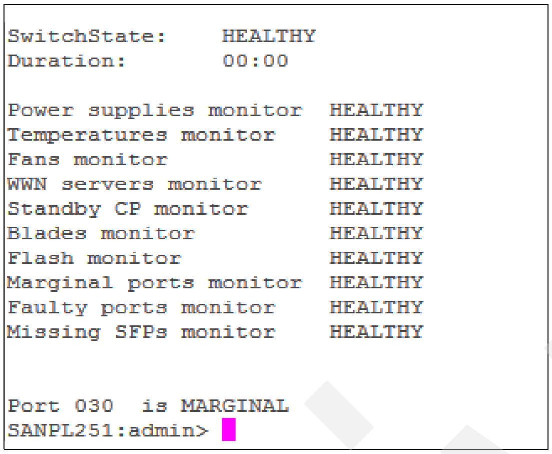

```
SwitchState:     HEALTHY
Duration:        00:00

Power supplies monitor  HEALTHY
Temperatures monitor    HEALTHY
Fans monitor            HEALTHY
WWN servers monitor     HEALTHY
Standby CP monitor      HEALTHY
Blades monitor          HEALTHY
Flash monitor           HEALTHY
Marginal ports monitor  HEALTHY
Faulty ports monitor    HEALTHY
Missing SFPs monitor    HEALTHY


Port 030  is MARGINAL
SANPL251:admin> 
```

*Figure 4-51 Switch status healthy*

Finally, after a CIMOM discovery, TPC shows that the switch is in an `OK` operational status and in normal health (Figure 4-52).

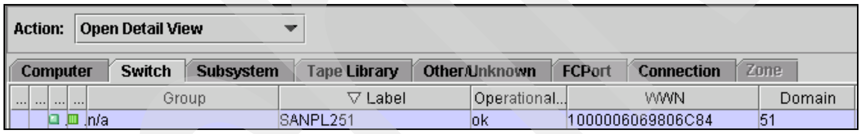

| ... | ... | ... | ... | Group | ▽ Label | Operational... | WWN | Domain |
|---|---|---|---|---|---|---|---|---|
| | | ■ | ■ | n/a | SANPL251 | ok | 1000006069806C84 | 51 |

*Figure 4-52 Tabular view of OK operational status*

## 4.9 Alerting mechanism (event)

TPC uses events to notify subscribers about alerts. The type of alert mechanism that is used differs from environment to environment. The notification mechanisms in Table 4-32 on page 180 are possible.

*Table 4-32 Events*

| Alert mechanism | Enables you to |
|---|---|
| SNMP trap | Generate an SNMP trap message to any network management station (NMS), console, or terminal to indicate the occurrence of an alert. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files in order to receive SNMP traps from IBM TPC. |
| TEC event | Send an alert to the Tivoli Enterprise Console® (TEC). The TEC administrator can write correlation and automation rules to analyze IBM TPC events according to the event definitions specified in the tivoliSRM.baroc and fabric.baroc files. It also performs responses, such as sending further notification, creating or updating trouble tickets, running programs, and so forth. |
| Login notification | Indicate that alerts should appear to a specified user upon accessing the **Alerting** → **Alert Log**. Specify the user to receive the alerts in the Login ID field. Use the **Preferences** → **Edit General** option on the IBM TPC menu bar to control which alerts a user views upon logging in to the product. |
| Windows event log or<br>UNIX Syslog | Write out alert messages to the OS log. If you already have an administrator monitoring OS logs, this is an easy way to have all of your priority messages centralized for quick notification and viewing.<br>Event Type (Windows Event Log only) - Indicate the type of event that will be recorded to the Windows Event Log.<br>Facility (UNIX Syslog only) - Select where the UNIX Syslog will be sent. You can select `User` or `Local`.<br>Level (UNIX Syslog only) - Select the level of UNIX Syslog event. You can select `Informational`, `Notice`, `Warning`, `Error`, `Critical`, or `Alert`. |
| Run script | Run a script in response to an alert. This "triggered action facility" enables you to run scripts that execute any third-party tools for actions, such as archiving, backup and recovery, or provisioning. |
| Archive/Backup<br>(Constraints only) | Define an IBM Tivoli Storage Manager backup or archive job that will run in response to an alert associated with a constraint. |
| E-mail | Receive an e-mail when a specific type of alert is triggered. IBM TPC can send this type of notification to you, the storage administrator, or the user. |

To use these alert mechanisms to define the event disposition, click **Administrative Services** → **Configuration** → **Alert Disposition**. See Figure 4-53.

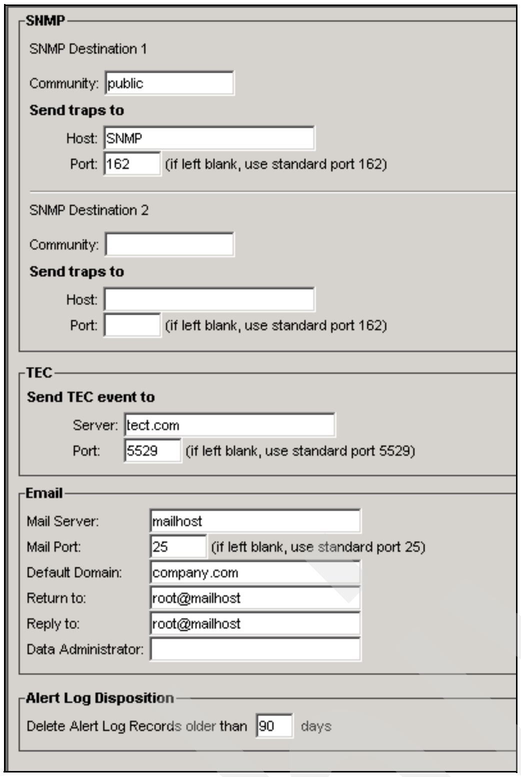

**SNMP**

SNMP Destination 1

Community: public

**Send traps to**

Host: SNMP

Port: 162 (if left blank, use standard port 162)

SNMP Destination 2

Community:

**Send traps to**

Host:

Port: (if left blank, use standard port 162)

**TEC**

**Send TEC event to**

Server: tect.com

Port: 5529 (if left blank, use standard port 5529)

**Email**

Mail Server: mailhost

Mail Port: 25 (if left blank, use standard port 25)

Default Domain: company.com

Return to: root@mailhost

Reply to: root@mailhost

Data Administrator:

**Alert Log Disposition**

Delete Alert Log Records older than 90 days

*Figure 4-53 Event notification configuration*

You also set up the Alert Log Disposition on this window. This option controls how long to keep records in the alert log. The default is `90` days. For most environments, 90 days is sufficient.

> **Note:** Most Web-based e-mail servers, for example, Yahoo! and Gmail, require authentication, commonly Secure Sockets Layer (SSL) and a username and password combination. This authentication methodology is currently unsupported. For best results, use an available internal Simple Mail Transfer Protocol (SMTP) mail server.

## Alert examples

One example is using e-mail notification for all failed monitoring jobs. If a job fails, the alert sends an e-mail to the storage management team mailbox. The responsible person can fix the problem and clean it from the log view in the TPC GUI. If several users work with TPC, you can assign the alert to a specific person by selecting the **Alert** tab → **Login Notification**.

Set up important storage hardware failures, such as San Volume Controller (SVC) Node Offline, to create a TEC event and perhaps create a ticket. Therefore, the responsible person can be notified immediately. Additionally, the ticket can be seen by a large community and is great for reporting and auditing tasks.

For performance threshold violations, use e-mail. But set the thresholds high; otherwise, you might get overwhelmed with e-mails. It is better to use high thresholds for notification and use reporting for performance analysis.

With TPC Version 3.1.2 and Version 3.1.3, you cannot use variable information, such as device name in the alert e-mail subject. You set a fixed subject text per alert.

For quota and constraint violations, you might send an automatic e-mail to the owner or to the system administrator. Refer to “Setting up quota and constraint e-mail” on page 134 for an example.

Figure 4-54, Figure 4-55, and Figure 4-56 show examples of alert notifications.

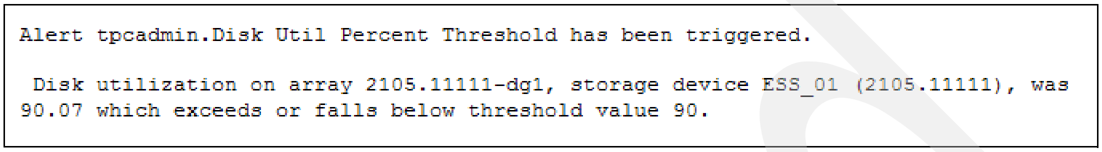

```
Alert tpcadmin.Disk Util Percent Threshold has been triggered.

 Disk utilization on array 2105.11111-dg1, storage device ESS_01 (2105.11111), was
90.07 which exceeds or falls below threshold value 90.
```

*Figure 4-54   Example e-mail of a threshold alert*

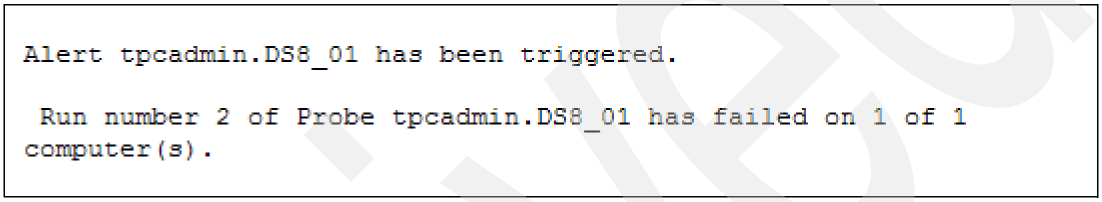

```
Alert tpcadmin.DS8_01 has been triggered.

 Run number 2 of Probe tpcadmin.DS8_01 has failed on 1 of 1
computer(s).
```

*Figure 4-55   Example e-mail of a failed monitor job*

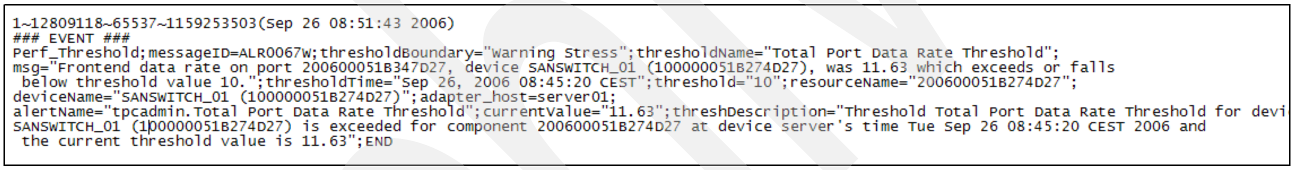

```
1~12809118~65537~1159253503(Sep 26 08:51:43 2006)
### EVENT ###
Perf_Threshold;messageID=ALR0067W;thresholdBoundary="Warning Stress";thresholdName="Total Port Data Rate Threshold";
msg="Frontend data rate on port 200600051B347D27, device SANSWITCH_01 (100000051B274D27), was 11.63 which exceeds or falls
 below threshold value 10.";thresholdTime="Sep 26, 2006 08:45:20 CEST";threshold="10";resourceName="200600051B274D27";
deviceName="SANSWITCH_01 (100000051B274D27)";adapter_host=server01;
alertName="tpcadmin.Total Port Data Rate Threshold";currentValue="11.63";threshDescription="Threshold Total Port Data Rate Threshold for devi
SANSWITCH_01 (100000051B274D27) is exceeded for component 200600051B274D27 at device server's time Tue Sep 26 08:45:20 CEST 2006 and
 the current threshold value is 11.63";END
```

*Figure 4-56   TEC event example*

## 4.10 Reporting

TPC provides reporting of its collected data in the TPC repository. Certain reports are predefined and you can modify them. You can define your own reports based on the standard definitions. Reports display in tables, graphs, and charts. Use these reports for:

- Asset reporting
- Availability reporting
- Capacity reporting
- Usage reporting
- Usage violation reporting
- Backup reporting
- Performance reporting
- Performance alert reporting

You can display the generated reports or you can export them to a printer or file, such as CSV, HTML, PDF, or text-formatted. In addition, you can schedule reports to run at certain times.

For a detailed description of reporting, see *Building and Scaling SAP Business Information Warehouse on DB2 UDB ESE*, SG24-7094. We provide a brief overview and several examples.

## 4.10.1 System reports

To see predefined system reports, select **IBM TotalStorage Productivity Center** → **My Reports** → **System Reports**. The report categories are Data, Disk, and Fabric. See Figure 4-57.

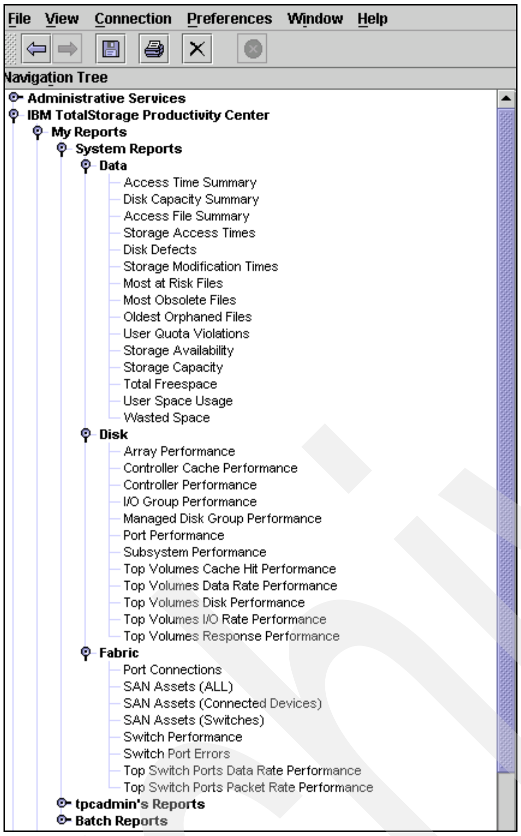

- Administrative Services
- IBM TotalStorage Productivity Center
  - My Reports
    - System Reports
      - Data
        - Access Time Summary
        - Disk Capacity Summary
        - Access File Summary
        - Storage Access Times
        - Disk Defects
        - Storage Modification Times
        - Most at Risk Files
        - Most Obsolete Files
        - Oldest Orphaned Files
        - User Quota Violations
        - Storage Availability
        - Storage Capacity
        - Total Freespace
        - User Space Usage
        - Wasted Space
      - Disk
        - Array Performance
        - Controller Cache Performance
        - Controller Performance
        - I/O Group Performance
        - Managed Disk Group Performance
        - Port Performance
        - Subsystem Performance
        - Top Volumes Cache Hit Performance
        - Top Volumes Data Rate Performance
        - Top Volumes Disk Performance
        - Top Volumes I/O Rate Performance
        - Top Volumes Response Performance
      - Fabric
        - Port Connections
        - SAN Assets (ALL)
        - SAN Assets (Connected Devices)
        - SAN Assets (Switches)
        - Switch Performance
        - Switch Port Errors
        - Top Switch Ports Data Rate Performance
        - Top Switch Ports Packet Rate Performance
    - tpcadmin's Reports
    - Batch Reports

*Figure 4-57 System Reports list*

Note that performance reports rely on the newest performance samples (most recently collected). For example, the `Top Volumes Data Rate Performance` report represents only the newest performance samples and not historical samples.

## 4.10.2 Reports by user name

Select reports by user name or <*username*> reports by clicking **IBM TotalStorage Productivity Center** → **My Reports** → **<username> Reports**. This node includes the saved reports.

For example, you can create a simple SAN asset list with switch label, IP, model, fabric label, and serial number. Select the predefined `SAN Asset (Switches)` report by clicking **IBM TotalStorage Productivity Center** → **My Reports** → **System Reports** → **Fabric** → **SAN Assets (Switches)** and adjust the **Selection** tab in the right pane. See Figure 4-58 on page 184.

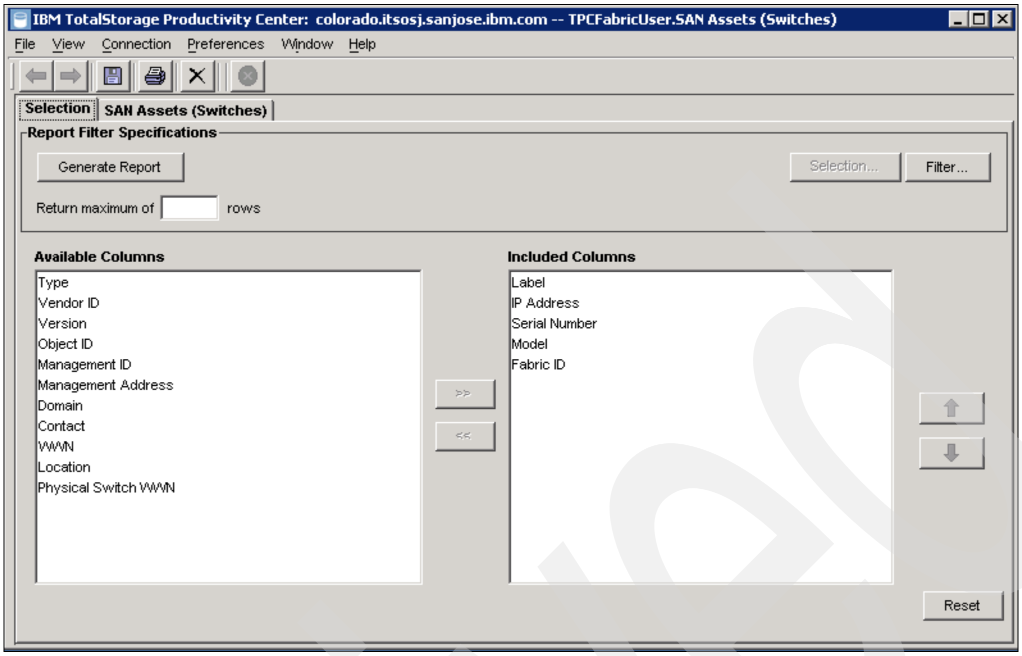
IBM TotalStorage Productivity Center: colorado.itsosj.sanjose.ibm.com -- TPCFabricUser.SAN Assets (Switches)
File
View
Connection
Preferences
Window
Help
Selection
SAN Assets (Switches)
Report Filter Specifications
Generate Report
Selection...
Filter...
Return maximum of
rows
>>
<<
Reset

**Available Columns**

Type
Vendor ID
Version
Object ID
Management ID
Management Address
Domain
Contact
WWN
Location
Physical Switch WWN

**Included Columns**

Label
IP Address
Serial Number
Model
Fabric ID

*Figure 4-58 Column selection*

Select the column identifier that you want and click the arrow at the right side to change the order. Click **Generate Report** and review the result. See Figure 4-59.

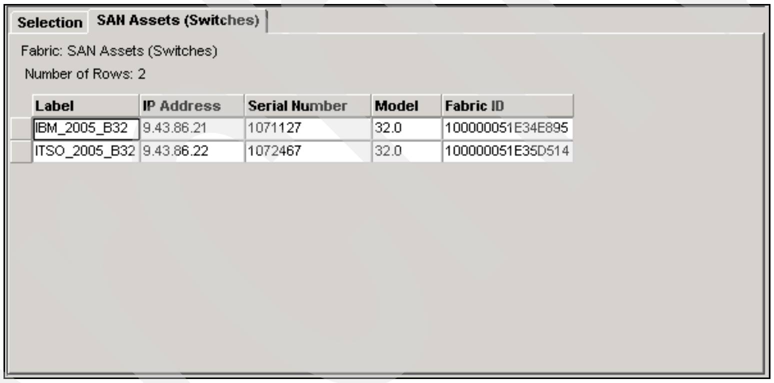

Fabric: SAN Assets (Switches)

Number of Rows: 2

| Label | IP Address | Serial Number | Model | Fabric ID |
|---|---|---|---|---|
| IBM_2005_B32 | 9.43.86.21 | 1071127 | 32.0 | 100000051E34E895 |
| ITSO_2005_B32 | 9.43.86.22 | 1072467 | 32.0 | 100000051E35D514 |

*Figure 4-59 Report*

Click **File** → **Save As** or use the Save icon to save the report with a new report name. You can now run the saved report at <*username*>'s Report.

### 4.10.3 Batch reports

The batch reporting feature enables you to run any report on a regularly scheduled basis. This enables you to conveniently run reports and gather data on a set schedule. The batch reporting feature provides a convenient and powerful way for you to save report definitions and schedule when to run reports. You can use automatic reports for report distribution, publishing on a Web server, or as input for another application, such as a Data Warehouse.

To create a batch report:

1. Select **IBM TotalStorage Productivity Center** → **My Reports** → **Batch Reports**.
2. Right-click **Create Batch Report**.
3. Select one of the existing reports.
4. Choose the columns to include.
5. Set options, such as where to run and store the batch report, the output file type, and whether to run a script after completing the report generation. The script can copy the report to another location or send it through e-mail to certain recipients.
6. Set the run time and frequency.
7. Set the triggering action (alert notification) if the report fails.
8. After a batch report runs, you can check the status and the log at the batch reports node.

The system saves the batch reports at:

For Windows:
<*installation path*>\IBM\TPC\ca\subagents\TPC\Data\log\<*computername*>\reports

For UNIX:
<*installation path*>/IBM/TPC/ca/subagents/TPC/Data/log/<*computername*>/reports

> **Note:** With the TPC Version 3.1.2, you cannot run predefined reports from My Reports (**System Reports** → **<username>'s report**) as a batch report.

### 4.10.4 Data Manager reports

Data Manager (**Data Manager** → **Reporting**) within TPC provides several reports.

Figure 4-60 on page 186 shows the report category overview.

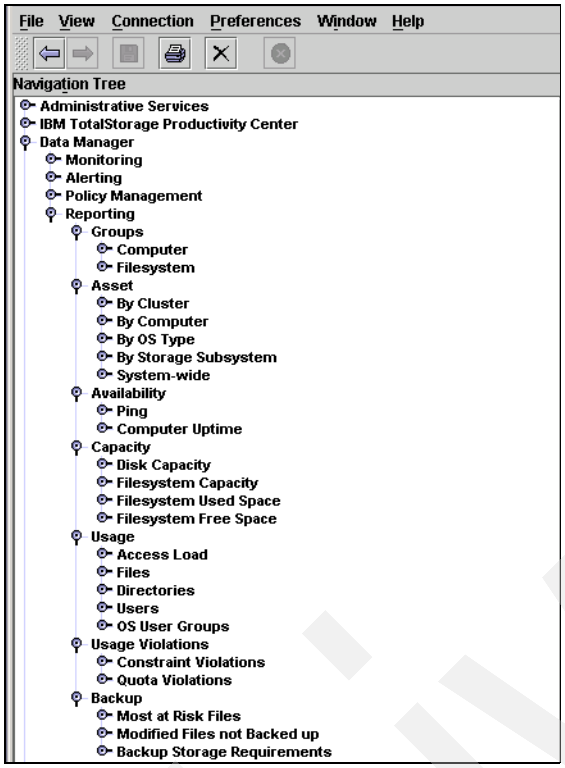

*Figure 4-60 Reporting category within Data Manager reporting*

Here are descriptions with several report examples for the reporting nodes.

### *Groups*

Within groups, you can define computer or file system groups.

### *Asset*

In the asset section, there are computers, clusters, storage subsystems, and more. You can drill down several levels. By selecting a node, you get the details in the right pane. Figure 4-61 on page 187 shows an example of a computer asset. We selected **File System E:\**. The properties display in the right pane and give you information about file system type, last scan, used space, free space, and more.

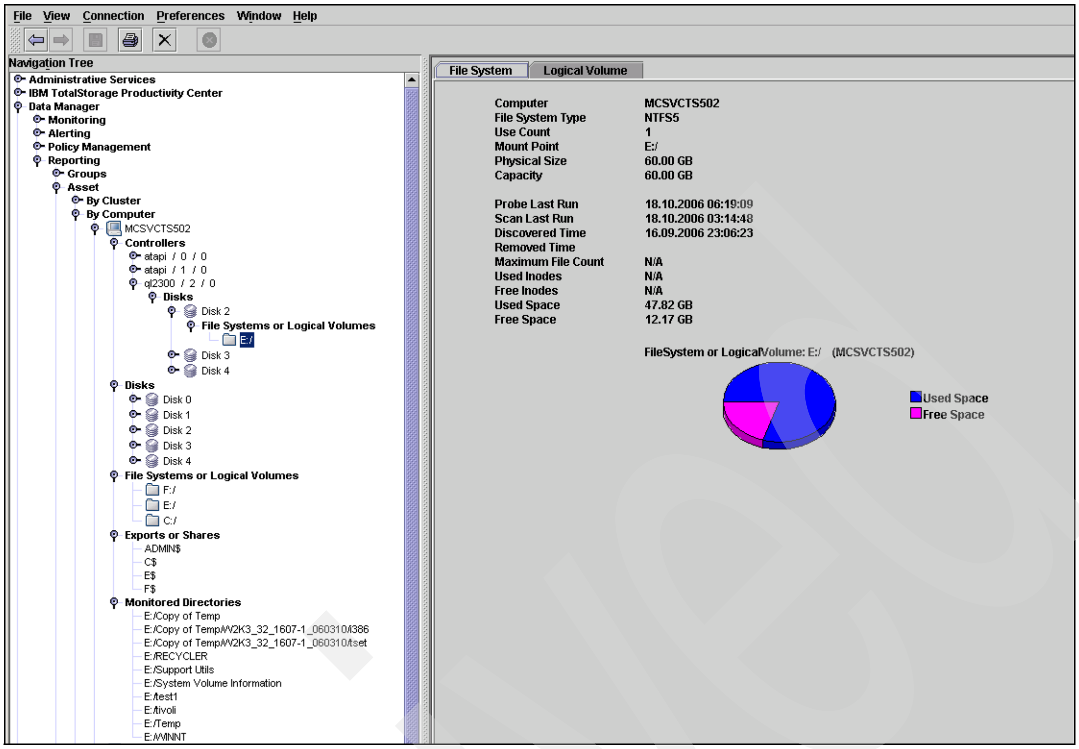

| | |
|---|---|
| Computer | MCSVCTS502 |
| File System Type | NTFS5 |
| Use Count | 1 |
| Mount Point | E:/ |
| Physical Size | 60.00 GB |
| Capacity | 60.00 GB |
| Probe Last Run | 18.10.2006 06:19:09 |
| Scan Last Run | 18.10.2006 03:14:48 |
| Discovered Time | 16.09.2006 23:06:23 |
| Removed Time | |
| Maximum File Count | N/A |
| Used Inodes | N/A |
| Free Inodes | N/A |
| Used Space | 47.82 GB |
| Free Space | 12.17 GB |

*Figure 4-61 Data Manager → Reporting → Asset → By Computer*

Figure 4-62 on page 188 is an example of a storage subsystem. At the right pane, you see the details of the selected extent pool. The details provide information, such as raid type or storage pool free space.

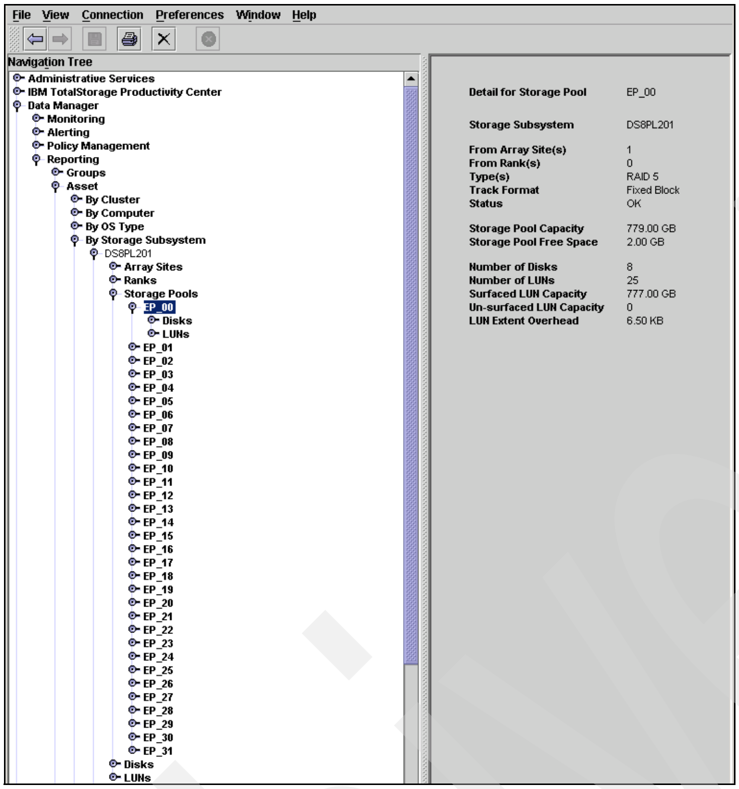

| | |
|---|---|
| Detail for Storage Pool | EP_00 |
| Storage Subsystem | DS8PL201 |
| From Array Site(s) | 1 |
| From Rank(s) | 0 |
| Type(s) | RAID 5 |
| Track Format | Fixed Block |
| Status | OK |
| Storage Pool Capacity | 779.00 GB |
| Storage Pool Free Space | 2.00 GB |
| Number of Disks | 8 |
| Number of LUNs | 25 |
| Surfaced LUN Capacity | 777.00 GB |
| Un-surfaced LUN Capacity | 0 |
| LUN Extent Overhead | 6.50 KB |

*Figure 4-62 Data Manager → Reporting → Asset → By Storage Subsystem*

Figure 4-63 shows an agent report, which is under the system-wide node. Figure 4-63 gives you information about the status, version, and connection errors of your Common Agents.

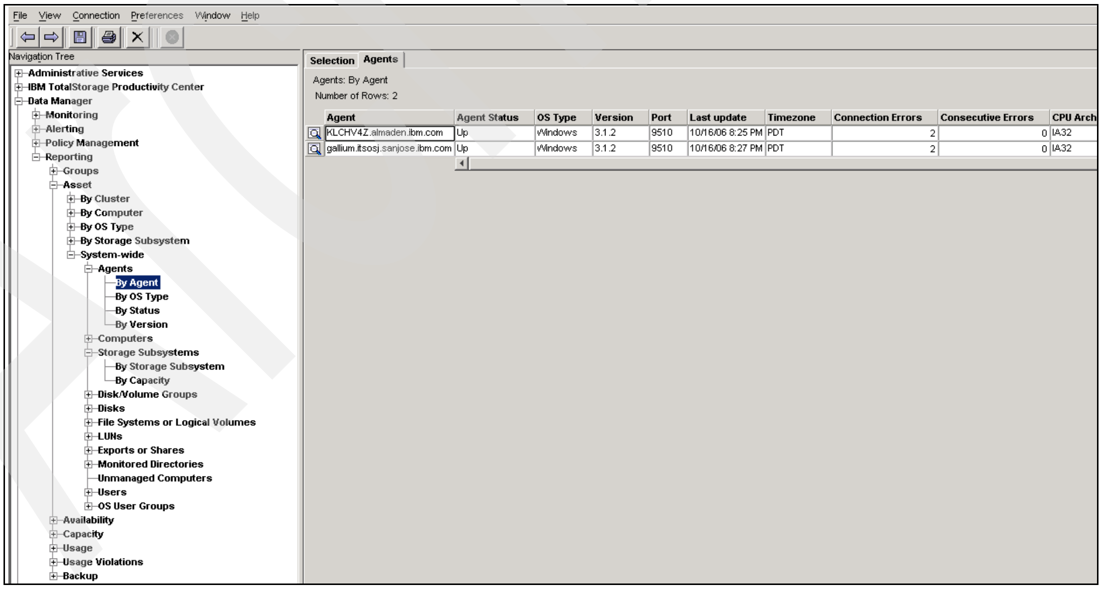

Agents: By Agent

Number of Rows: 2

| Agent | Agent Status | OS Type | Version | Port | Last update | Timezone | Connection Errors | Consecutive Errors | CPU Arch |
|---|---|---|---|---|---|---|---|---|---|
| KLCHV4Z.almaden.ibm.com | Up | Windows | 3.1.2 | 9510 | 10/16/06 8:25 PM | PDT | 2 | 0 | IA32 |
| gallium.itsosj.sanjose.ibm.com | Up | Windows | 3.1.2 | 9510 | 10/16/06 8:27 PM | PDT | 2 | 0 | IA32 |

*Figure 4-63 Agent report*

Figure 4-64 is a computer overview report.

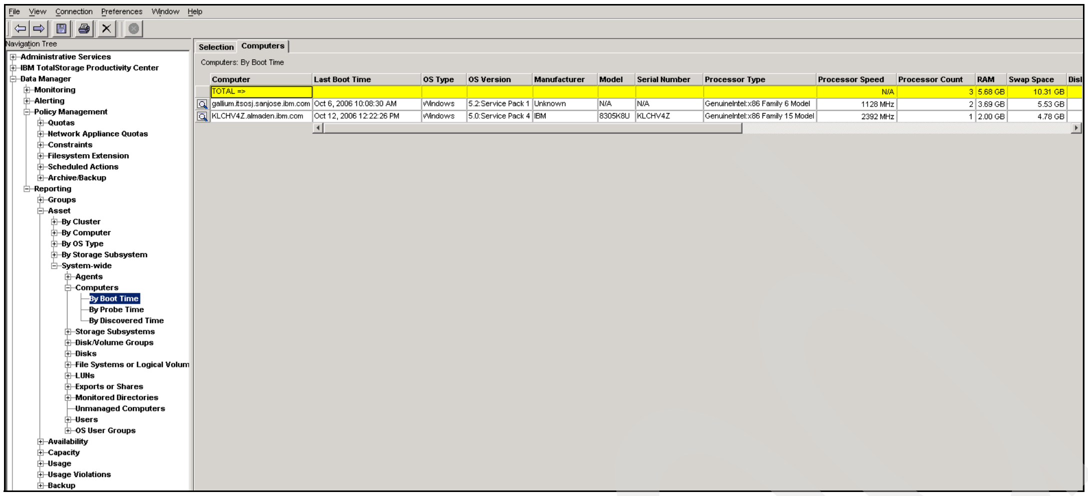

Computers: By Boot Time

| Computer | Last Boot Time | OS Type | OS Version | Manufacturer | Model | Serial Number | Processor Type | Processor Speed | Processor Count | RAM | Swap Space | Disl |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| TOTAL => | | | | | | | | N/A | 3 | 5.68 GB | 10.31 GB | |
| gallium.itsosj.sanjose.ibm.com | Oct 6, 2006 10:08:30 AM | Windows | 5.2:Service Pack 1 | Unknown | N/A | N/A | GenuineIntel:x86 Family 6 Model | 1128 MHz | 2 | 3.69 GB | 5.53 GB | |
| KLCHV4Z.almaden.ibm.com | Oct 12, 2006 12:22:26 PM | Windows | 5.0:Service Pack 4 | IBM | 8305K8U | KLCHV4Z | GenuineIntel:x86 Family 15 Model | 2392 MHz | 1 | 2.00 GB | 4.78 GB | |

Figure 4-64   Computer By Boot Time

Figure 4-65 gives you an overview of your storage subsystems and provides you information about capacities, serial numbers, firmware versions, cache, and more.

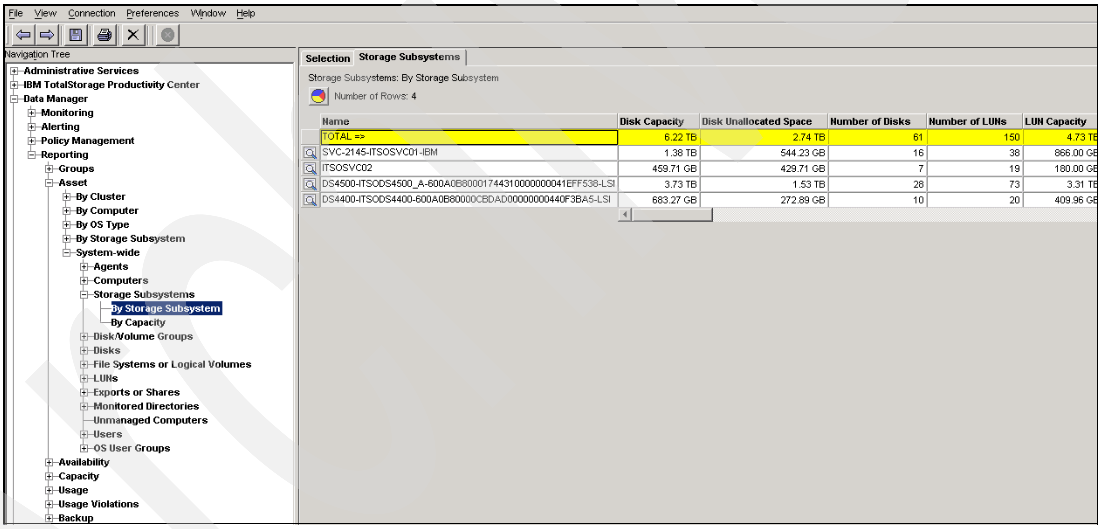

Storage Subsystems: By Storage Subsystem

Number of Rows: 4

| Name | Disk Capacity | Disk Unallocated Space | Number of Disks | Number of LUNs | LUN Capacity |
|---|---|---|---|---|---|
| TOTAL => | 6.22 TB | 2.74 TB | 61 | 150 | 4.73 TB |
| SVC-2145-ITSOSVC01-IBM | 1.38 TB | 544.23 GB | 16 | 38 | 866.00 GB |
| ITSOSVC02 | 459.71 GB | 429.71 GB | 7 | 19 | 180.00 GB |
| DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | 3.73 TB | 1.53 TB | 28 | 73 | 3.31 TB |
| DS4400-ITSODS4400-600A0B80000CBDAD0000000440F3BA5-LSI | 683.27 GB | 272.89 GB | 10 | 20 | 409.96 GB |

Figure 4-65   Report By Storage Subsystem

### Availability

Within the Availability node, you can create availability reports for computers based on pings and uptime information.

### Capacity

At this node, there are capacity-related reports for storage subsystems and computers, which might be useful to display storage trends over time. Figure 4-66 on page 190 shows the current capacity overview of two storage subsystems.

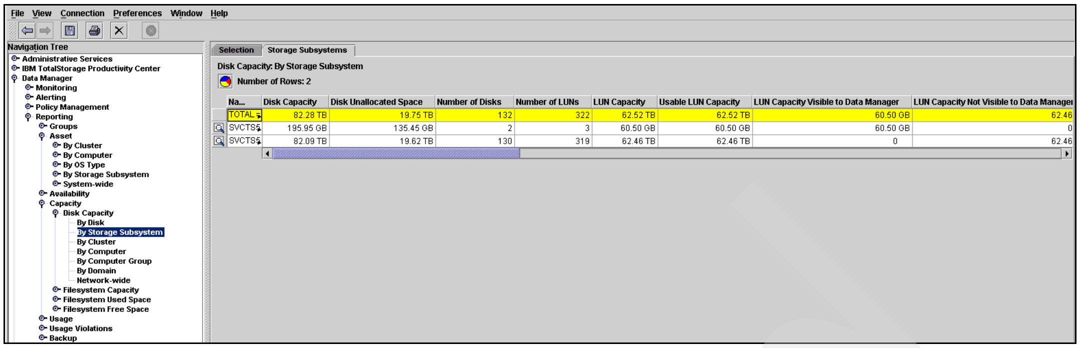

Disk Capacity: By Storage Subsystem

Number of Rows: 2

| Na... | Disk Capacity | Disk Unallocated Space | Number of Disks | Number of LUNs | LUN Capacity | Usable LUN Capacity | LUN Capacity Visible to Data Manager | LUN Capacity Not Visible to Data Manager |
|---|---|---|---|---|---|---|---|---|
| TOTAL | 82.28 TB | 19.75 TB | 132 | 322 | 62.52 TB | 62.52 TB | 60.50 GB | 62.46 |
| SVCTS5 | 195.95 GB | 135.45 GB | 2 | 3 | 60.50 GB | 60.50 GB | 60.50 GB | 0 |
| SVCTS5 | 82.09 TB | 19.62 TB | 130 | 319 | 62.46 TB | 62.46 TB | 0 | 62.46 |

*Figure 4-66 Storage Capacity By Storage Subsystem*

For historical data, select one or more subsystems (rows) and click ◕ and choose, for example, **History Chart: Free Space for selected**. Figure 4-67 shows the free space over time and the future trend, which is the dotted line.

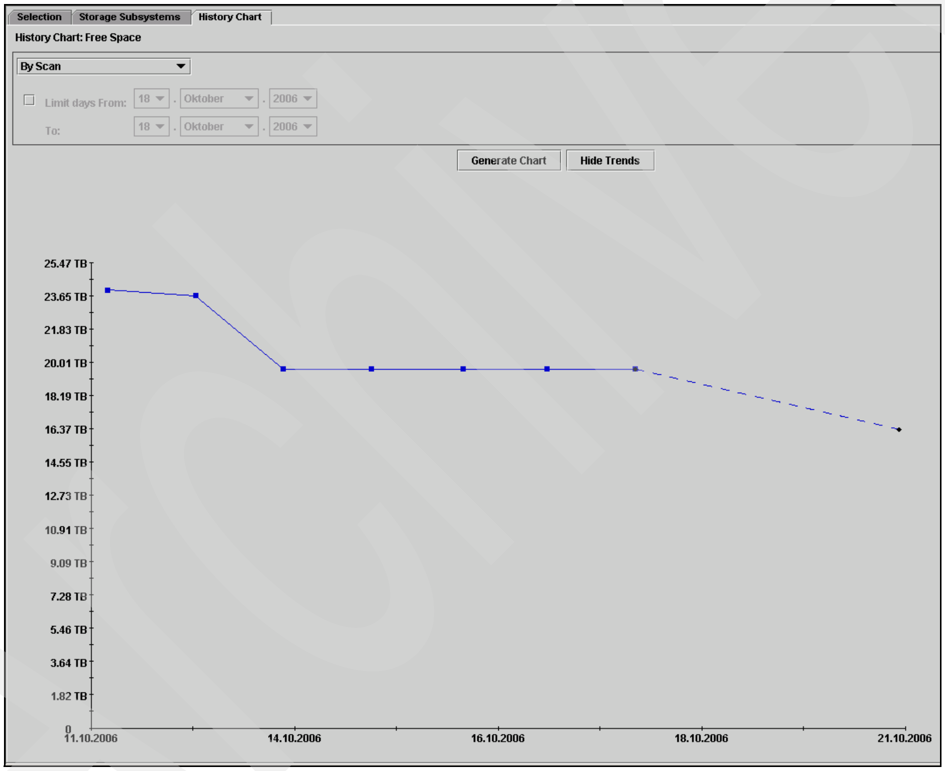

History Chart: Free Space

By Scan

□ Limit days From: 18 . Oktober . 2006

To: 18 . Oktober . 2006

*Figure 4-67 Free Space on a Subsystem*

## Usage

Usage reports are available for filesystems, directories, and users and present information, such as a file's access time, modification time, or creation time. You can display reports about the largest files in your filesystems, duplicate files, file types, orphan files, and much more.

The report example in Figure 4-68 on page 191 shows the file access distribution of filesystems (**Data Manager** → **Reporting** → **Usage** → **Access Load** → **Access Time** → **By Filesystem**).

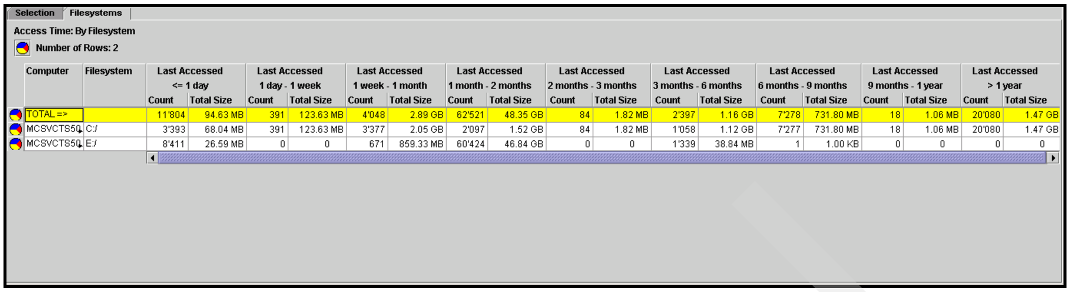

Selection | Filesystems

Access Time: By Filesystem

Number of Rows: 2

| Computer | Filesystem | Last Accessed<br><= 1 day | | Last Accessed<br>1 day - 1 week | | Last Accessed<br>1 week - 1 month | | Last Accessed<br>1 month - 2 months | | Last Accessed<br>2 months - 3 months | | Last Accessed<br>3 months - 6 months | | Last Accessed<br>6 months - 9 months | | Last Accessed<br>9 months - 1 year | | Last Accessed<br>> 1 year | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Count | Total Size | Count | Total Size | Count | Total Size | Count | Total Size | Count | Total Size | Count | Total Size | Count | Total Size | Count | Total Size | Count | Total Size |
| TOTAL => | | 11'804 | 94.63 MB | 391 | 123.63 MB | 4'048 | 2.89 GB | 62'521 | 48.35 GB | 84 | 1.82 MB | 2'397 | 1.16 GB | 7'278 | 731.80 MB | 18 | 1.06 MB | 20'080 | 1.47 GB |
| MCSVCTS50 | C:/ | 3'393 | 68.04 MB | 391 | 123.63 MB | 3'377 | 2.05 GB | 2'097 | 1.52 GB | 84 | 1.82 MB | 1'058 | 1.12 GB | 7'277 | 731.80 MB | 18 | 1.06 MB | 20'080 | 1.47 GB |
| MCSVCTS50 | E:/ | 8'411 | 26.59 MB | 0 | 0 | 671 | 859.33 MB | 60'424 | 46.84 GB | 0 | 0 | 1'339 | 38.84 MB | 1 | 1.00 KB | 0 | 0 | 0 | 0 |

*Figure 4-68 Usage access files report*

Click ● to get a pie chart as shown in Figure 4-69.

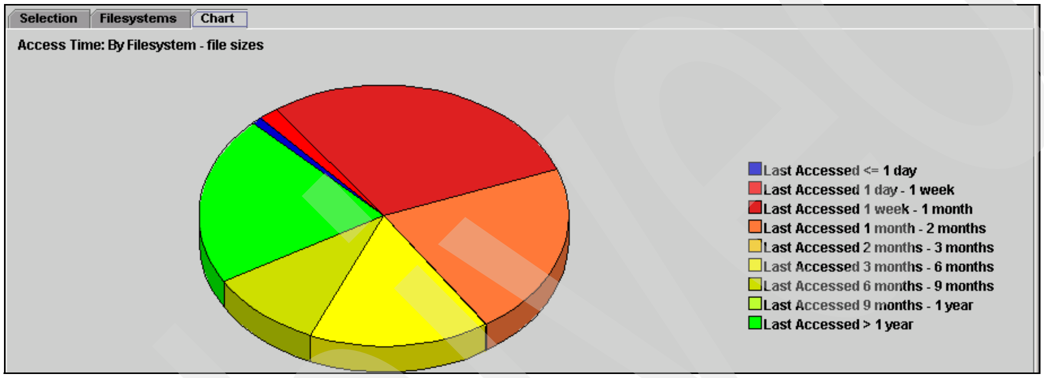

*Figure 4-69 Pie chart report*

You might use this report to decide about using hierarchical storage management (HSM) or to move rarely accessed files, directories, or filesystems to slower storage.

Look at additional useful detailed examples of the reports within Data Manager in Chapter 2, “Data management techniques” on page 13.

### 4.10.5 Disk Manager reports

The node Disk Manager provides two types of disk reports. Click **Disk Manager** → **Reporting** and then choose between **Storage Subsystems** and **Storage Subsystems Performance**. See Figure 4-70 on page 192.

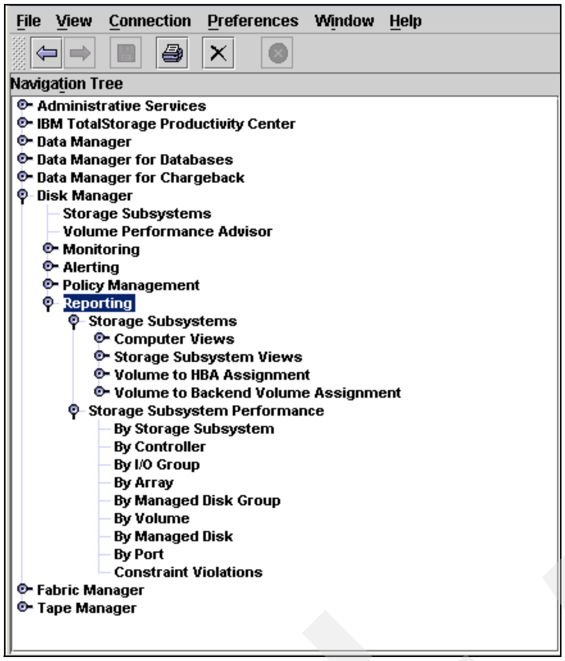

*Figure 4-70 Disk Manager reports*

## Storage subsystems

Within this category, you can choose reports about disk relationships between subsystems and computers.

> **Note:** Reports within **Disk Manager** → **Reporting** → **Computer Views** and **Disk Manager** → **Reporting** → **Storage Subsystem Views** are only available for computers with a Common Agent.

The following report example displays which disk subsystem and disk (volume) is related to a computer filesystem. Select **Disk Manager** → **Reporting** → **Storage Subsystems** → **Computer Views** → **By Filesystem/Logical Volume** as shown in Figure 4-71 on page 193.

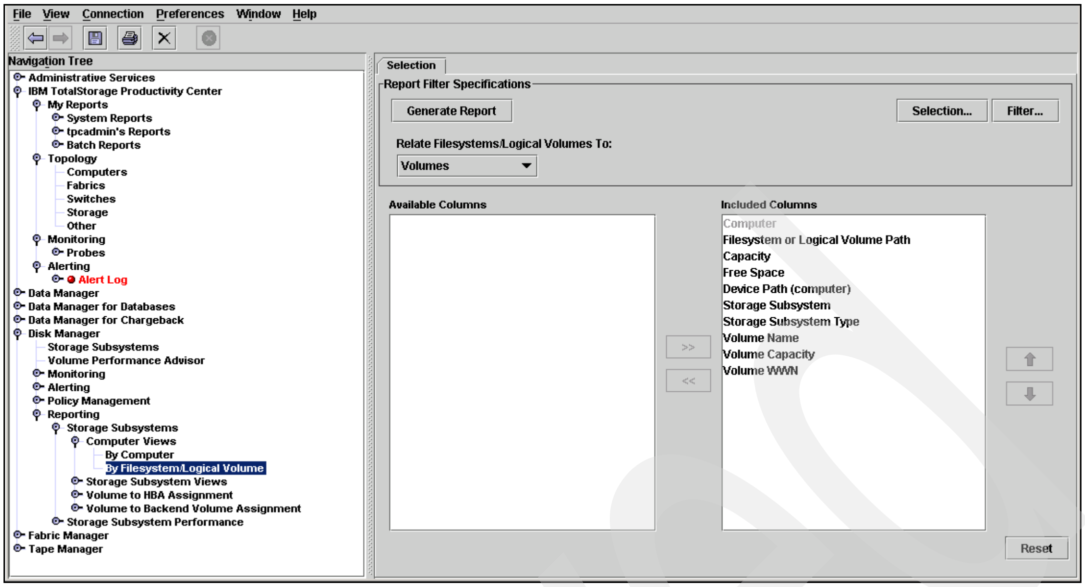

Relate Filesystems/Logical Volumes To: Volumes

Available Columns

Included Columns
Computer
Filesystem or Logical Volume Path
Capacity
Free Space
Device Path (computer)
Storage Subsystem
Storage Subsystem Type
Volume Name
Volume Capacity
Volume WWN

*Figure 4-71 Disk Manager → Reporting → Storage Subsystems → Computer Views → By Filesystem*

In the right pane, choose the relationship to **Volumes** in the **Relate Filesystems/Logical Volumes To:** drop-down menu. If you choose **Storage Subsystem**, the report shows the relationship between a filesystem and a storage subsystem. To shorten the report, you can use **Selection** to choose a set of filesystems that you want for the report. With **Filter**, you can specify more restrictions. In this example, we only want to report the relationship between a filesystem and an SVC volume (VDisk). Therefore, we select the columns **Storage Subsystem** and **SVC**, because the subsystem name starts with SVC and click **OK** (see Figure 4-72).

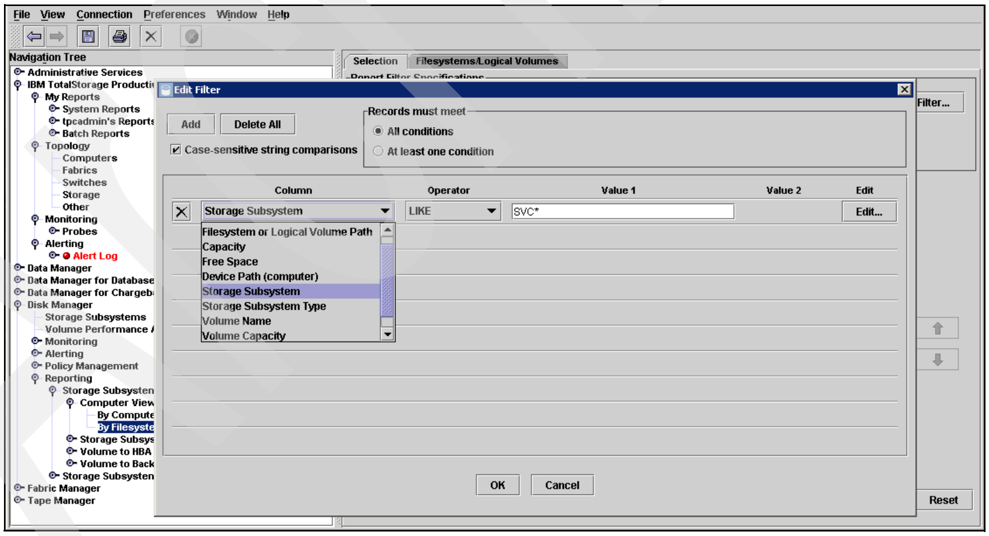

☑ Case-sensitive string comparisons

Records must meet
● All conditions
○ At least one condition

| Column | Operator | Value 1 | Value 2 | Edit |
|---|---|---|---|---|
| Storage Subsystem | LIKE | SVC* |  | Edit... |

Filesystem or Logical Volume Path
Capacity
Free Space
Device Path (computer)
Storage Subsystem
Storage Subsystem Type
Volume Name
Volume Capacity

*Figure 4-72 Edit Filter*

Finally, we generate the report by clicking **Generate Report**. The report in Figure 4-73 on page 194 shows the computer name, each filesystem, the internal disk name, the subsystem name, and finally, on which volume (VDisk) the filesystem resides.

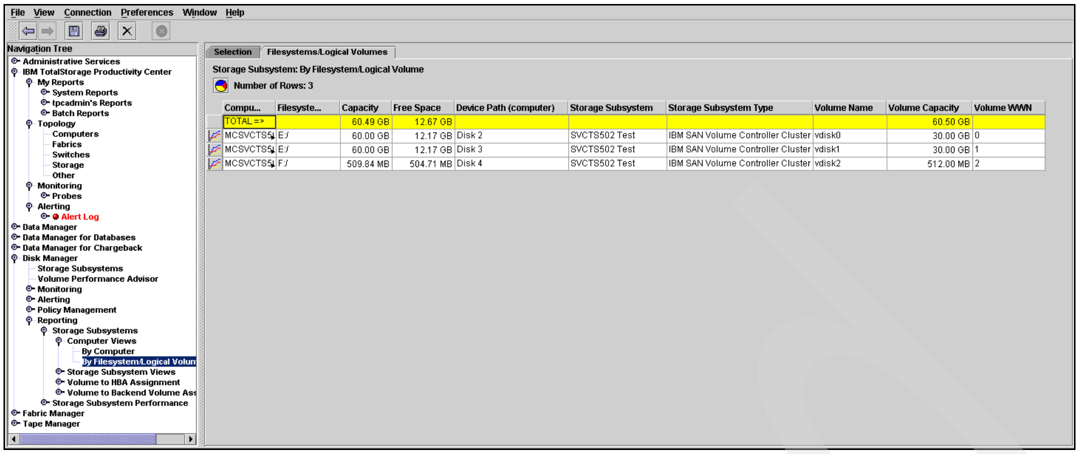

Storage Subsystem: By Filesystem/Logical Volume

Number of Rows: 3

| Compu... | Filesyste... | Capacity | Free Space | Device Path (computer) | Storage Subsystem | Storage Subsystem Type | Volume Name | Volume Capacity | Volume WWN |
|---|---|---|---|---|---|---|---|---|---|
| TOTAL => | | 60.49 GB | 12.67 GB | | | | | 60.50 GB | |
| MCSVCTS5 | E:/ | 60.00 GB | 12.17 GB | Disk 2 | SVCTS502 Test | IBM SAN Volume Controller Cluster | vdisk0 | 30.00 GB | 0 |
| MCSVCTS5 | E:/ | 60.00 GB | 12.17 GB | Disk 3 | SVCTS502 Test | IBM SAN Volume Controller Cluster | vdisk1 | 30.00 GB | 1 |
| MCSVCTS5 | F:/ | 509.84 MB | 504.71 MB | Disk 4 | SVCTS502 Test | IBM SAN Volume Controller Cluster | vdisk2 | 512.00 MB | 2 |

*Figure 4-73 Disk Manager → Reporting → Storage Subsystems → Computer Views → By Filesystem*

From this report, you can display graphical reports about the filesystems. Select a filesystem and click ◔ for historical data about filesystem capacity, or click 📈 for a filesystem free space history.

## Storage Subsystem reports

The node **Volume to HBA Assignment** contains storage subsystem-related reports. A very popular report is **Disk Manager** → **Reporting** → **Storage Subsystems** → **Volume to HBA Assignment** → **By Storage Subsystem**. The report displays storage subsystem, volume name, volume capacity, and the host assignment. For host assignment, TPC uses the **SMIS Host Alias**. This is the host name defined in the Storage Subsystem. In Figure 4-74, we select the column for the report and run **Generate Report**.

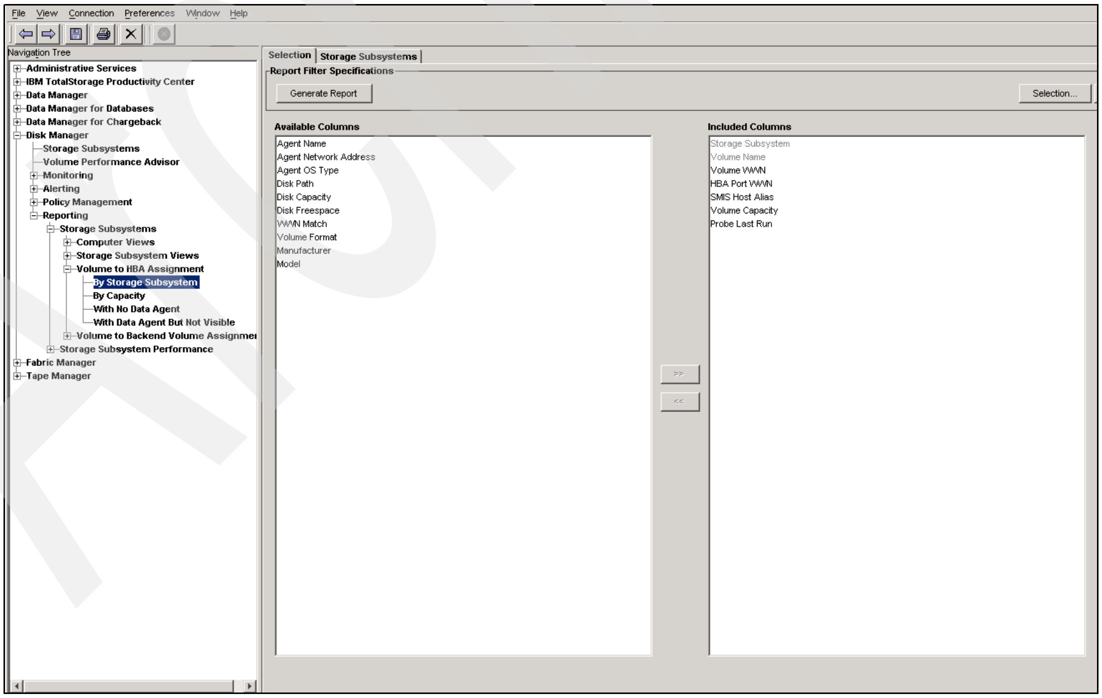

Available Columns

Agent Name
Agent Network Address
Agent OS Type
Disk Path
Disk Capacity
Disk Freespace
WWN Match
Volume Format
Manufacturer
Model

Included Columns

Storage Subsystem
Volume Name
Volume WWN
HBA Port WWN
SMIS Host Alias
Volume Capacity
Probe Last Run

*Figure 4-74 Disk Manager → Reporting → Volume to HBA Assignment → By Storage Subsystem*

Figure 4-75 displays the report. Note that for every LUN masking, there is one row. Therefore, there is one row for each computer HBA.

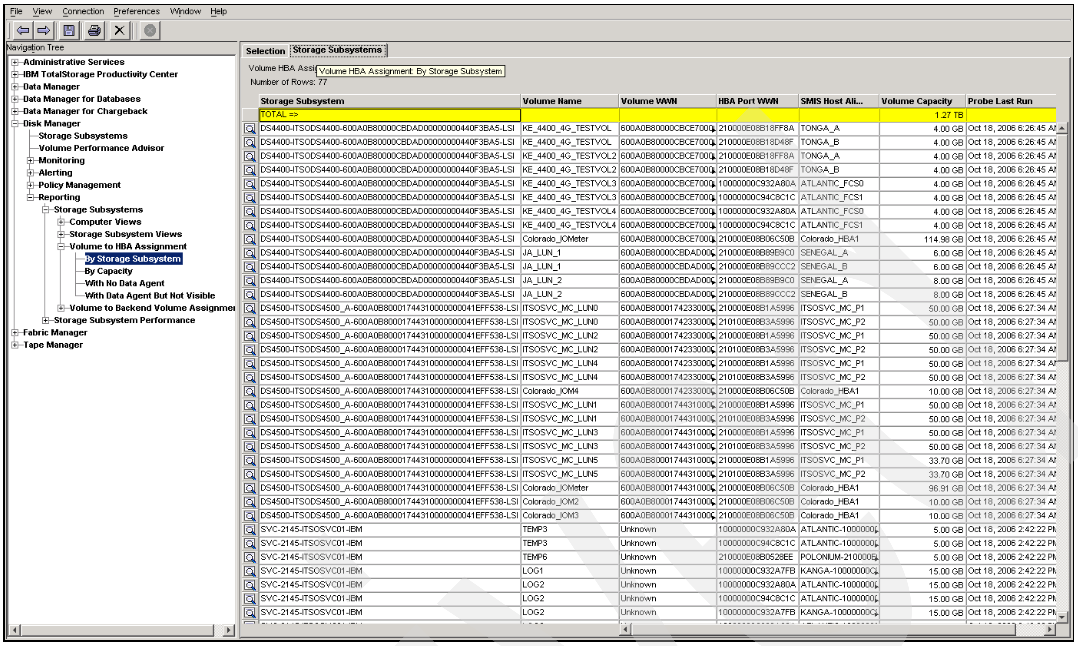

Volume HBA Assig Volume HBA Assignment: By Storage Subsystem

Number of Rows: 77

| Storage Subsystem | Volume Name | Volume WWN | HBA Port WWN | SMIS Host Ali... | Volume Capacity | Probe Last Run |
|---|---|---|---|---|---|---|
| TOTAL => | | | | | 1.27 TB | |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL | 600A0B80000CBCE7000 | 210000E08B18FF8A | TONGA_A | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL | 600A0B80000CBCE7000 | 210000E08B18D48F | TONGA_B | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL2 | 600A0B80000CBCE7000 | 210000E08B18FF8A | TONGA_A | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL2 | 600A0B80000CBCE7000 | 210000E08B18D48F | TONGA_B | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL3 | 600A0B80000CBCE7000 | 10000000C932A80A | ATLANTIC_FCS0 | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL3 | 600A0B80000CBCE7000 | 10000000C94C8C1C | ATLANTIC_FCS1 | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL4 | 600A0B80000CBCE7000 | 10000000C932A80A | ATLANTIC_FCS0 | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | KE_4400_4G_TESTVOL4 | 600A0B80000CBCE7000 | 10000000C94C8C1C | ATLANTIC_FCS1 | 4.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | Colorado_IOMeter | 600A0B80000CBCE7000 | 210000E08B06C50B | Colorado_HBA1 | 114.98 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | JA_LUN_1 | 600A0B80000CBDAD00 | 210000E08B89B9C0 | SENEGAL_A | 6.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | JA_LUN_1 | 600A0B80000CBDAD00 | 210000E08B89CCC2 | SENEGAL_B | 6.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | JA_LUN_2 | 600A0B80000CBDAD00 | 210000E08B89B9C0 | SENEGAL_A | 8.00 GB | Oct 18, 2006 6:26:45 A |
| DS4400-ITSODS4400-600A0B80000CBDAD00000000440F3BA5-LSI | JA_LUN_2 | 600A0B80000CBDAD00 | 210000E08B89CCC2 | SENEGAL_B | 8.00 GB | Oct 18, 2006 6:26:45 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN0 | 600A0B8000174233000 | 210000E08B1A5996 | ITSOSVC_MC_P1 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN0 | 600A0B8000174233000 | 210100E08B3A5996 | ITSOSVC_MC_P2 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN2 | 600A0B8000174233000 | 210000E08B1A5996 | ITSOSVC_MC_P1 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN2 | 600A0B8000174233000 | 210100E08B3A5996 | ITSOSVC_MC_P2 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN4 | 600A0B8000174233000 | 210000E08B1A5996 | ITSOSVC_MC_P1 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN4 | 600A0B8000174233000 | 210100E08B3A5996 | ITSOSVC_MC_P2 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | Colorado_IOM4 | 600A0B8000174233000 | 210000E08B06C50B | Colorado_HBA1 | 10.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN1 | 600A0B8000174431000 | 210000E08B1A5996 | ITSOSVC_MC_P1 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN1 | 600A0B8000174431000 | 210100E08B3A5996 | ITSOSVC_MC_P2 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN3 | 600A0B8000174431000 | 210000E08B1A5996 | ITSOSVC_MC_P1 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN3 | 600A0B8000174431000 | 210100E08B3A5996 | ITSOSVC_MC_P2 | 50.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN5 | 600A0B8000174431000 | 210000E08B1A5996 | ITSOSVC_MC_P1 | 33.70 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | ITSOSVC_MC_LUN5 | 600A0B8000174431000 | 210100E08B3A5996 | ITSOSVC_MC_P2 | 33.70 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | Colorado_IOMeter | 600A0B8000174431000 | 210000E08B06C50B | Colorado_HBA1 | 96.91 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | Colorado_IOM2 | 600A0B8000174431000 | 210000E08B06C50B | Colorado_HBA1 | 10.00 GB | Oct 18, 2006 6:27:34 A |
| DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI | Colorado_IOM3 | 600A0B8000174431000 | 210000E08B06C50B | Colorado_HBA1 | 10.00 GB | Oct 18, 2006 6:27:34 A |
| SVC-2145-ITSOSVC01-IBM | TEMP3 | Unknown | 10000000C932A80A | ATLANTIC-1000000 | 5.00 GB | Oct 18, 2006 2:42:22 P |
| SVC-2145-ITSOSVC01-IBM | TEMP3 | Unknown | 10000000C94C8C1C | ATLANTIC-1000000 | 5.00 GB | Oct 18, 2006 2:42:22 P |
| SVC-2145-ITSOSVC01-IBM | TEMP6 | Unknown | 210000E08B0528EE | POLONIUM-210000E | 5.00 GB | Oct 18, 2006 2:42:22 P |
| SVC-2145-ITSOSVC01-IBM | LOG1 | Unknown | 10000000C932A7FB | KANGA-10000000C | 15.00 GB | Oct 18, 2006 2:42:22 P |
| SVC-2145-ITSOSVC01-IBM | LOG2 | Unknown | 10000000C932A80A | ATLANTIC-1000000 | 15.00 GB | Oct 18, 2006 2:42:22 P |
| SVC-2145-ITSOSVC01-IBM | LOG2 | Unknown | 10000000C94C8C1C | ATLANTIC-1000000 | 15.00 GB | Oct 18, 2006 2:42:22 P |
| SVC-2145-ITSOSVC01-IBM | LOG2 | Unknown | 10000000C932A7FB | KANGA-10000000C | 15.00 GB | Oct 18, 2006 2:42:22 P |

*Figure 4-75 Disk Manager → Reporting → Volume to HBA Assignment → By Storage Subsystem*

After you generate a report, you can save it (**File** → **Save As**) or export the report (**File** → **Export Data**) to a delimited text or HTML file. Or, you can use the print function (**File** → **Print**) as shown in Figure 4-76. The output destination can be a printer, HTML file, CSV file, or a text-formatted file. You have these choices with every report.

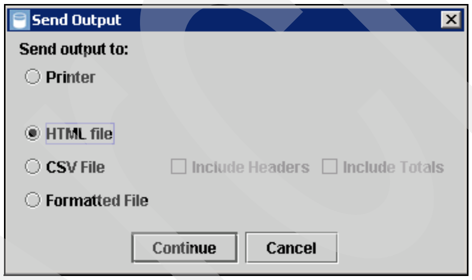

Send output to:

○ Printer

● HTML file

○ CSV File □ Include Headers □ Include Totals

○ Formatted File

*Figure 4-76 File → Print*

Selecting **Disk Manager** → **Reporting** → **Storage Subsystems** → **Volume to Backend Volume Assignment** → **By Volume** gives you a report with information about the relationship between a VDisk (SVC) to MDisk Group (SVC) and MDisk (SVC) as shown in Figure 4-77 on page 196.

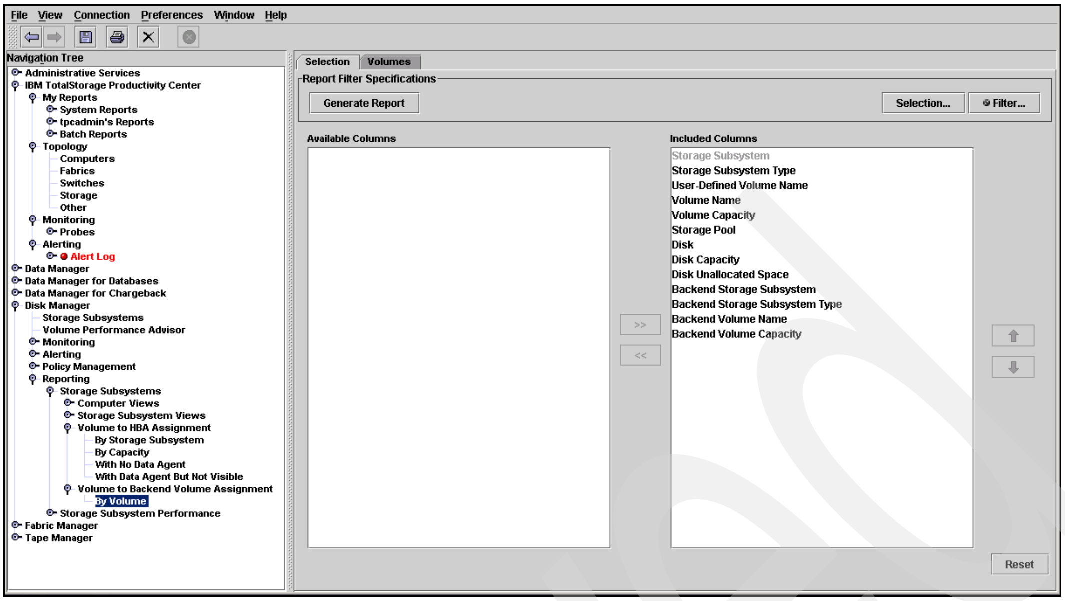

Included Columns

- Storage Subsystem
- Storage Subsystem Type
- User-Defined Volume Name
- Volume Name
- Volume Capacity
- Storage Pool
- Disk
- Disk Capacity
- Disk Unallocated Space
- Backend Storage Subsystem
- Backend Storage Subsystem Type
- Backend Volume Name
- Backend Volume Capacity

*Figure 4-77 Disk Manager → Reporting → Volume to Backend Volume Assignment*

Refine your report by clicking **Filter** (see Figure 4-78) and clicking **OK**.

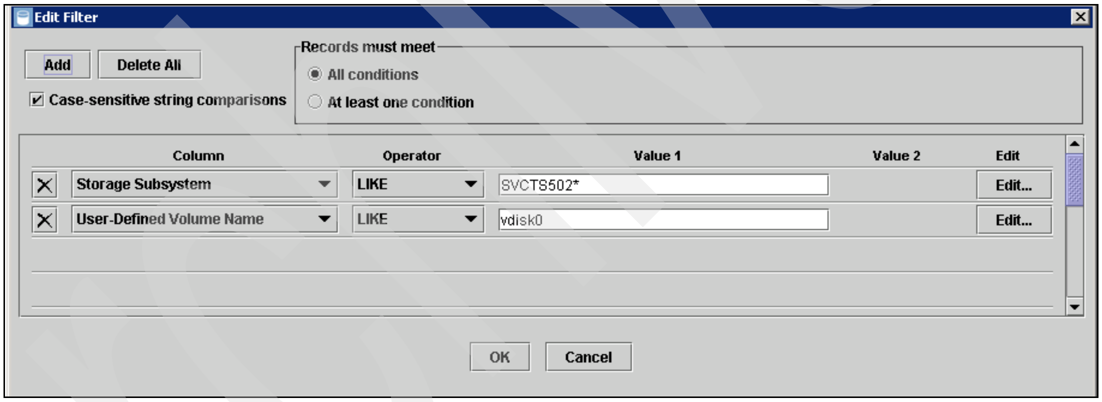

☑ Case-sensitive string comparisons

Records must meet
● All conditions
○ At least one condition

| | Column | Operator | Value 1 | Value 2 | Edit |
|---|---|---|---|---|---|
| ☒ | Storage Subsystem | LIKE | SVCTS502* | | Edit... |
| ☒ | User-Defined Volume Name | LIKE | vdisk0 | | Edit... |

*Figure 4-78 Edit Filter*

Generate the report. Figure 4-79 shows the report. The capacity of vdisk0 in the storage subsystem SVCTS502 Test is 30 GB. The VDisk is spread over mdisk0 and mdisk1. Both MDisks have a capacity of 98.68 GB; mdisk0 has 67.98 GB of free space and mdisk1 has 67.48 GB of free space.

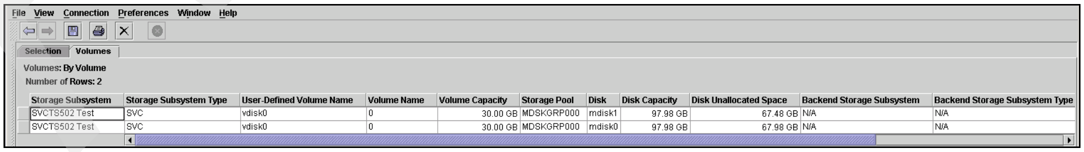

Volumes: By Volume
Number of Rows: 2

| Storage Subsystem | Storage Subsystem Type | User-Defined Volume Name | Volume Name | Volume Capacity | Storage Pool | Disk | Disk Capacity | Disk Unallocated Space | Backend Storage Subsystem | Backend Storage Subsystem Type |
|---|---|---|---|---|---|---|---|---|---|---|
| SVCTS502 Test | SVC | vdisk0 | 0 | 30.00 GB | MDSKGRP000 | mdisk1 | 97.98 GB | 67.48 GB | N/A | N/A |
| SVCTS502 Test | SVC | vdisk0 | 0 | 30.00 GB | MDSKGRP000 | mdisk0 | 97.98 GB | 67.98 GB | N/A | N/A |

*Figure 4-79 Report: Storage Subsystems → Volume to Backend Volume Assignment*

### *Storage subsystem performance*

Within the node **Disk Manager** → **Reporting** → **Storage Subsystem Performance**, you can create reports about subsystem performance on several levels. Figure 4-80 shows the performance reporting levels. With TPC 3.1.3, you additionally have the node **By Node**.

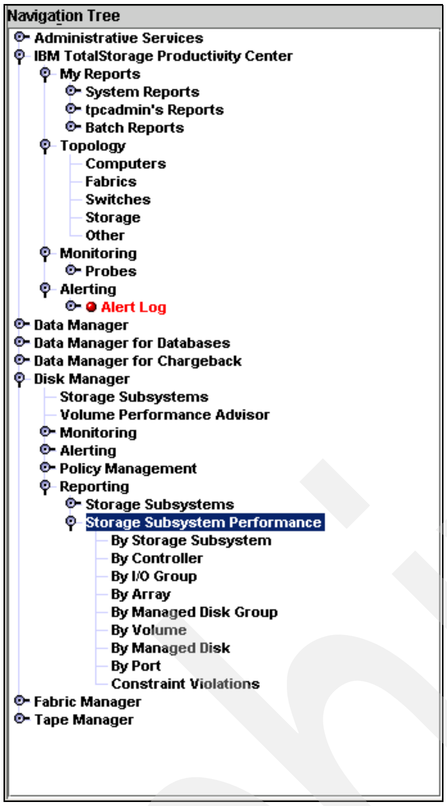

*Figure 4-80 Storage Subsystem reporting overview*

The nodes **By I/O Group**, **By Node**, **By Managed Disk Group**, and **By Managed Disk** belong to SVC.

The following example report, `By I/O Group (SVC)`, shows the read performance from an SVC I/O Group. Select **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **By I/O Group**.

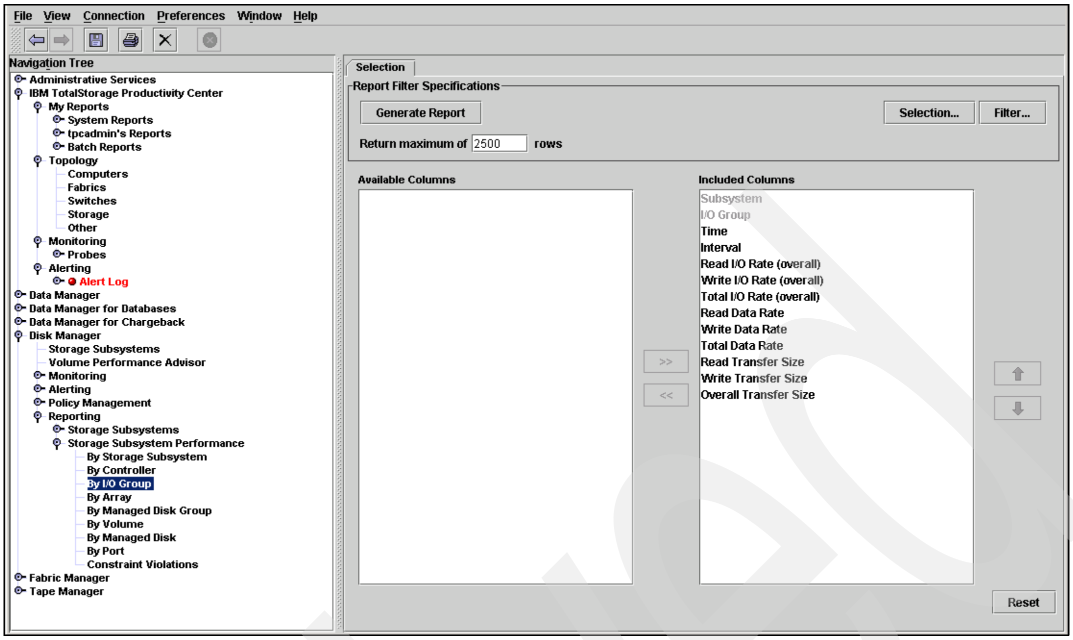

Return maximum of 2500 rows

Available Columns

Included Columns
Subsystem
I/O Group
Time
Interval
Read I/O Rate (overall)
Write I/O Rate (overall)
Total I/O Rate (overall)
Read Data Rate
Write Data Rate
Total Data Rate
Read Transfer Size
Write Transfer Size
Overall Transfer Size

*Figure 4-81 Disk Manager → Reporting → Storage Subsystem Performance → By I/O Group*

Click **Generate Report** and select the first row. The performance data that you see in the table in Figure 4-82 is based on the newest collected performance sample. The sample time is shown in the column `Time`. The column `Interval` shows the collecting interval time. In this case, the samples are collected in a 15 minute interval (900 seconds), which also means that the performance samples are based on a average value over 15 minutes.

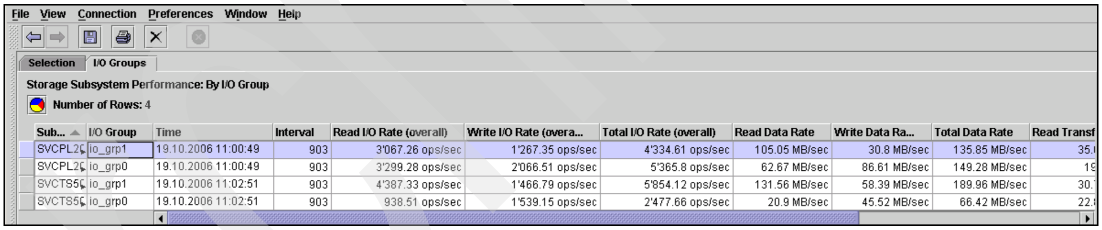

Storage Subsystem Performance: By I/O Group

Number of Rows: 4

| Sub... | I/O Group | Time | Interval | Read I/O Rate (overall) | Write I/O Rate (overa... | Total I/O Rate (overall) | Read Data Rate | Write Data Ra... | Total Data Rate | Read Transf |
|---|---|---|---|---|---|---|---|---|---|---|
| SVCPL2( | io_grp1 | 19.10.2006 11:00:49 | 903 | 3'067.26 ops/sec | 1'267.35 ops/sec | 4'334.61 ops/sec | 105.05 MB/sec | 30.8 MB/sec | 135.85 MB/sec | 35.( |
| SVCPL2( | io_grp0 | 19.10.2006 11:00:49 | 903 | 3'299.28 ops/sec | 2'066.51 ops/sec | 5'365.8 ops/sec | 62.67 MB/sec | 86.61 MB/sec | 149.28 MB/sec | 19 |
| SVCTS5( | io_grp1 | 19.10.2006 11:02:51 | 903 | 4'387.33 ops/sec | 1'466.79 ops/sec | 5'854.12 ops/sec | 131.56 MB/sec | 58.39 MB/sec | 189.96 MB/sec | 30.7 |
| SVCTS5( | io_grp0 | 19.10.2006 11:02:51 | 903 | 938.51 ops/sec | 1'539.15 ops/sec | 2'477.66 ops/sec | 20.9 MB/sec | 45.52 MB/sec | 66.42 MB/sec | 22.! |

*Figure 4-82 Disk Manager → Reporting → Storage Subsystem Performance → By I/O Group*

In this example, we want to see the read data rate over the time. To generate a historical performance report, select the desired I/O Group and click ◕ . Select **Read Data Rate** and click **OK** (see Figure 4-83 on page 199).

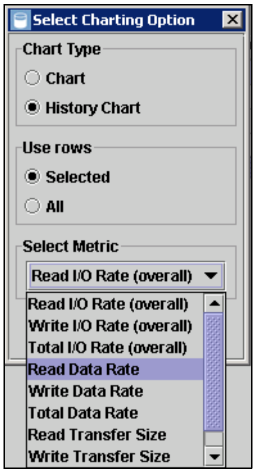
Select Charting Option

Chart Type
- ○ Chart
- ● History Chart

Use rows
- ● Selected
- ○ All

Select Metric: Read I/O Rate (overall)

*Figure 4-83 Performance report: select metric Read Data Rate*

Figure 4-84 shows the performance report **Read Data Rate** for an I/O Group. Every collected sample is marked with a blue quad. The time range is limited from 16 October 2006 to 19 October 2006. The grid line is the calculated trend projected in the future. Select the drop-down menu to change the resolution from **By Sample** to **Hourly** and click **Generate Chart** to refresh the graph.

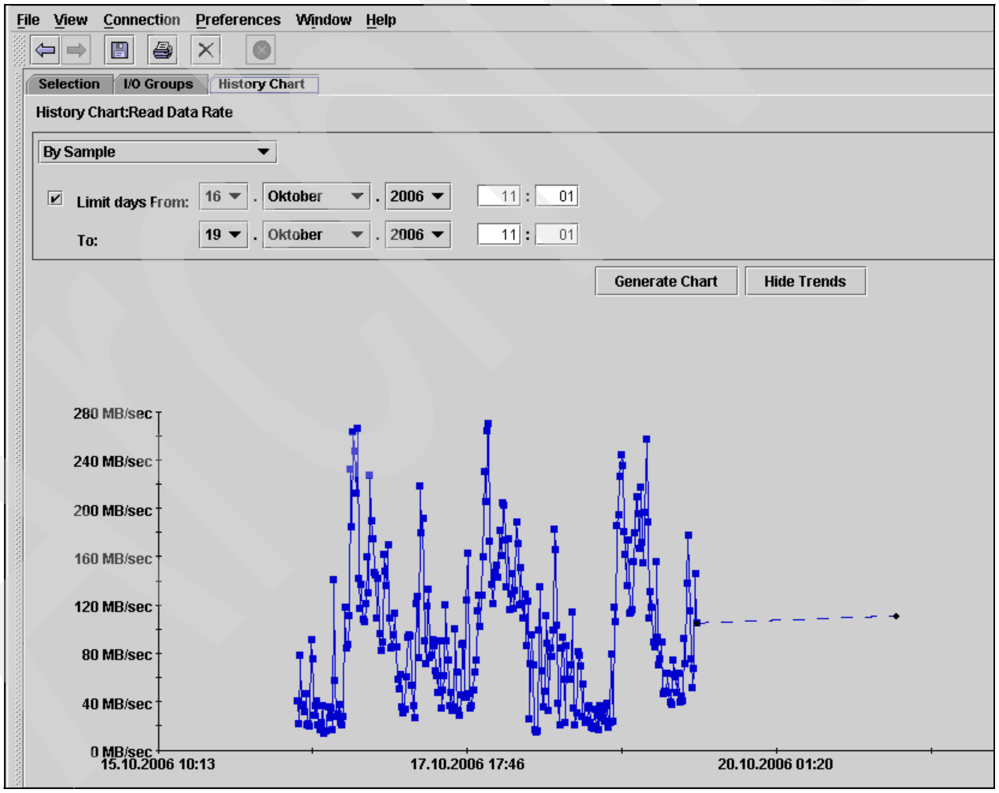

History Chart:Read Data Rate

By Sample

☑ Limit days From: 16 . Oktober . 2006 11 : 01

To: 19 . Oktober . 2006 11 : 01

*Figure 4-84 Performance graph with sample resolution*

Figure 4-85 on page 200 shows the graph with **Hourly** performance values, which means that every blue quad in the graph is an average of data read performance over an hour. By hovering over a quad with your mouse, you get the performance value and time.

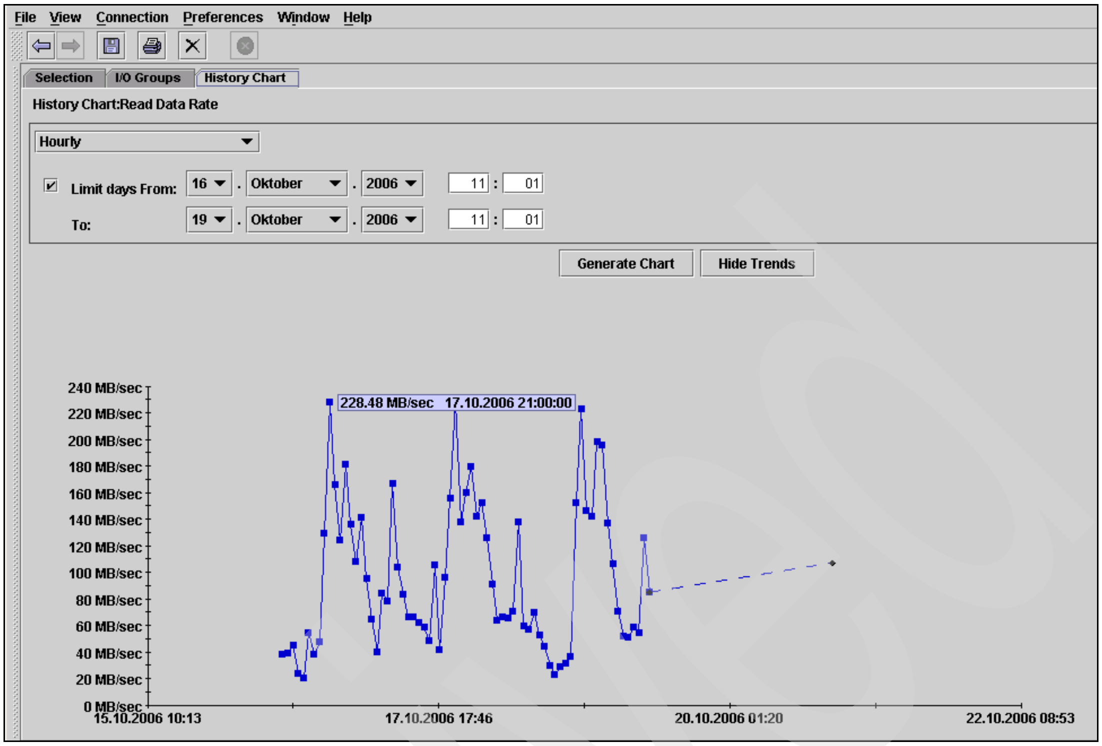

History Chart:Read Data Rate

Hourly

☑ Limit days From: 16 . Oktober . 2006 11 : 01

To: 19 . Oktober . 2006 11 : 01

*Figure 4-85 Performance graph with hourly resolution*

The largest resolution is the **Daily** performance data (average read performance over a day), which you might use for long term monitoring. Therefore, we disabled the time limitation. See Figure 4-86.

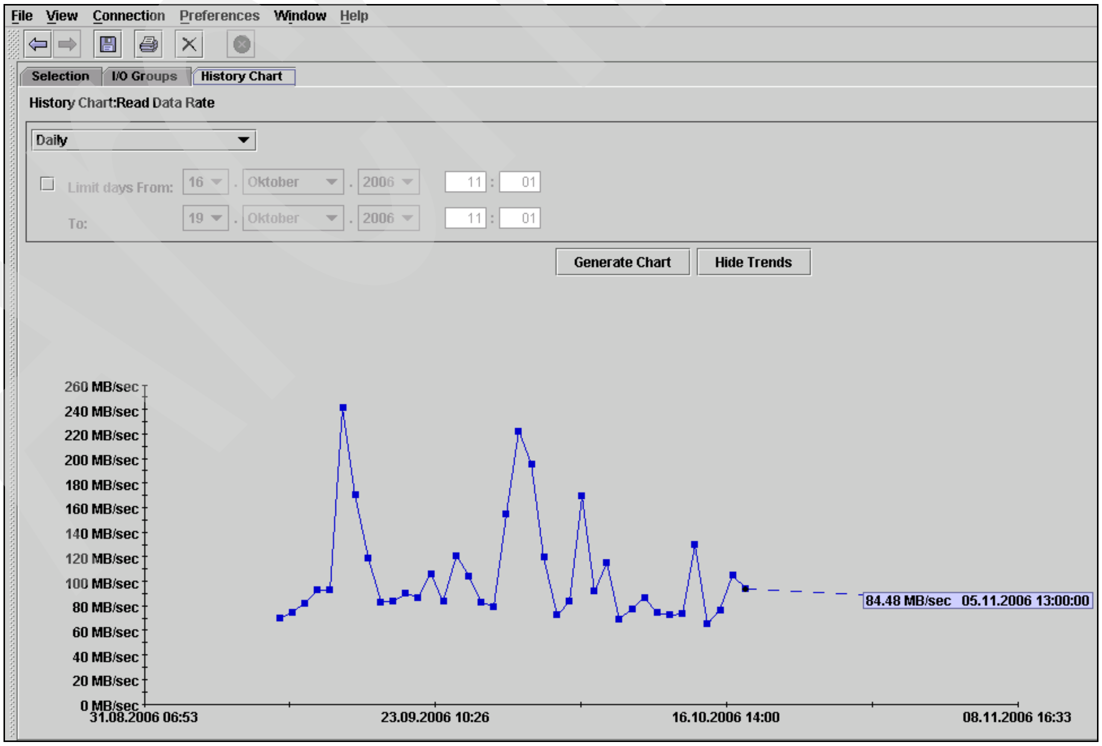

History Chart:Read Data Rate

Daily

☐ Limit days From: 16 . Oktober . 2006 11 : 01

To: 19 . Oktober . 2006 11 : 01

*Figure 4-86 Performance graph with daily resolution*

To export the report, you can use **File** → **Export Data**. The export file is a comma-delimited file or comma-separated value text file (CSV) with optional headers. Or, you can export the report with **File** → **Print**. You can choose between graphical export, such as PDF or HTML, or text export, such as CSV or plain text (see Figure 4-87).

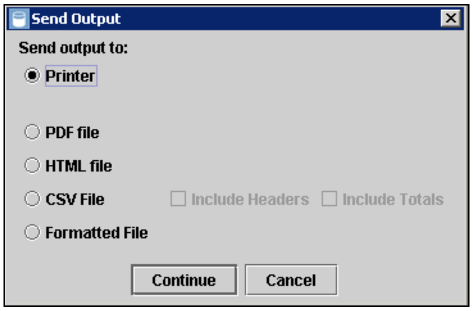

Send output to:

● Printer

○ PDF file

○ HTML file

○ CSV File □ Include Headers □ Include Totals

○ Formatted File

*Figure 4-87 Choose output device*

### Violation Performance Report

Another example is a Violation Performance Report. To see this report, select **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **Constraint Violations**. This report is based on performance threshold violations. You can define the thresholds (alerts) at **Disk Manager** → **Alerting** → **Storage Subsystem Alerts**. Every violation of a performance alert triggers a constraint violation. You can define the following performance alerts (TPC Version 3.1.2):

- Overall Port Response Time Threshold
- Total Port Data Rate Threshold
- Total Port I/O Rate Threshold
- Cache Holding Time Threshold
- NVS Full Percentage Threshold
- Total Data Rate Threshold
- Total I/O Rate Threshold
- Overall Backend Response Time Threshold
- Total Backend Data Rate Threshold
- Total Backend I/O Rate Threshold
- Disk Utilization Percentage Threshold

To generate the report, select **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **Constraint Violations**.

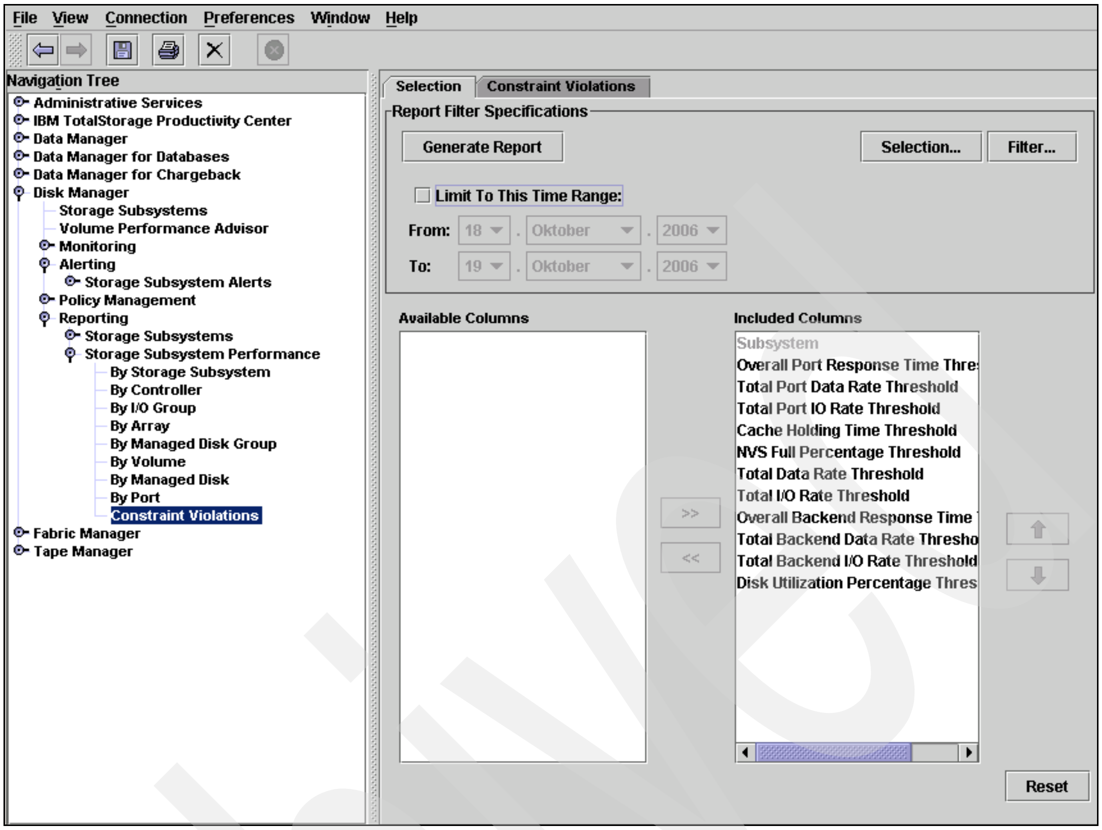

Navigation Tree

- Administrative Services
- IBM TotalStorage Productivity Center
- Data Manager
- Data Manager for Databases
- Data Manager for Chargeback
- Disk Manager
  - Storage Subsystems
  - Volume Performance Advisor
  - Monitoring
  - Alerting
    - Storage Subsystem Alerts
  - Policy Management
  - Reporting
    - Storage Subsystems
    - Storage Subsystem Performance
      - By Storage Subsystem
      - By Controller
      - By I/O Group
      - By Array
      - By Managed Disk Group
      - By Volume
      - By Managed Disk
      - By Port
      - Constraint Violations
- Fabric Manager
- Tape Manager

Report Filter Specifications

□ Limit To This Time Range:

From: 18 . Oktober . 2006

To: 19 . Oktober . 2006

Available Columns

Included Columns

Subsystem
Overall Port Response Time Thre:
Total Port Data Rate Threshold
Total Port IO Rate Threshold
Cache Holding Time Threshold
NVS Full Percentage Threshold
Total Data Rate Threshold
Total I/O Rate Threshold
Overall Backend Response Time
Total Backend Data Rate Thresho
Total Backend I/O Rate Threshold
Disk Utilization Percentage Thres

*Figure 4-88 Performance report: Constraint Violations*

Click **Generate Report**. Figure 4-89 shows a summary of constraint violations. The first column shows the storage subsystem, which is to the left of the fields shown in Figure 4-89. Every row belongs to one subsystem. Every column represents a performance alert threshold. The numbers in the cells are the counts of performance samples that violated the constraint. For example, the collected performance samples for the subsystem in the first row violated the `Overall Backend Response Time Threshold` 24 times.

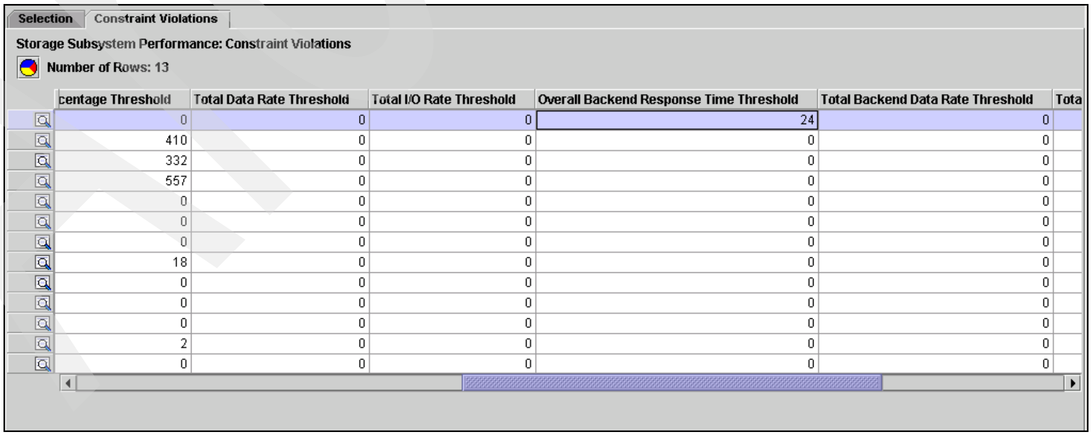

Storage Subsystem Performance: Constraint Violations

Number of Rows: 13

| | centage Threshold | Total Data Rate Threshold | Total I/O Rate Threshold | Overall Backend Response Time Threshold | Total Backend Data Rate Threshold | Tota |
|---|---|---|---|---|---|---|
| 🔍 | 0 | 0 | 0 | 24 | 0 | |
| 🔍 | 410 | 0 | 0 | 0 | 0 | |
| 🔍 | 332 | 0 | 0 | 0 | 0 | |
| 🔍 | 557 | 0 | 0 | 0 | 0 | |
| 🔍 | 0 | 0 | 0 | 0 | 0 | |
| 🔍 | 0 | 0 | 0 | 0 | 0 | |
| 🔍 | 0 | 0 | 0 | 0 | 0 | |
| 🔍 | 18 | 0 | 0 | 0 | 0 | |
| 🔍 | 0 | 0 | 0 | 0 | 0 | |
| 🔍 | 0 | 0 | 0 | 0 | 0 | |
| 🔍 | 0 | 0 | 0 | 0 | 0 | |
| 🔍 | 2 | 0 | 0 | 0 | 0 | |
| 🔍 | 0 | 0 | 0 | 0 | 0 | |

*Figure 4-89 Constraint violations overview*

By clicking ◕, you can display a graphical overview (chart) of constraint violations for the selected subsystem as shown in Figure 4-90 on page 203.

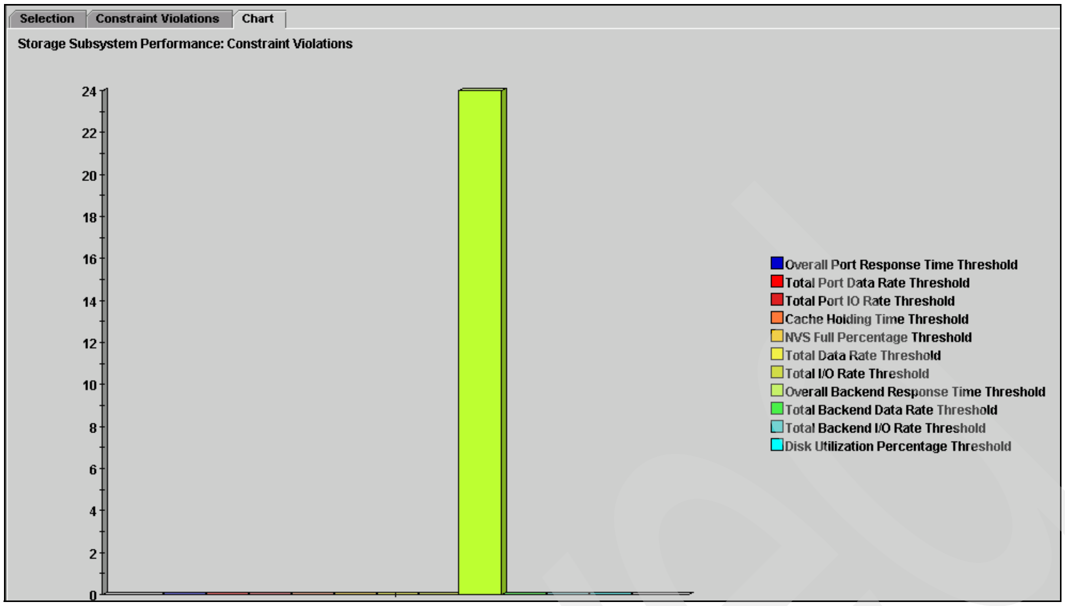

*Figure 4-90 Graphical overview of constraint violations of a selected subsystem*

Click the magnifying glass icon 🔍 on the **Constraint Violations** tab to get a list of all constraint violations for the selected subsystem. Figure 4-91 displays a detailed list with timestamps and components. For example in the first row, we see that the `mdisk_098` (SVC) at `08.10.2006 14:15:56` took an Overall Backend Response Time of `35.5 ms`. If the response time is longer than 35 ms, it creates an alert and a constraint violation.

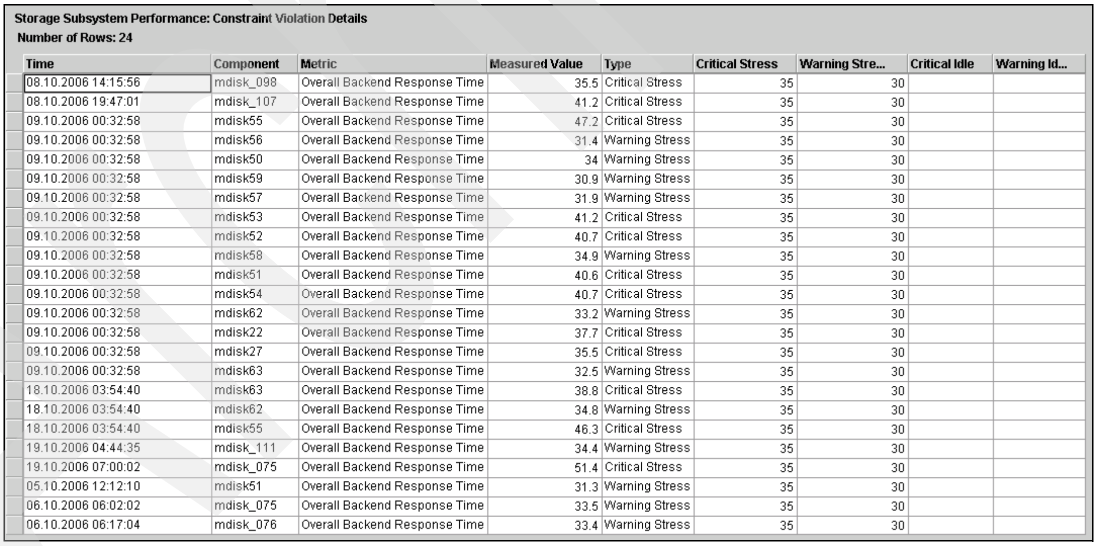

Storage Subsystem Performance: Constraint Violation Details

Number of Rows: 24

| Time | Component | Metric | Measured Value | Type | Critical Stress | Warning Stre... | Critical Idle | Warning Id... |
|---|---|---|---|---|---|---|---|---|
| 08.10.2006 14:15:56 | mdisk_098 | Overall Backend Response Time | 35.5 | Critical Stress | 35 | 30 | | |
| 08.10.2006 19:47:01 | mdisk_107 | Overall Backend Response Time | 41.2 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk55 | Overall Backend Response Time | 47.2 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk56 | Overall Backend Response Time | 31.4 | Warning Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk50 | Overall Backend Response Time | 34 | Warning Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk59 | Overall Backend Response Time | 30.9 | Warning Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk57 | Overall Backend Response Time | 31.9 | Warning Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk53 | Overall Backend Response Time | 41.2 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk52 | Overall Backend Response Time | 40.7 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk58 | Overall Backend Response Time | 34.9 | Warning Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk51 | Overall Backend Response Time | 40.6 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk54 | Overall Backend Response Time | 40.7 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk62 | Overall Backend Response Time | 33.2 | Warning Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk22 | Overall Backend Response Time | 37.7 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk27 | Overall Backend Response Time | 35.5 | Critical Stress | 35 | 30 | | |
| 09.10.2006 00:32:58 | mdisk63 | Overall Backend Response Time | 32.5 | Warning Stress | 35 | 30 | | |
| 18.10.2006 03:54:40 | mdisk63 | Overall Backend Response Time | 38.8 | Critical Stress | 35 | 30 | | |
| 18.10.2006 03:54:40 | mdisk62 | Overall Backend Response Time | 34.8 | Warning Stress | 35 | 30 | | |
| 18.10.2006 03:54:40 | mdisk55 | Overall Backend Response Time | 46.3 | Critical Stress | 35 | 30 | | |
| 19.10.2006 04:44:35 | mdisk_111 | Overall Backend Response Time | 34.4 | Warning Stress | 35 | 30 | | |
| 19.10.2006 07:00:02 | mdisk_075 | Overall Backend Response Time | 51.4 | Critical Stress | 35 | 30 | | |
| 05.10.2006 12:12:10 | mdisk51 | Overall Backend Response Time | 31.3 | Warning Stress | 35 | 30 | | |
| 06.10.2006 06:02:02 | mdisk_075 | Overall Backend Response Time | 33.5 | Warning Stress | 35 | 30 | | |
| 06.10.2006 06:17:04 | mdisk_076 | Overall Backend Response Time | 33.4 | Warning Stress | 35 | 30 | | |

*Figure 4-91 Report of all constraint violations for a subsystem*

Certain metrics, such as NVS Full Percentage Threshold and Disk Utilization Percentage Threshold, have an additional report. Figure 4-92 on page 204 is an example of a Constraint Violation Details report for an ESS.

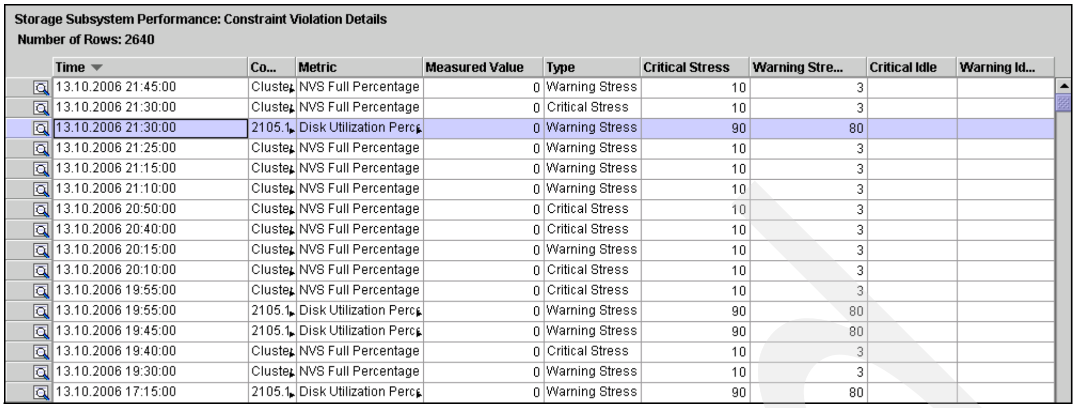

Storage Subsystem Performance: Constraint Violation Details
Number of Rows: 2640

| | Time ▼ | Co... | Metric | Measured Value | Type | Critical Stress | Warning Stre... | Critical Idle | Warning Id... |
|---|---|---|---|---|---|---|---|---|---|
| 🔍 | 13.10.2006 21:45:00 | Cluste | NVS Full Percentage | 0 | Warning Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 21:30:00 | Cluste | NVS Full Percentage | 0 | Critical Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 21:30:00 | 2105.1 | Disk Utilization Perc | 0 | Warning Stress | 90 | 80 | | |
| 🔍 | 13.10.2006 21:25:00 | Cluste | NVS Full Percentage | 0 | Warning Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 21:15:00 | Cluste | NVS Full Percentage | 0 | Warning Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 21:10:00 | Cluste | NVS Full Percentage | 0 | Warning Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 20:50:00 | Cluste | NVS Full Percentage | 0 | Critical Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 20:40:00 | Cluste | NVS Full Percentage | 0 | Critical Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 20:15:00 | Cluste | NVS Full Percentage | 0 | Warning Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 20:10:00 | Cluste | NVS Full Percentage | 0 | Critical Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 19:55:00 | Cluste | NVS Full Percentage | 0 | Critical Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 19:55:00 | 2105.1 | Disk Utilization Perc | 0 | Warning Stress | 90 | 80 | | |
| 🔍 | 13.10.2006 19:45:00 | 2105.1 | Disk Utilization Perc | 0 | Warning Stress | 90 | 80 | | |
| 🔍 | 13.10.2006 19:40:00 | Cluste | NVS Full Percentage | 0 | Critical Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 19:30:00 | Cluste | NVS Full Percentage | 0 | Warning Stress | 10 | 3 | | |
| 🔍 | 13.10.2006 17:15:00 | 2105.1 | Disk Utilization Perc | 0 | Warning Stress | 90 | 80 | | |

*Figure 4-92 Report of all constraint violations for a subsystem*

There is a magnifying glass icon 🔍 to the left of the threshold violations. By clicking the icon, you get the following pane. From this pane (Figure 4-93), you can now generate a report on the specific component (specific rank).

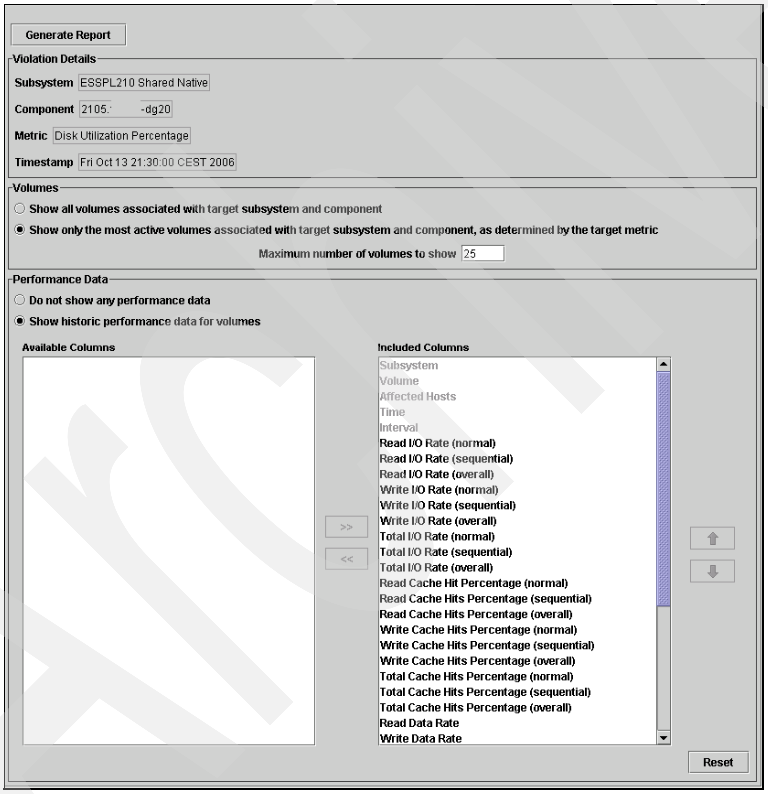

Violation Details

Subsystem: ESSPL210 Shared Native
Component: 2105.   -dg20
Metric: Disk Utilization Percentage
Timestamp: Fri Oct 13 21:30:00 CEST 2006

Volumes

○ Show all volumes associated with target subsystem and component
● Show only the most active volumes associated with target subsystem and component, as determined by the target metric

Maximum number of volumes to show: 25

Performance Data

○ Do not show any performance data
● Show historic performance data for volumes

Available Columns

Included Columns
Subsystem
Volume
Affected Hosts
Time
Interval
Read I/O Rate (normal)
Read I/O Rate (sequential)
Read I/O Rate (overall)
Write I/O Rate (normal)
Write I/O Rate (sequential)
Write I/O Rate (overall)
Total I/O Rate (normal)
Total I/O Rate (sequential)
Total I/O Rate (overall)
Read Cache Hit Percentage (normal)
Read Cache Hits Percentage (sequential)
Read Cache Hits Percentage (overall)
Write Cache Hits Percentage (normal)
Write Cache Hits Percentage (sequential)
Write Cache Hits Percentage (overall)
Total Cache Hits Percentage (normal)
Total Cache Hits Percentage (sequential)
Total Cache Hits Percentage (overall)
Read Data Rate
Write Data Rate

*Figure 4-93 Additional report based on the constraint violation*

The resulting report (Figure 4-94 on page 205) shows the most active volumes (a maximum of 25 in this example) on this rank.

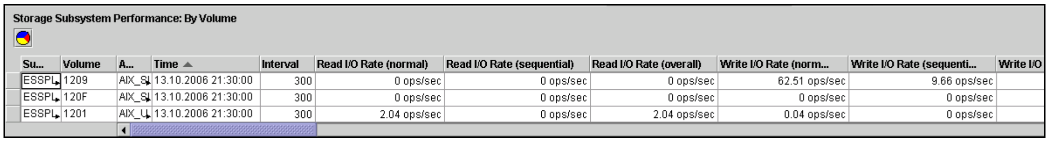

Storage Subsystem Performance: By Volume

| Su... | Volume | A... | Time | Interval | Read I/O Rate (normal) | Read I/O Rate (sequential) | Read I/O Rate (overall) | Write I/O Rate (norm... | Write I/O Rate (sequenti... | Write I/O |
|---|---|---|---|---|---|---|---|---|---|---|
| ESSPL | 1209 | AIX_S | 13.10.2006 21:30:00 | 300 | 0 ops/sec | 0 ops/sec | 0 ops/sec | 62.51 ops/sec | 9.66 ops/sec | |
| ESSPL | 120F | AIX_S | 13.10.2006 21:30:00 | 300 | 0 ops/sec | 0 ops/sec | 0 ops/sec | 0 ops/sec | 0 ops/sec | |
| ESSPL | 1201 | AIX_U | 13.10.2006 21:30:00 | 300 | 2.04 ops/sec | 0 ops/sec | 2.04 ops/sec | 0.04 ops/sec | 0 ops/sec | |

*Figure 4-94 Affected volumes on the specific rank*

Click to get the historical performance data of those volumes. The Disk Utilization Threshold was violated at `13. October` at `21:30`. Therefore, we adjusted the reporting time around this window of time. In Figure 4-94, we can see that at the time of the threshold violation, many more writes than reads exist. Therefore, we suspect that the number of writes makes the rank busy. The following graph (Figure 4-95) shows the Cache to Disk Transfer Rate. We can see a peak at 21:30. The volume `1209` produces the most I/Os at this time.

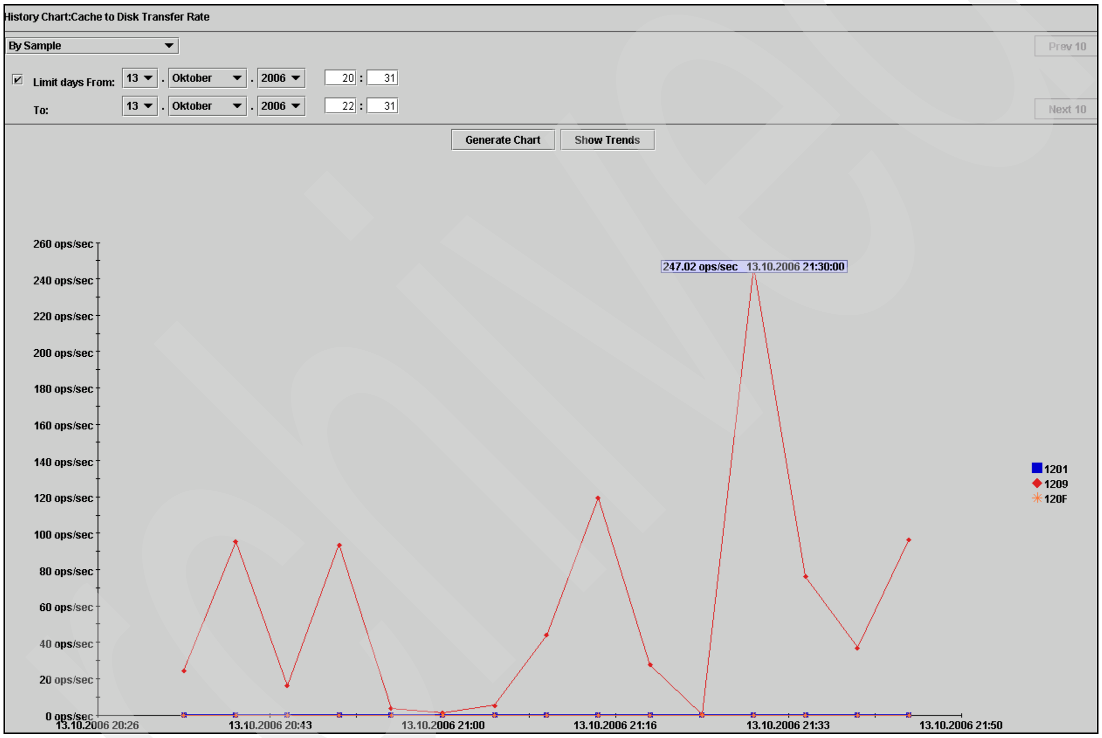

History Chart:Cache to Disk Transfer Rate

By Sample

☑ Limit days From: 13 . Oktober . 2006 20 : 31

To: 13 . Oktober . 2006 22 : 31

*Figure 4-95 Cache to Disk Transfer Rate*

### NVS Full Percentage report

Another useful report is the NVS Full Percentage report.

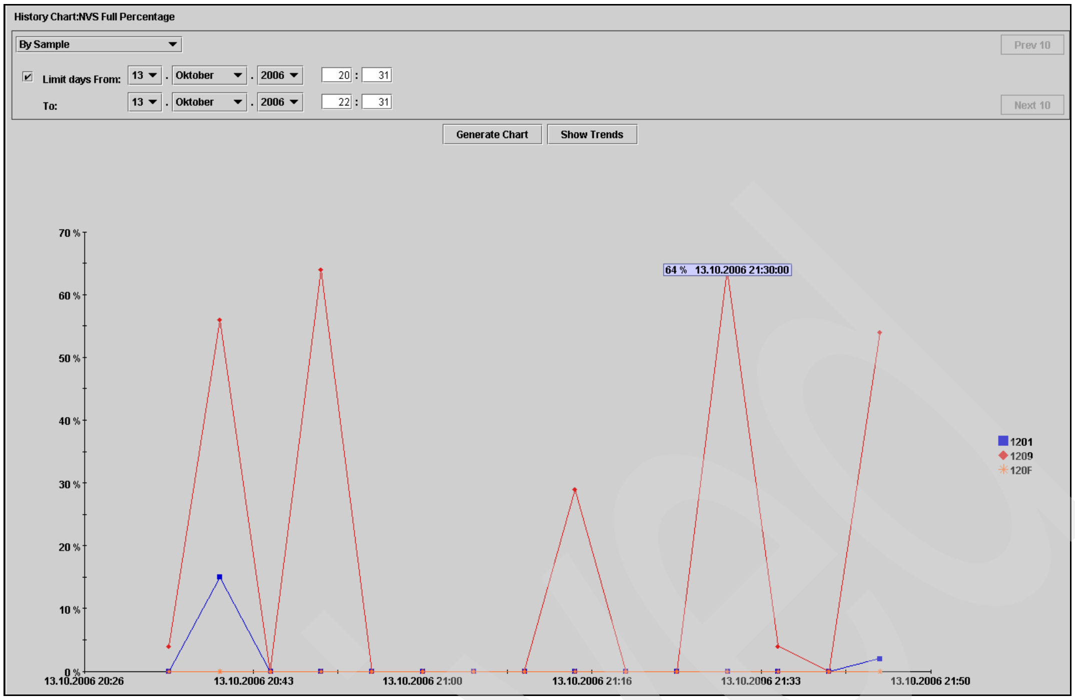

History Chart:NVS Full Percentage

By Sample

☑ Limit days From: 13 . Oktober . 2006 20 : 31

To: 13 . Oktober . 2006 22 : 31

*Figure 4-96 NVS Full Percentage*

We know now that, in this example, the volumes 1209, 120F, and 1201 were affected at 13. October 21:30 with a increased rank (disk) utilization and the volume 1209 was the most active volume. Furthermore, we can analyze many more performance metrics. But ultimately, the storage subsystem clients decide what is sufficient performance.

Performance analysis is a challenging topic because of the influence of many diverse factors. TPC provides reports to support your analysis of performance issues.

### 4.10.6 Fabric Manager reports

The node **Fabric Manager** → **Reporting** → **Switch Performance** provides two types of reporting. You can generate performance reports, including error statistics, that are based on the SAN switch port performance samples, or you can generate a report of constraint violations. Every time that a performance threshold on a switchport is violated, a constraint violation is created (see Figure 4-97 on page 207).

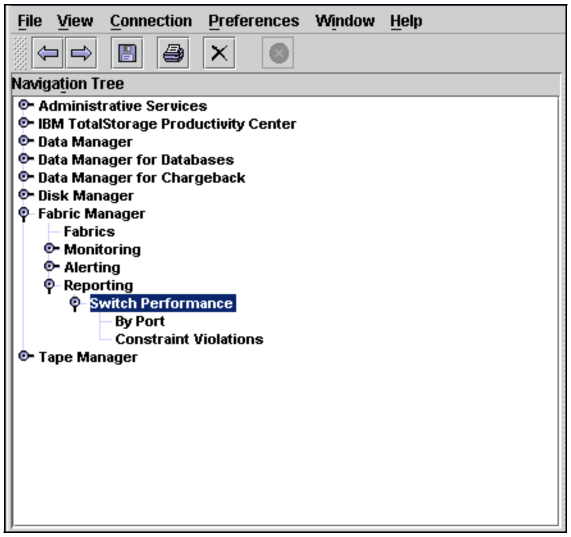

*Figure 4-97 Switch Performance reporting*

Figure 4-98 is an example of a Switch Performance report (**Fabric Manager** → **Reporting** → **Switch Performance** → **By Port**).

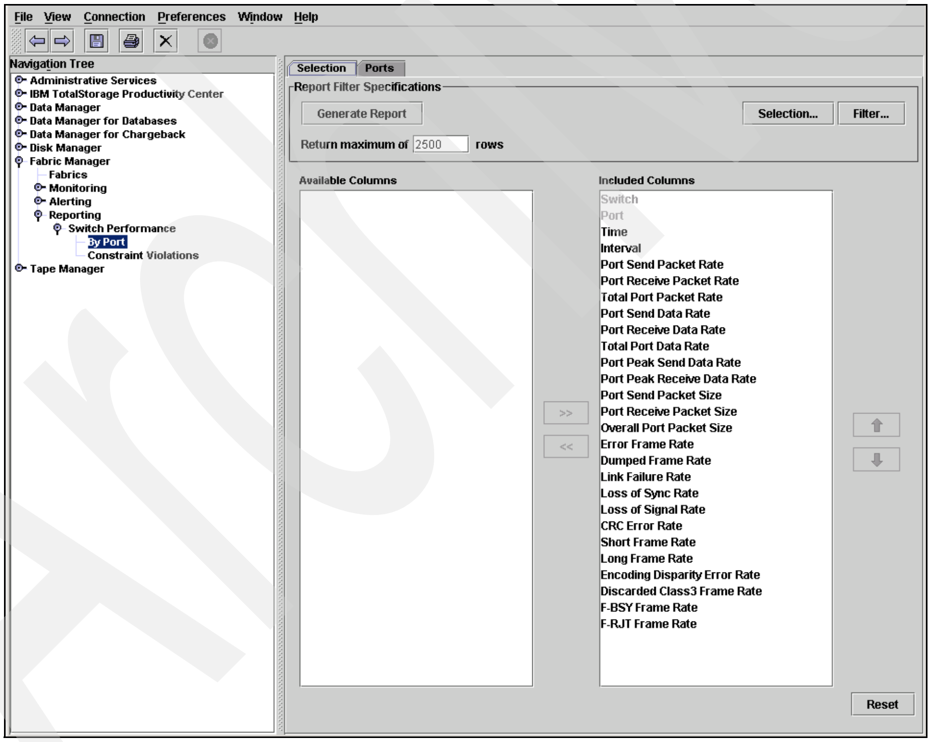

Return maximum of 2500 rows

Available Columns

Included Columns

Switch
Port
Time
Interval
Port Send Packet Rate
Port Receive Packet Rate
Total Port Packet Rate
Port Send Data Rate
Port Receive Data Rate
Total Port Data Rate
Port Peak Send Data Rate
Port Peak Receive Data Rate
Port Send Packet Size
Port Receive Packet Size
Overall Port Packet Size
Error Frame Rate
Dumped Frame Rate
Link Failure Rate
Loss of Sync Rate
Loss of Signal Rate
CRC Error Rate
Short Frame Rate
Long Frame Rate
Encoding Disparity Error Rate
Discarded Class3 Frame Rate
F-BSY Frame Rate
F-RJT Frame Rate

*Figure 4-98 Performance Monitor By Port*

The right pane (see Figure 4-99 on page 208) displays the available performance metrics for switch ports. Click **Generate Report** to create the report. This report is a tabular view containing the last performance samples at the specific time that you see in the column `Time`.

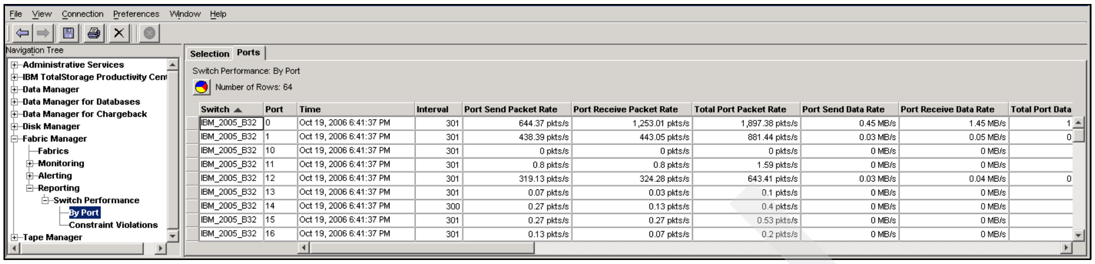

Switch Performance: By Port

Number of Rows: 64

| Switch ▲ | Port | Time | Interval | Port Send Packet Rate | Port Receive Packet Rate | Total Port Packet Rate | Port Send Data Rate | Port Receive Data Rate | Total Port Data |
|---|---|---|---|---|---|---|---|---|---|
| IBM_2005_B32 | 0 | Oct 19, 2006 6:41:37 PM | 301 | 644.37 pkts/s | 1,253.01 pkts/s | 1,897.38 pkts/s | 0.45 MB/s | 1.45 MB/s | 1 |
| IBM_2005_B32 | 1 | Oct 19, 2006 6:41:37 PM | 301 | 438.39 pkts/s | 443.05 pkts/s | 881.44 pkts/s | 0.03 MB/s | 0.05 MB/s | 0 |
| IBM_2005_B32 | 10 | Oct 19, 2006 6:41:37 PM | 301 | 0 pkts/s | 0 pkts/s | 0 pkts/s | 0 MB/s | 0 MB/s | |
| IBM_2005_B32 | 11 | Oct 19, 2006 6:41:37 PM | 301 | 0.8 pkts/s | 0.8 pkts/s | 1.59 pkts/s | 0 MB/s | 0 MB/s | |
| IBM_2005_B32 | 12 | Oct 19, 2006 6:41:37 PM | 301 | 319.13 pkts/s | 324.28 pkts/s | 643.41 pkts/s | 0.03 MB/s | 0.04 MB/s | 0 |
| IBM_2005_B32 | 13 | Oct 19, 2006 6:41:37 PM | 301 | 0.07 pkts/s | 0.03 pkts/s | 0.1 pkts/s | 0 MB/s | 0 MB/s | |
| IBM_2005_B32 | 14 | Oct 19, 2006 6:41:37 PM | 300 | 0.27 pkts/s | 0.13 pkts/s | 0.4 pkts/s | 0 MB/s | 0 MB/s | |
| IBM_2005_B32 | 15 | Oct 19, 2006 6:41:37 PM | 301 | 0.27 pkts/s | 0.27 pkts/s | 0.53 pkts/s | 0 MB/s | 0 MB/s | |
| IBM_2005_B32 | 16 | Oct 19, 2006 6:41:37 PM | 301 | 0.13 pkts/s | 0.07 pkts/s | 0.2 pkts/s | 0 MB/s | 0 MB/s | |

*Figure 4-99 SAN Switch Performance By Port report*

Select the port of interest and click to generate the historical data of one or more ports. In this example, we display the historical **Port Send Data Rate** for all ports (see Figure 4-100).

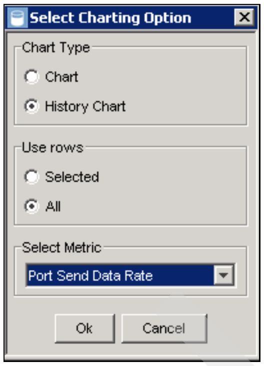

Chart Type

○ Chart

● History Chart

Use rows

○ Selected

● All

Select Metric

Port Send Data Rate

*Figure 4-100 Select Charting Option*

Click **OK**. Figure 4-101 on page 209 shows performance graphs (Port Send Data Rate) for 10 ports. Every square is a performance sample (measuring point). Longer parts of the graph without a square mean that there are no performance samples available at this time.

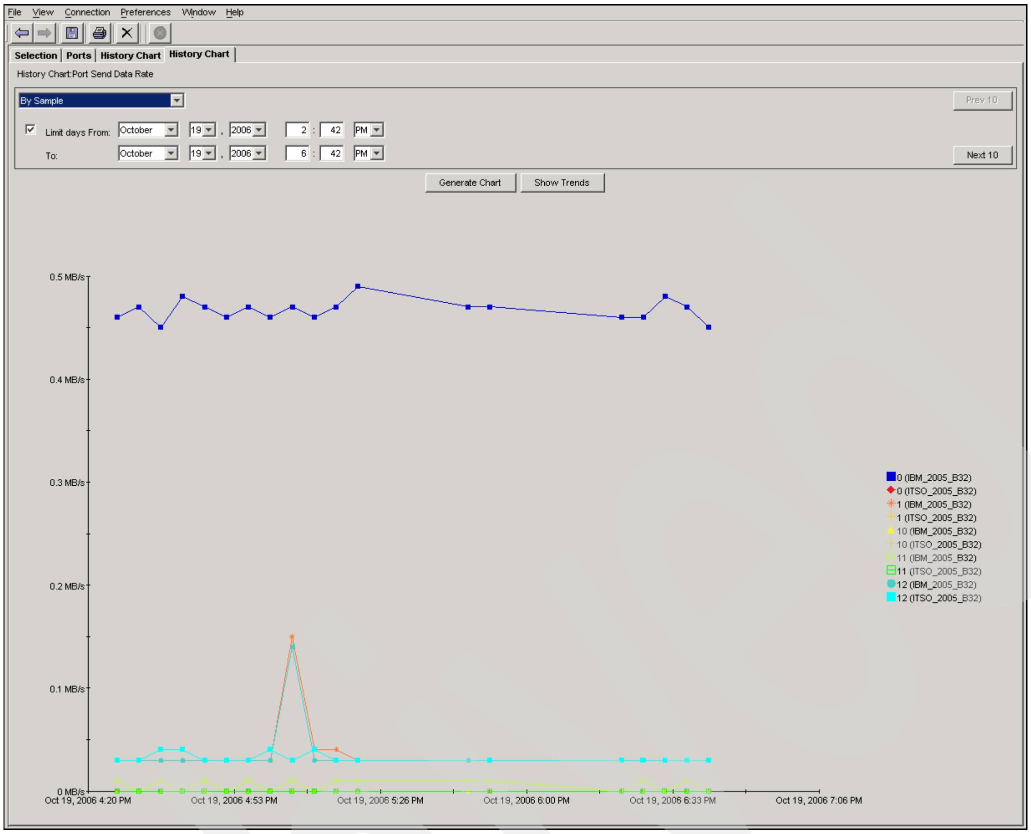

History Chart:Port Send Data Rate

By Sample

☑ Limit days From: October 19 , 2006 2 : 42 PM

To: October 19 , 2006 6 : 42 PM

*Figure 4-101 Performance Graph of Port Send Data Rate*

The other type of performance reporting is **Fabric Manager → Reporting → Switch Performance → Constraint Violations**. Every performance alert threshold violation causes a constraint violation. **Included Columns** on the right lists the potential constraints (Figure 4-102).

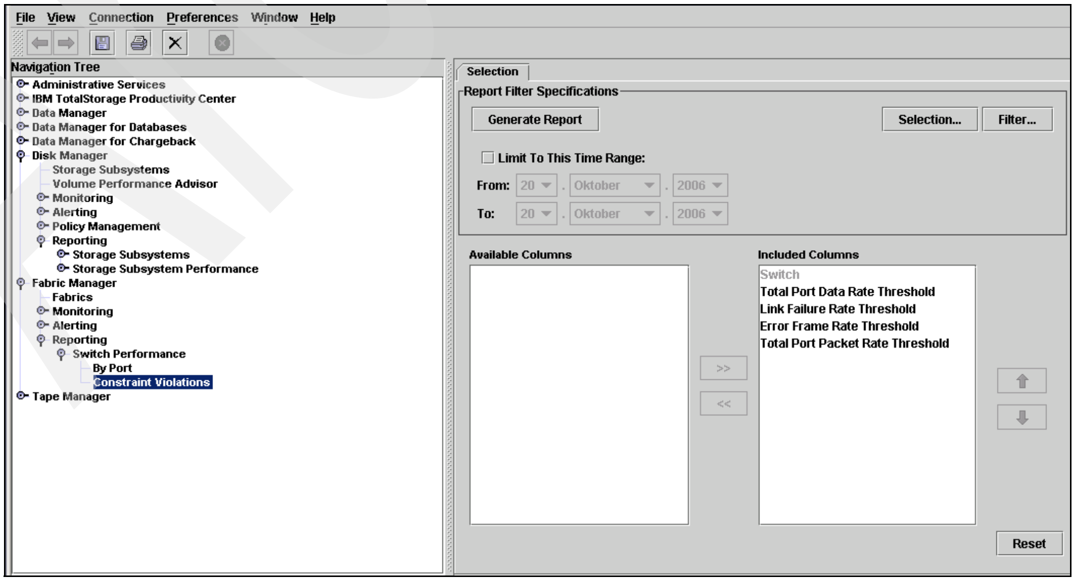

Navigation Tree

- Administrative Services
- IBM TotalStorage Productivity Center
- Data Manager
- Data Manager for Databases
- Data Manager for Chargeback
- Disk Manager
  - Storage Subsystems
  - Volume Performance Advisor
  - Monitoring
  - Alerting
  - Policy Management
  - Reporting
    - Storage Subsystems
    - Storage Subsystem Performance
- Fabric Manager
  - Fabrics
  - Monitoring
  - Alerting
  - Reporting
    - Switch Performance
      - By Port
      - Constraint Violations
- Tape Manager

Report Filter Specifications

□ Limit To This Time Range:

From: 20 . Oktober . 2006

To: 20 . Oktober . 2006

Available Columns

Included Columns

Switch
Total Port Data Rate Threshold
Link Failure Rate Threshold
Error Frame Rate Threshold
Total Port Packet Rate Threshold

*Figure 4-102 Constraint Violations report*

Click **Generate Report** to get the Constraint Violation overview. Each row represents one SAN switch (Figure 4-103).

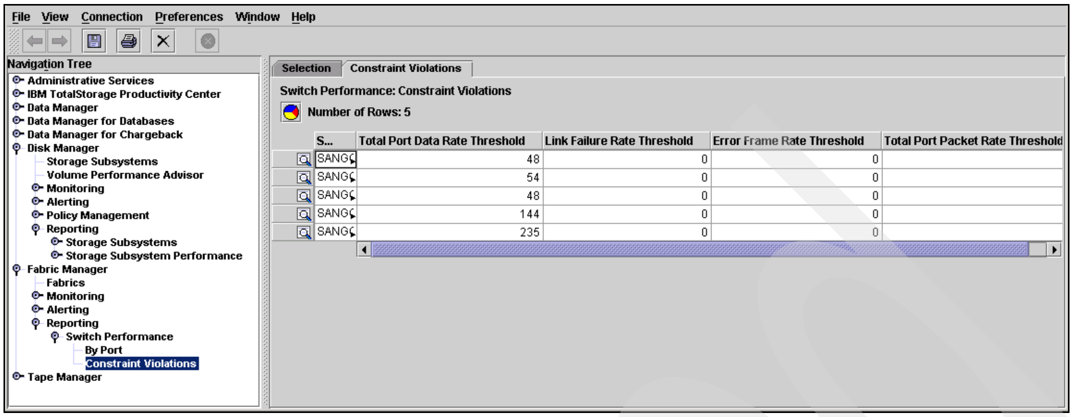

Switch Performance: Constraint Violations

Number of Rows: 5

| S... | Total Port Data Rate Threshold | Link Failure Rate Threshold | Error Frame Rate Threshold | Total Port Packet Rate Threshold |
|---|---|---|---|---|
| SANG( | 48 | 0 | 0 | |
| SANG( | 54 | 0 | 0 | |
| SANG( | 48 | 0 | 0 | |
| SANG( | 144 | 0 | 0 | |
| SANG( | 235 | 0 | 0 | |

*Figure 4-103 Report: Constraint Violation summary*

Click to get a chart of one or more SAN switches. The chart represents the disposition of all four types of constraint violation for one switch. In Figure 4-104, the constraint Total Port Data Rate Threshold was violated most frequently.

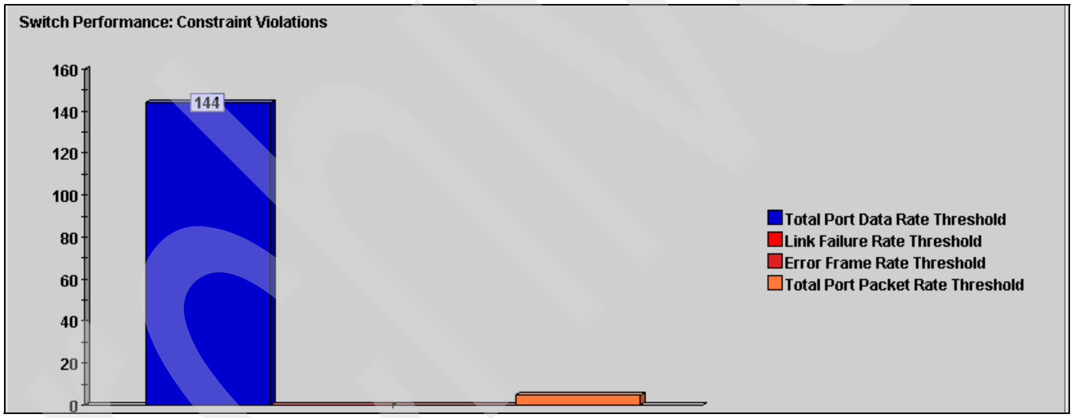

*Figure 4-104 Constraint Violation chart for a switch*

Click the magnifying glass icon to get a complete table of all constraint violations for the selected row (a specific SAN switch) as shown in Figure 4-105.

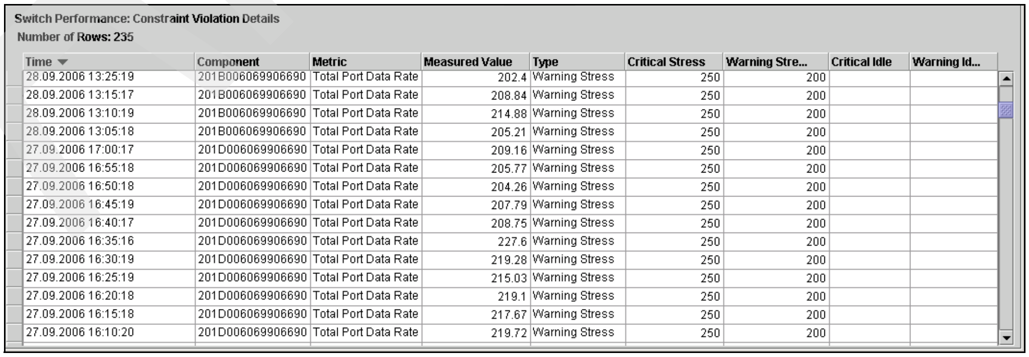

Switch Performance: Constraint Violation Details

Number of Rows: 235

| Time ▼ | Component | Metric | Measured Value | Type | Critical Stress | Warning Stre... | Critical Idle | Warning Id... |
|---|---|---|---|---|---|---|---|---|
| 28.09.2006 13:25:19 | 201B006069906690 | Total Port Data Rate | 202.4 | Warning Stress | 250 | 200 | | |
| 28.09.2006 13:15:17 | 201B006069906690 | Total Port Data Rate | 208.84 | Warning Stress | 250 | 200 | | |
| 28.09.2006 13:10:19 | 201B006069906690 | Total Port Data Rate | 214.88 | Warning Stress | 250 | 200 | | |
| 28.09.2006 13:05:18 | 201B006069906690 | Total Port Data Rate | 205.21 | Warning Stress | 250 | 200 | | |
| 27.09.2006 17:00:17 | 201D006069906690 | Total Port Data Rate | 209.16 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:55:18 | 201D006069906690 | Total Port Data Rate | 205.77 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:50:18 | 201D006069906690 | Total Port Data Rate | 204.26 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:45:19 | 201D006069906690 | Total Port Data Rate | 207.79 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:40:17 | 201D006069906690 | Total Port Data Rate | 208.75 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:35:16 | 201D006069906690 | Total Port Data Rate | 227.6 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:30:19 | 201D006069906690 | Total Port Data Rate | 219.28 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:25:19 | 201D006069906690 | Total Port Data Rate | 215.03 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:20:18 | 201D006069906690 | Total Port Data Rate | 219.1 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:15:18 | 201D006069906690 | Total Port Data Rate | 217.67 | Warning Stress | 250 | 200 | | |
| 27.09.2006 16:10:20 | 201D006069906690 | Total Port Data Rate | 219.72 | Warning Stress | 250 | 200 | | |

*Figure 4-105 List of all constraint violations for a specific SAN switch*

The Time column shows the time of the sample. The Component column identifies the switchport. With accurate performance thresholds (constraints), you can export and publish this report. It shows in a five minute interval, which depends on the sample interval, if the SAN switch has a constraint violation (a performance alert) over the time period.

5

# Security considerations

Firewalls, passwords, and user IDs follow IT security policies. This chapter documents security considerations and how to work with them.

## 5.1 Data and access considerations

IBM Total Productivity Center (TPC) stores several types of information about storage subsystems, SAN switches, computers, clusters, filesystems, and more. TPC stores only structural and informational data and never touches client data. The most sensitive data that is stored is:

- IP addresses
- Host names
- Usernames and passwords of your CIM Agent and out-of-band agents
- File Names
- Directory names
- User account names

The access to TPC is restricted with role-based user management and uses the multi-client management capability from the operating system. You can assign an operation system group to each role. Table 5-1 lists the TPC roles.

*Table 5-1 Roles within IBM TPC*

| Role | Authorization level |
|---|---|
| Superuser | This role has full access to all TPC functions. |
| Productivity center administrator | The role has full access to operations in the administration section of the GUI. |
| Disk administrator | This role has full access to TPC disk functions. |
| Disk operator | This role has access to reports only for TPC disk functions, which includes reports on tape devices. |
| Fabric administrator | This role has full access to IBM TotalStorage Productivity Center for Fabric (TPC for Fabric) functions. |
| Fabric operator | This role has access to reports only for TPC for Fabric functions. |
| Data administrator | This role has full access to IBM TotalStorage Productivity Center for Data (TPC for Data) functions. |
| Data operator | The data operator has access to reports only for data functions. |

## 5.2 Common data agent

When using TPC for Data, you must install the common data agent on each server client. By default, the agents can run scripts from the TPC server. If you do not need this feature or it is not permitted in your enterprise, you can deactivate it during the installation process.

During the installation, select **Data agent options** (Figure 5-1 on page 215) and remove the check to the left of **Agent may run scripts sent by server** (Figure 5-2 on page 215).

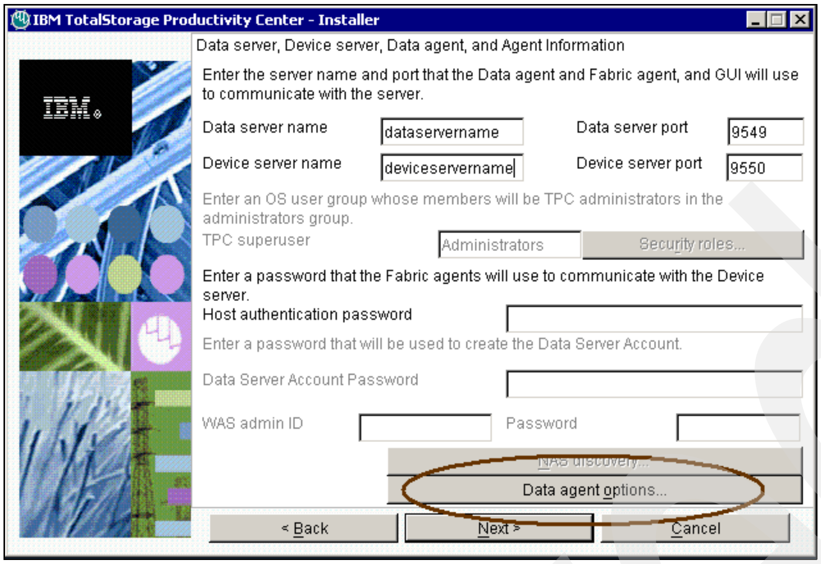

IBM TotalStorage Productivity Center - Installer

Data server, Device server, Data agent, and Agent Information

Enter the server name and port that the Data agent and Fabric agent, and GUI will use to communicate with the server.

Data server name: dataservername

Data server port: 9549

Device server name: deviceservername

Device server port: 9550

Enter an OS user group whose members will be TPC administrators in the administrators group.

TPC superuser: Administrators

Enter a password that the Fabric agents will use to communicate with the Device server.

Host authentication password:

Enter a password that will be used to create the Data Server Account.

Data Server Account Password:

WAS admin ID:

Password:

*Figure 5-1 Installation of the Common Agent*

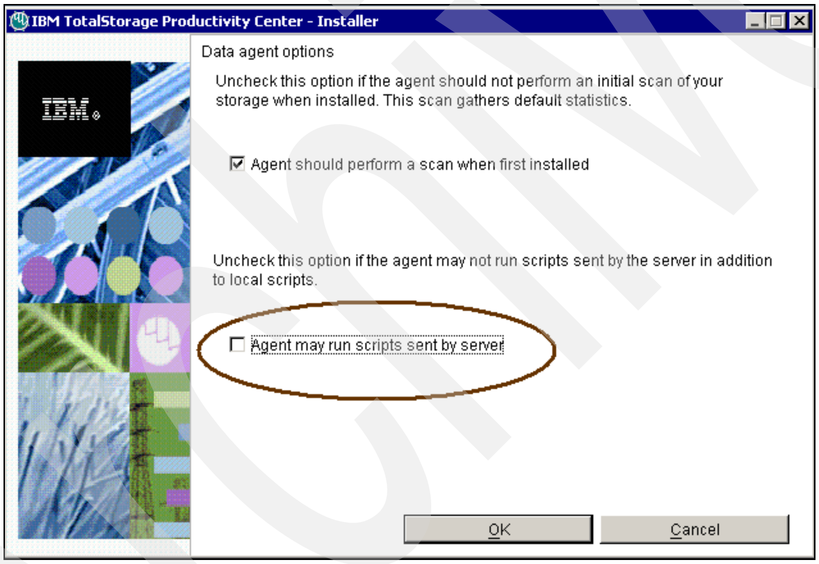

IBM TotalStorage Productivity Center - Installer

Data agent options

Uncheck this option if the agent should not perform an initial scan of your storage when installed. This scan gathers default statistics.

☑ Agent should perform a scan when first installed

Uncheck this option if the agent may not run scripts sent by the server in addition to local scripts.

☐ Agent may run scripts sent by server

*Figure 5-2 Disable agent to run script sent by TPC server*

## 5.2.1 Disable running scripts

You also can disable this setting after the installation by editing the Common Agent configuration file on your computer.

1. Open the file:

   ...\IBM\TPC\config\InstallVariable.properties

2. Look for the entry `varDataAgtScripts=true`.
3. Change this entry to `varDataAgtScripts=false` and restart the TPC server process.

## 5.3 Default security certificates

During TPC installation, we strongly recommend that you generate your own key pairs instead of using the default keys from the installation CD (Figure 5-3).

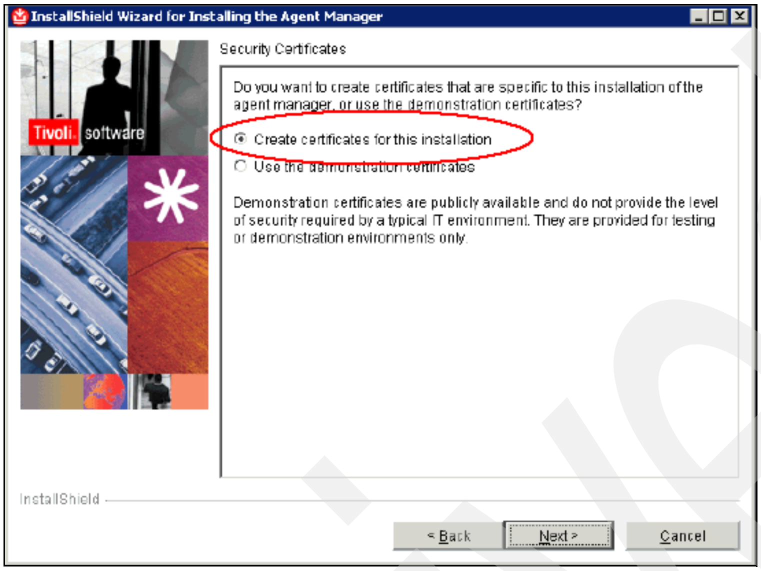

InstallShield Wizard for Installing the Agent Manager

Security Certificates

Do you want to create certificates that are specific to this installation of the agent manager, or use the demonstration certificates?

● Create certificates for this installation

○ Use the demonstration certificates

Demonstration certificates are publicly available and do not provide the level of security required by a typical IT environment. They are provided for testing or demonstration environments only.

*Figure 5-3 Creating security certificates*

Using the default keys bypasses security. Changing the certificates after the installation takes time and is a more complex task. Therefore, we recommend that you never use the default certificates.

## 5.4 How to change passwords

We reference various documents that contain instructions for you to change passwords for certain TPC components. Be careful about changing passwords because changing passwords might require that you change passwords on other components. Therefore, it is always a good practice to back up before you change a password.

### 5.4.1 DB2 password

The following documents provide information about how to change the TPC DB2 repository password.

- IC49190 *Need procedures for changes to IBM TotalStorage Productivity Center: DB2 password changes; uninstallation of Data agents; uninstallation of Tivoli Common Agent*

  You can locate this document at:

  `http://www-1.ibm.com/support/docview.wss?rs=1133&context=SS8JB5&context=SSWQP2&dc=DB500&q1=password&uid=swg21236490&loc=en_US&cs=utf-8&lang=en`

- *IBM TotalStorage Productivity Center for Data V3 - Best Practices*:
  Chapter: Changing the DB2 password for the Data Server
  Chapter: Changing the DB2 password for the Device Server
  Chapter: Changing the host authentication password for the Device Server

You can locate this document at:

http://www-1.ibm.com/support/docview.wss?rs=597&context=STCRLM4&context=SSMMUP&dc=DA4A10&uid=ssg1S7001491&loc=en_US&cs=utf-8&lang=en

## 5.4.2 TCP/IP ports

Table 5-2 lists the TCP/IP ports that are used by the Agent Manager component. You can change any of the port numbers except for port 80, which is used by the agent recovery service.

If there is a firewall between the Agent Manager and the agents and resource managers in your deployment, you must open the Agent Manager ports for inbound TCP/IP traffic to the Agent Manager.

*Table 5-2 TCP/IP ports used by the Agent Manager*

| Port | Usage | Connection security |
|---|---|---|
| 9511 | Register agents.<br>Register resource manager. | Secure SSL. |
| 9512 | Provide configuration updates.<br>Renew and revoke certificates.<br>Query the registry for agent information.<br>Requesting ID resets. | Secure SSL with client authentication. |
| 9513 | Requesting updates to the certificate revocation list.<br>Requesting Agent Manager information.<br>Downloading the trust store file.<br>This is the alternate port for the agent recovery service. | This is a public port and is insecure. |
| 80 | Recovery service (optional). | Insecure. |

Table 5-3 on page 218 lists the ports for the TPC components.

Table 5-3   TCP/IP ports used by TPC

| Component | Session initiator (server perspective) | Inbound/ outbound (server perspective) | Firewall port | Inbound/ outbound (agent perspective) | Session initiator (agent perspective) |
|---|---|---|---|---|---|
| Data server | N/A | Both | 9549 | N/A | N/A |
| Device server | N/A | Both | 9550 | N/A | N/A |
| Common Agent | Yes | Outbound | 9510 | Inbound | No |
| Agent Manager | No | Inbound | 9511 | Outbound | Yes |
| Agent Manager | Yes | Both | 9512 | Both | Yes |
| Agent Manager | No | Inbound | 9513 | Outbound | Yes |
| Common Agent (no access needed) | N/A | N/A | 9514 | Local to server | N/A |
| Common Agent (no access needed) | N/A | N/A | 9515 | Local to server | N/A |
| Agent Manager Recovery Service | No | Inbound | 80 | Outbound | Yes |
| PUSH UNIX | Yes | Outbound | SSH (22) | Both | No |
| PUSH WINDOWS | Yes | Outbound | NetBIOS sessions service 139 | N/A | N/A |
| PUSH UNIX | Yes | Outbound | RSH (514) | Both | No |
| PUSH UNIX | Yes | Outbound | REXEC (512) | Both | No |
| PUSH UNIX | Yes | Inbound | 601 | N/A | N/A |
| PUSH UNIX | Yes | Inbound | High ports 3000+ | Both | No |
| PUSH ALL | Yes | Inbound | TPCD Server 2078 | N/A | N/A |
| SLP | N/A | Both | 427 | N/A | N/A |
| SNMP Listener Port | N/A | Inbound | 162 | N/A | N/A |

## 5.4.3 Agent Manager

The *Agent Manager* is a network service that provides authentication and authorization using X.509 certificates and the Secure Sockets Layer (SSL) protocol. It also processes queries about its registry of configuration information about the agents and management applications (which are also called resource managers). A *resource manager* is the server component of a management application product that manages the agents. Examples of resource managers are IBM TPC for Fabric (server component) and IBM TPC for Data (server component).

Resource managers and agents must register with the Agent Manager before they can use its services to communicate with each other. Registration is password-protected, with separate passwords for the registration of agents and resource managers.

To control the access from the Resource Manager to the Common Agent, you use certificates to make sure that only an authorized Resource Manager can install and run code on a computer system. This certificate is stored in the agentTrust.jks and locked with the agent registration password.

You use certificates to validate if a requester is allowed to establish a communication.

You can use the ikeyman utility in the java\jre subdirectory to verify your password. Figure 5-4 shows you an overview of the Agent Manager.

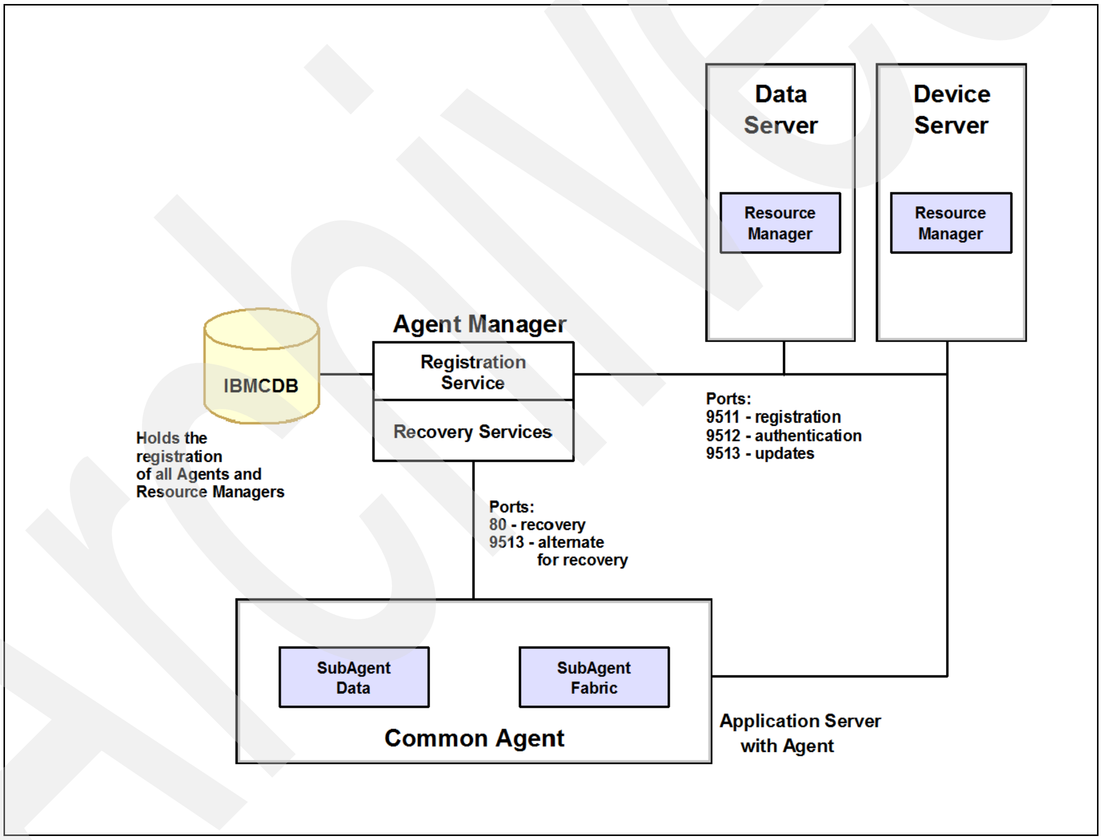

*Figure 5-4 Diagram of Agent Manager and its services*

The Agent Manager also provides an agent recovery service, which is a network service for error logging for agents that cannot communicate with other Agent Manager services. Agents use an unsecured HTTP connection to communicate with the agent recovery service. Because the connection is unsecured, an agent can always communicate with the agent recovery service, even if the agent is incorrectly configured or has expired or revoked certificates. The agent recovery service is a WebSphere® servlet container that runs on the Agent Manager server.

Agents locate the agent recovery service using the unqualified host name TivoliAgentRecovery. Your Domain Name System (DNS) server must map the host name TivoliAgentRecovery to the computer system where you installed the Agent Manager. The normal DNS lookup mechanism iterates through the domain search list for the agent, appends each domain in the list to the unqualified host name, and then performs a DNS lookup to attempt to resolve the name.

The agent recovery service listens for recovery requests on two ports: port 80 and a configurable port (by default, 9513). Using port 80 makes the request more likely to pass through a firewall between the agent and the agent recovery service. However, if the Agent Manager is on the same system as the HTTP server, port 80 is unavailable. The configurable second port provides an alternate communication port, in case you need to disable the use of port 80 by the agent recovery service.

# 6

# IBM Total Productivity Center tools

This chapter is a guide that is designed to help IBM TotalStorage Productivity Center (TPC) administrators and Storage Area Network (SAN) administrators debug, manage, and troubleshoot their TPC environments with the tools that are provided by TPC. We discuss these tools:

- Service tools to help you gather required information and logs necessary for analysis
- HealthCheck tools to verify your installation
- Tools to help you manage your environment by using a command line interface (CLI)

Using these tools requires that storage administrators understand their environment and have basic knowledge of TPC and SAN functionality.

This chapter also provides you with recommendations and best practices for managing your TPC environment.

# 6.1 Service tool

Serviceability is also known as supportability. It refers to the ability of technical support personnel to debug or perform root cause analysis in pursuit of solving a problem with a product.

If you have problems running TPC, the service tool helps collect information for all installed TPC components (for example, Data server, Device server, agents, and GUI) from a centralized location. This tool detects the configuration of the system on which it runs and collects the appropriate information. The tools places the information in a .zip file.

This tool installs when you install the TPC server. IBM support personnel will ask you to run the service tool when they need you to send them a PMR.

You do not need to specify what type of logs to collect, because TPC collects all of the logs that are required for analysis.

> **Note:** The service tool does not collection information from the repository database, see **`repocopy`**.

## 6.1.1 Location of the service tool

The service tool is located in the following directories after the code is installed:

- For Windows:

  **C:\Program Files\IBM\TPC\service\**

- For UNIX or Linux:

  **/opt/IBM/TPC/service/**

## 6.1.2 Output of the service tool

This tool collects the following information:

- Host name of the TPC server
- IP address of the TPC server
- Operating system name, version, and architecture
- Java™ home, version, and classpath
- JVM™ implementation name and version
- Full IP configuration information
- Protocol statistics and current TCP/IP network connections (including listening ports)
- Diagnostic information regarding the system and its services
- Listing of all library files (for example, server and library, agent and library, GUI library)
- Installation logs (these logs include the agent deployment log and the endpoint log status, agent logs, TPC server logs, all probes, all discoveries, and so forth)
- Required logs including the subdirectories (these logs are the system logs, trace logs, audit logs, and message logs)
- The required configuration files (the agent configuration files, data, and device configuration files)

- TPC version and build information
- Output from the **ipconfig /all** command (for Windows)
- Output from the **ipconfig -a** command (for UNIX or Linux)

### 6.1.3 How to run the service tool

To run the service tool for the Data server and Device server (to collect all information about the TPC environment), follow these steps:

1. Log on to the system. You must have administrator authority on Windows or root authority on UNIX or Linux.
2. If you used the default directory, go to the following directory (see Figure 6-1):
   – For Windows:
     C:\Program Files\IBM\TPC\service\
   – For UNIX or Linux:
     /opt/IBM/TPC/service/

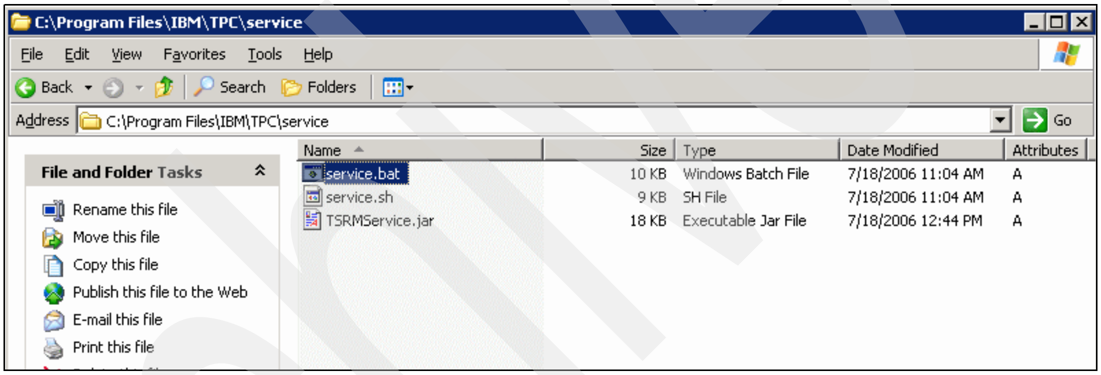

| Name | Size | Type | Date Modified | Attributes |
|---|---|---|---|---|
| service.bat | 10 KB | Windows Batch File | 7/18/2006 11:04 AM | A |
| service.sh | 9 KB | SH File | 7/18/2006 11:04 AM | A |
| TSRMService.jar | 18 KB | Executable Jar File | 7/18/2006 12:44 PM | A |

*Figure 6-1 Service Tool default directory*

3. Run the following program:
   – **service.bat** (for Windows)
   – **service.sh** (for UNIX or Linux)

   In Figure 6-2 on page 224, notice that a command window runs in the foreground when you execute the service tool. The command window verifies the information collected by the tool.

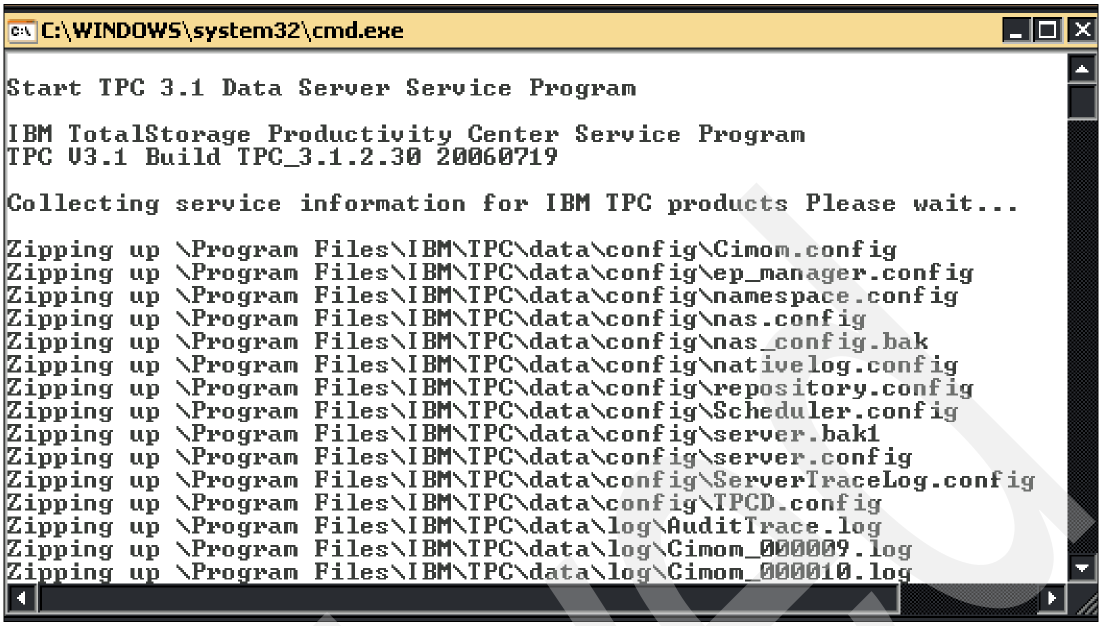

```
Start TPC 3.1 Data Server Service Program

IBM TotalStorage Productivity Center Service Program
TPC V3.1 Build TPC_3.1.2.30 20060719

Collecting service information for IBM TPC products Please wait...

Zipping up \Program Files\IBM\TPC\data\config\Cimom.config
Zipping up \Program Files\IBM\TPC\data\config\ep_manager.config
Zipping up \Program Files\IBM\TPC\data\config\namespace.config
Zipping up \Program Files\IBM\TPC\data\config\nas.config
Zipping up \Program Files\IBM\TPC\data\config\nas_config.bak
Zipping up \Program Files\IBM\TPC\data\config\nativelog.config
Zipping up \Program Files\IBM\TPC\data\config\repository.config
Zipping up \Program Files\IBM\TPC\data\config\Scheduler.config
Zipping up \Program Files\IBM\TPC\data\config\server.bak1
Zipping up \Program Files\IBM\TPC\data\config\server.config
Zipping up \Program Files\IBM\TPC\data\config\ServerTraceLog.config
Zipping up \Program Files\IBM\TPC\data\config\TPCD.config
Zipping up \Program Files\IBM\TPC\data\log\AuditTrace.log
Zipping up \Program Files\IBM\TPC\data\log\Cimom_000009.log
Zipping up \Program Files\IBM\TPC\data\log\Cimom_000010.log
```

*Figure 6-2 Run the service.bat command*

4. One or more .zip files are created in the directory where you ran the service tool. If you have a Data agent or Fabric agent installed on the same computer as the Data server or Device server, that information is collected when you run the tool for the Data server or Device server. The files should be sent to the IBM support center.

   The following .zip files are created:

   – TPCDATAservice.zip (for the Data server)
   – TPCDEVservice.zip (for the Device server)
   – TPCDATACAservice.zip (for the Data agent)
   – TPCDEVCAservice.zip (for the Fabric agent)

### 6.1.4 Data agent and Fabric agent service tool

Running the service tool only for the Data agent and Fabric agent only collects information about the agents that you have deployed in your environment. Follow these steps:

1. Log on to the system. You must have administrator authority on Windows or root authority on UNIX or Linux.
2. If you used the default directory, go to the following directory (see Figure 6-3 on page 225):
   – Windows:

     **C:\Program Files\IBM\TPC\ca\subagents\TPC\service**

   – UNIX or Linux:

     **/opt/IBM/TPC/ca/subagents/TPC/service**

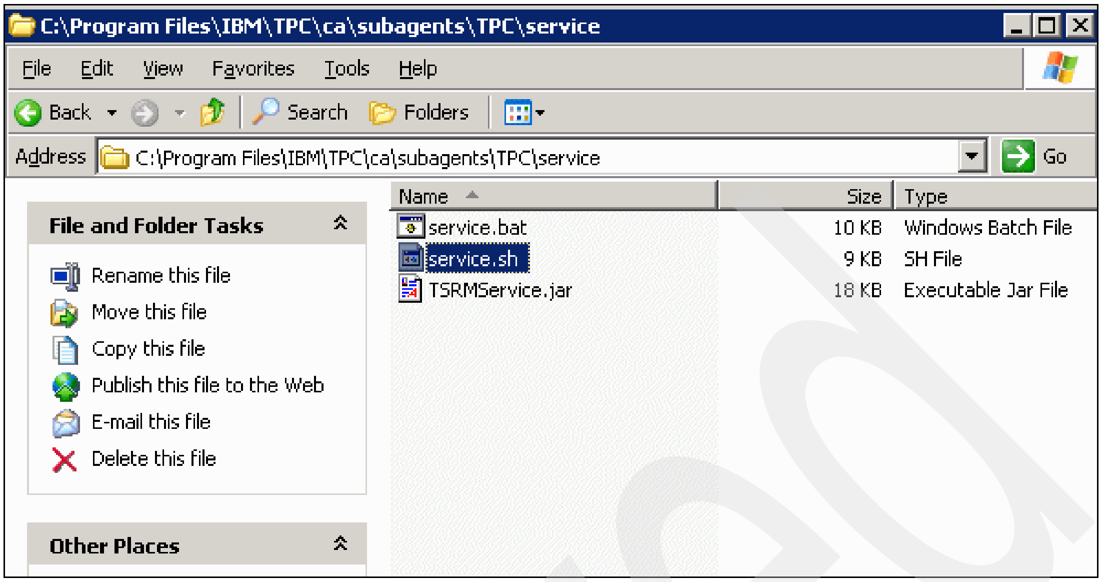

| Name | Size | Type |
|---|---|---|
| service.bat | 10 KB | Windows Batch File |
| service.sh | 9 KB | SH File |
| TSRMService.jar | 18 KB | Executable Jar File |

*Figure 6-3 Data agent and Fabric agent service tool default directory*

> **Note:** The service tool collects information in a different place for deployed agents compared to the whole TPC environment.

3. Run the following program:

   – `service.bat` (for Windows)

   – `service.sh` (for UNIX or Linux)

   Figure 6-4 shows the command prompt for the Windows environment.

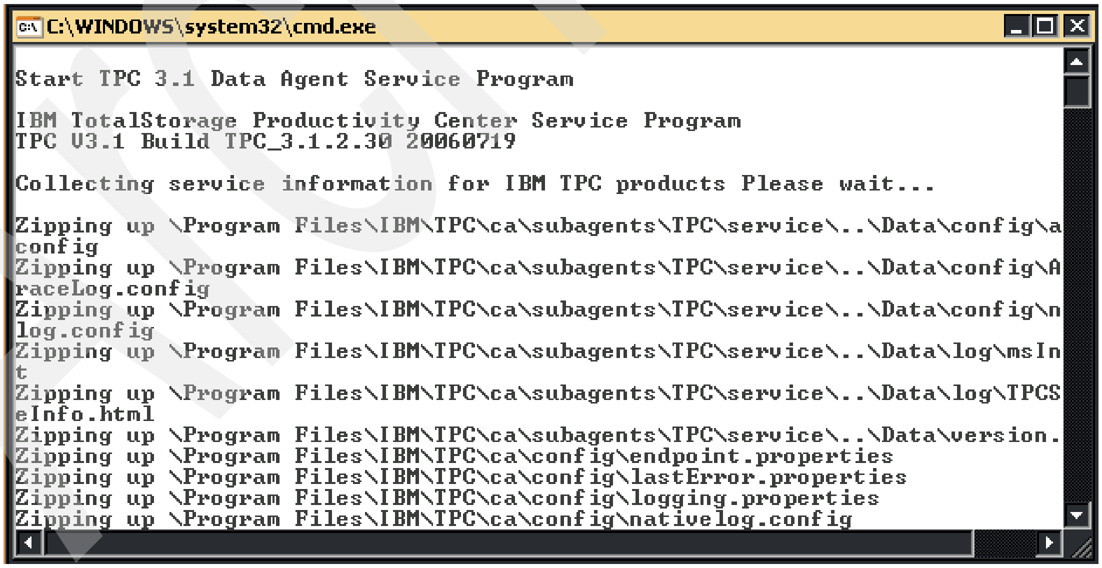

```
Start TPC 3.1 Data Agent Service Program

IBM TotalStorage Productivity Center Service Program
TPC V3.1 Build TPC_3.1.2.30 20060719

Collecting service information for IBM TPC products Please wait...

Zipping up \Program Files\IBM\TPC\ca\subagents\TPC\service\..\Data\config\a
config
Zipping up \Program Files\IBM\TPC\ca\subagents\TPC\service\..\Data\config\A
raceLog.config
Zipping up \Program Files\IBM\TPC\ca\subagents\TPC\service\..\Data\config\n
log.config
Zipping up \Program Files\IBM\TPC\ca\subagents\TPC\service\..\Data\log\msIn
t
Zipping up \Program Files\IBM\TPC\ca\subagents\TPC\service\..\Data\log\TPCS
eInfo.html
Zipping up \Program Files\IBM\TPC\ca\subagents\TPC\service\..\Data\version.
Zipping up \Program Files\IBM\TPC\ca\config\endpoint.properties
Zipping up \Program Files\IBM\TPC\ca\config\lastError.properties
Zipping up \Program Files\IBM\TPC\ca\config\logging.properties
Zipping up \Program Files\IBM\TPC\ca\config\nativelog.config
```

*Figure 6-4 Run the service.bat command for Data agents and Fabric agents*

4. A .zip file is created in the directory where you ran the service tool. Send the files to the IBM support center.

The following .zip files are created:

- TPCDATACAservice.zip (for the Data agent)
- TPCDEVCAservice.zip (for the Fabric agent)

## 6.2 Agent Manager tools

Many Tivoli management solutions require you to deploy client or agent code in the IT infrastructure across multiple user machines or application servers. The deployed agent code collects data from and performs operations on managed resources on behalf of a TPC instance. The machines on which this code is deployed are referred to as *agents*. Common Agent Services (CAS), comprised of Common Agent and Agent Manager, provides a manageable and extensible infrastructure for TPC agents.

The human effort required to install, configure, maintain, and monitor the various agent implementations becomes a management task when the environment consists of hundreds of machines.

*Common Agent Services* (CAS) was created in response to client concerns about having to install and care for multiple TPC agents on every machine. CAS reduces the effort required for developing and delivering management solutions that involve deployed agent technology. IBM has developed tools to ensure that your Agent Manager server is installed successfully, and if problems arise, these tools assist with debugging these problems.

The Agent Manager provides a toolkit of “as-is” administration tools. The toolkit is located in the AM_installdir/toolkit directory.

### 6.2.1 HealthCheck

*HealthCheck* is a command to verify the state of the basic functions of the Agent Manager. This command indicates if the Agent Manager is operational or not. You must run this command from the machine on which the Agent Manager is installed. HealthCheck contacts the Agent Manager through Web services and queries the Agent Manager property settings.

The HealthCheck readme file has instructions about running the commands and the parameters that are required for these commands. It also lists known issues and limitations.

You might run the `HealthCheck` command when you are having problems deploying agents or agents cannot register with Agent Manager.

HealthCheck is installed as part of the Agent Manager toolkit. The files are:

- For Windows:

  toolkit\bin\**HealthCheck.bat**

- For Linux and UNIX:

  toolkit/bin/**HealthCheck.sh**

#### How to run Healthcheck

On Windows:

1. Start a command window (CMD).
2. Make sure that your PATH environment variable points to the directory containing the Java command.

3. From the AgentManager\Toolkit\bin directory, type **`HealthCheck`**

   See the readme file for more options.

For UNIX and Linux:

1. Start a terminal session.
2. Make sure that your PATH environment variable points to the directory containing the Java command.
3. From the AgentManager/Toolkit/bin directory, type **`HealthCheck.sh`**

## Command line options

The following options are required only if the Agent Manager is configured for other than the installation defaults:

- **-RegistrationPW** <*Agent registration password*>

  Specify the agent registration password. This option is required if you are not using the default password. The default password is `changeMe` and is case sensitive.

- **Port** <*Registration port*>

  The agent registration port. By default, this is port `9511`.

- **-TrustStoreName** <*Name of trust store*>

  The name of the trust store from which to import trusted certificates.

- **-TrustStorePW** <*Password for trust store*>

  The password for the imported trust store.

- **-RegistrationURL** <*Registration URL*>

  The Agent Manager context root. The value must end with a forward slash (/). By default, the RegistrationURL is /AgentMgr/Registration.

An example of the **`HealthCheck`** command is to verify that the Common Agent password that you use is still valid and to verify the health of your Agent Manager server (see Figure 6-5 on page 228). The default agent registration password is `changeMe`, which is created when you install the Agent Manager code.

The syntax of the **`HealthCheck`** command is:

- For Windows:

```
healthcheck -registrationPw changeMe
```

- For UNIX:

```
/healthcheck.sh -registrationPw changeMe
```

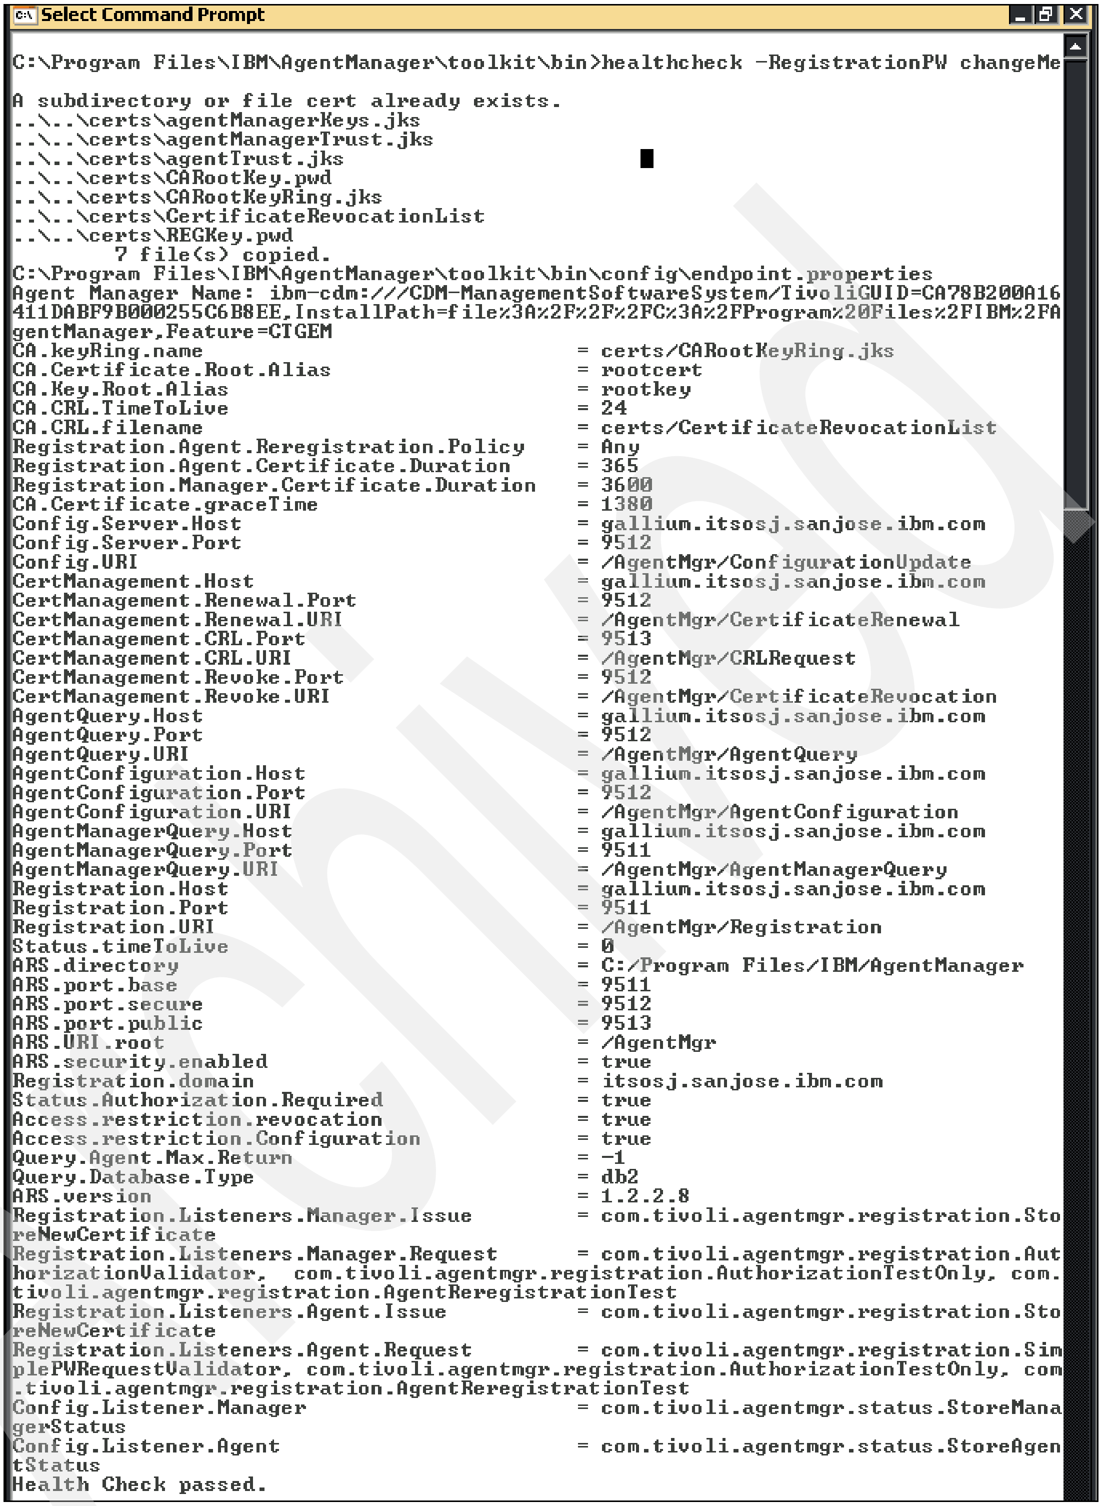

```
C:\Program Files\IBM\AgentManager\toolkit\bin>healthcheck -RegistrationPW changeMe

A subdirectory or file cert already exists.
..\..\certs\agentManagerKeys.jks
..\..\certs\agentManagerTrust.jks
..\..\certs\agentTrust.jks
..\..\certs\CARootKey.pwd
..\..\certs\CARootKeyRing.jks
..\..\certs\CertificateRevocationList
..\..\certs\REGKey.pwd
        7 file(s) copied.
C:\Program Files\IBM\AgentManager\toolkit\bin\config\endpoint.properties
Agent Manager Name: ibm-cdm:///CDM-ManagementSoftwareSystem/TivoliGUID=CA78B200A16
411DABF9B000255C6B8EE,InstallPath=file%3A%2F%2F%2FC%3A%2FProgram%20Files%2FIBM%2FA
gentManager,Feature=CTGEM
CA.keyRing.name                             = certs/CARootKeyRing.jks
CA.Certificate.Root.Alias                   = rootcert
CA.Key.Root.Alias                           = rootkey
CA.CRL.TimeToLive                           = 24
CA.CRL.filename                             = certs/CertificateRevocationList
Registration.Agent.Reregistration.Policy    = Any
Registration.Agent.Certificate.Duration     = 365
Registration.Manager.Certificate.Duration   = 3600
CA.Certificate.graceTime                    = 1380
Config.Server.Host                          = gallium.itsosj.sanjose.ibm.com
Config.Server.Port                          = 9512
Config.URI                                  = /AgentMgr/ConfigurationUpdate
CertManagement.Host                         = gallium.itsosj.sanjose.ibm.com
CertManagement.Renewal.Port                 = 9512
CertManagement.Renewal.URI                  = /AgentMgr/CertificateRenewal
CertManagement.CRL.Port                     = 9513
CertManagement.CRL.URI                      = /AgentMgr/CRLRequest
CertManagement.Revoke.Port                  = 9512
CertManagement.Revoke.URI                   = /AgentMgr/CertificateRevocation
AgentQuery.Host                             = gallium.itsosj.sanjose.ibm.com
AgentQuery.Port                             = 9512
AgentQuery.URI                              = /AgentMgr/AgentQuery
AgentConfiguration.Host                     = gallium.itsosj.sanjose.ibm.com
AgentConfiguration.Port                     = 9512
AgentConfiguration.URI                      = /AgentMgr/AgentConfiguration
AgentManagerQuery.Host                      = gallium.itsosj.sanjose.ibm.com
AgentManagerQuery.Port                      = 9511
AgentManagerQuery.URI                       = /AgentMgr/AgentManagerQuery
Registration.Host                           = gallium.itsosj.sanjose.ibm.com
Registration.Port                           = 9511
Registration.URI                            = /AgentMgr/Registration
Status.timeToLive                           = 0
ARS.directory                               = C:/Program Files/IBM/AgentManager
ARS.port.base                               = 9511
ARS.port.secure                             = 9512
ARS.port.public                             = 9513
ARS.URI.root                                = /AgentMgr
ARS.security.enabled                        = true
Registration.domain                         = itsosj.sanjose.ibm.com
Status.Authorization.Required               = true
Access.restriction.revocation               = true
Access.restriction.Configuration            = true
Query.Agent.Max.Return                      = -1
Query.Database.Type                         = db2
ARS.version                                 = 1.2.2.8
Registration.Listeners.Manager.Issue        = com.tivoli.agentmgr.registration.Sto
reNewCertificate
Registration.Listeners.Manager.Request      = com.tivoli.agentmgr.registration.Aut
horizationValidator, com.tivoli.agentmgr.registration.AuthorizationTestOnly, com.
tivoli.agentmgr.registration.AgentReregistrationTest
Registration.Listeners.Agent.Issue          = com.tivoli.agentmgr.registration.Sto
reNewCertificate
Registration.Listeners.Agent.Request        = com.tivoli.agentmgr.registration.Sim
plePWRequestValidator, com.tivoli.agentmgr.registration.AuthorizationTestOnly, com
.tivoli.agentmgr.registration.AgentReregistrationTest
Config.Listener.Manager                     = com.tivoli.agentmgr.status.StoreMana
gerStatus
Config.Listener.Agent                       = com.tivoli.agentmgr.status.StoreAgen
tStatus
Health Check passed.
```

*Figure 6-5 HealthCheck with registration password changeMe*

The current system is healthy based on the message, "`Health Check passed`." Now, let us look at a system where the agent registration password is not the default password `changeMe`, as shown in Figure 6-6 on page 229. The syntax used is:

- For Windows: **`healthcheck`**
- For UNIX: **`/healthcheck.sh`**

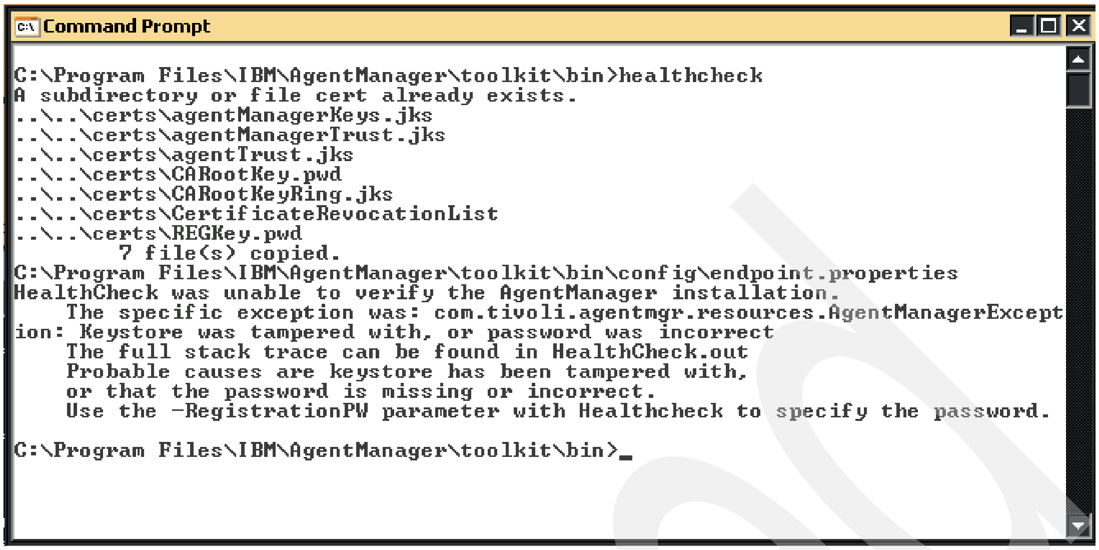

```
C:\Program Files\IBM\AgentManager\toolkit\bin>healthcheck
A subdirectory or file cert already exists.
..\..\certs\agentManagerKeys.jks
..\..\certs\agentManagerTrust.jks
..\..\certs\agentTrust.jks
..\..\certs\CARootKey.pwd
..\..\certs\CARootKeyRing.jks
..\..\certs\CertificateRevocationList
..\..\certs\REGKey.pwd
        7 file(s) copied.
C:\Program Files\IBM\AgentManager\toolkit\bin\config\endpoint.properties
HealthCheck was unable to verify the AgentManager installation.
    The specific exception was: com.tivoli.agentmgr.resources.AgentManagerExcept
ion: Keystore was tampered with, or password was incorrect
    The full stack trace can be found in HealthCheck.out
    Probable causes are keystore has been tampered with,
    or that the password is missing or incorrect.
    Use the -RegistrationPW parameter with Healthcheck to specify the password.

C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-6 HealthCheck with a registration password that is not the default*

The HealthCheck utility displays a message that your request was unsuccessful and that it cannot verify your Agent Manager installation. Health Check also recommends that, in this scenario, you use the **-RegistrationPW**. Now, use the correct registration password (**-RegistrationPW**) to verify the health of the Agent Manager system as shown in Figure 6-7 on page 230.

The syntax that we used is:

- For Windows: **`healthcheck -RegistrationPw tpcadmin`**
- For UNIX: **`/healthcheck.sh -RegistrationPw tpcadmin`**

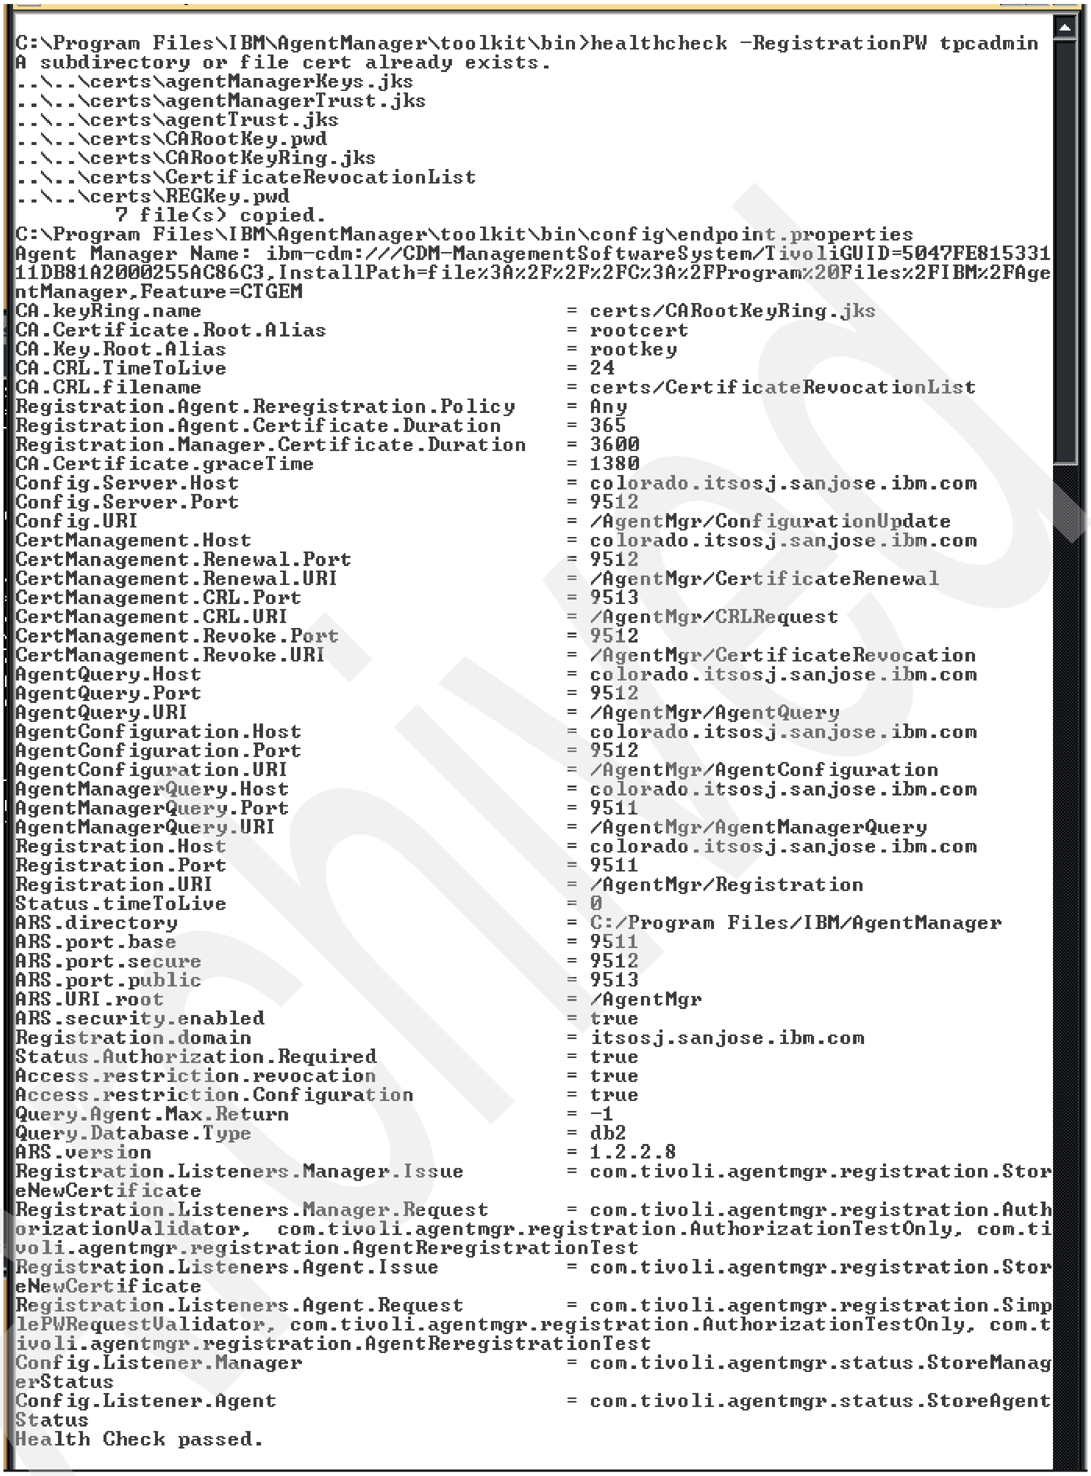

```
C:\Program Files\IBM\AgentManager\toolkit\bin>healthcheck -RegistrationPW tpcadmin
A subdirectory or file cert already exists.
..\..\certs\agentManagerKeys.jks
..\..\certs\agentManagerTrust.jks
..\..\certs\agentTrust.jks
..\..\certs\CARootKey.pwd
..\..\certs\CARootKeyRing.jks
..\..\certs\CertificateRevocationList
..\..\certs\REGKey.pwd
        7 file(s) copied.
C:\Program Files\IBM\AgentManager\toolkit\bin\config\endpoint.properties
Agent Manager Name: ibm-cdm:///CDM-ManagementSoftwareSystem/TivoliGUID=5047FE815331
11DB81A2000255AC86C3,InstallPath=file%3A%2F%2F%2FC%3A%2FProgram%20Files%2FIBM%2FAge
ntManager,Feature=CTGEM
CA.keyRing.name                             = certs/CARootKeyRing.jks
CA.Certificate.Root.Alias                   = rootcert
CA.Key.Root.Alias                           = rootkey
CA.CRL.TimeToLive                           = 24
CA.CRL.filename                             = certs/CertificateRevocationList
Registration.Agent.Reregistration.Policy    = Any
Registration.Agent.Certificate.Duration     = 365
Registration.Manager.Certificate.Duration   = 3600
CA.Certificate.graceTime                    = 1380
Config.Server.Host                          = colorado.itsosj.sanjose.ibm.com
Config.Server.Port                          = 9512
Config.URI                                  = /AgentMgr/ConfigurationUpdate
CertManagement.Host                         = colorado.itsosj.sanjose.ibm.com
CertManagement.Renewal.Port                 = 9512
CertManagement.Renewal.URI                  = /AgentMgr/CertificateRenewal
CertManagement.CRL.Port                     = 9513
CertManagement.CRL.URI                      = /AgentMgr/CRLRequest
CertManagement.Revoke.Port                  = 9512
CertManagement.Revoke.URI                   = /AgentMgr/CertificateRevocation
AgentQuery.Host                             = colorado.itsosj.sanjose.ibm.com
AgentQuery.Port                             = 9512
AgentQuery.URI                              = /AgentMgr/AgentQuery
AgentConfiguration.Host                     = colorado.itsosj.sanjose.ibm.com
AgentConfiguration.Port                     = 9512
AgentConfiguration.URI                      = /AgentMgr/AgentConfiguration
AgentManagerQuery.Host                      = colorado.itsosj.sanjose.ibm.com
AgentManagerQuery.Port                      = 9511
AgentManagerQuery.URI                       = /AgentMgr/AgentManagerQuery
Registration.Host                           = colorado.itsosj.sanjose.ibm.com
Registration.Port                           = 9511
Registration.URI                            = /AgentMgr/Registration
Status.timeToLive                           = 0
ARS.directory                               = C:/Program Files/IBM/AgentManager
ARS.port.base                               = 9511
ARS.port.secure                             = 9512
ARS.port.public                             = 9513
ARS.URI.root                                = /AgentMgr
ARS.security.enabled                        = true
Registration.domain                         = itsosj.sanjose.ibm.com
Status.Authorization.Required               = true
Access.restriction.revocation               = true
Access.restriction.Configuration            = true
Query.Agent.Max.Return                      = -1
Query.Database.Type                         = db2
ARS.version                                 = 1.2.2.8
Registration.Listeners.Manager.Issue        = com.tivoli.agentmgr.registration.Stor
eNewCertificate
Registration.Listeners.Manager.Request      = com.tivoli.agentmgr.registration.Auth
orizationValidator,  com.tivoli.agentmgr.registration.AuthorizationTestOnly, com.ti
voli.agentmgr.registration.AgentReregistrationTest
Registration.Listeners.Agent.Issue          = com.tivoli.agentmgr.registration.Stor
eNewCertificate
Registration.Listeners.Agent.Request        = com.tivoli.agentmgr.registration.Simp
lePWRequestValidator, com.tivoli.agentmgr.registration.AuthorizationTestOnly, com.t
ivoli.agentmgr.registration.AgentReregistrationTest
Config.Listener.Manager                     = com.tivoli.agentmgr.status.StoreManag
erStatus
Config.Listener.Agent                       = com.tivoli.agentmgr.status.StoreAgent
Status
Health Check passed.
```

*Figure 6-7 HealthCheck with the registration password stipulated*

## 6.2.2 Verifying the installation

If the Agent Manager is operational, the **`HeathCheck`** command displays the Agent Manager configuration, followed by a “`Passed`” message as shown in Figure 6-8 on page 231.

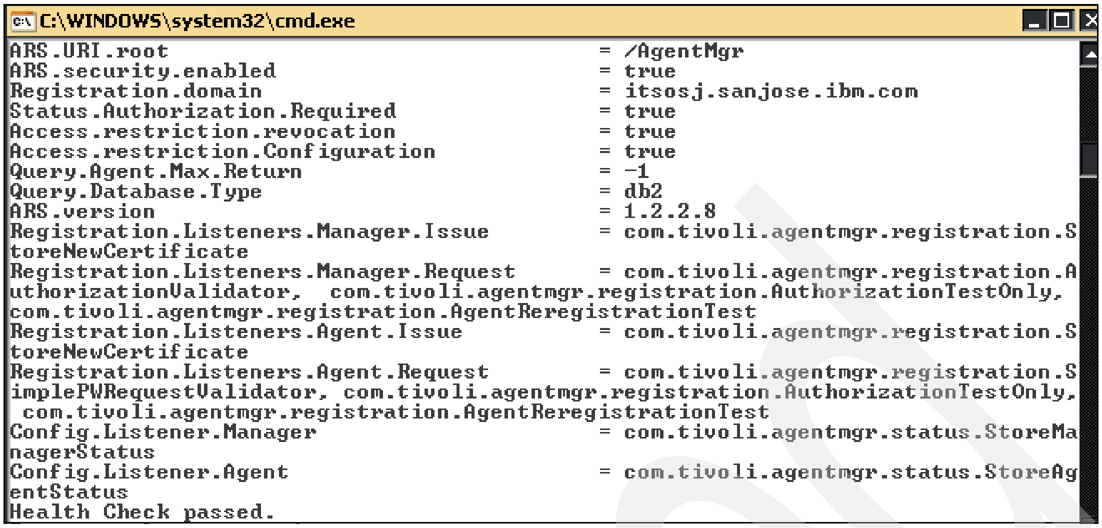

```
ARS.URI.root                                 = /AgentMgr
ARS.security.enabled                         = true
Registration.domain                          = itsosj.sanjose.ibm.com
Status.Authorization.Required                = true
Access.restriction.revocation                = true
Access.restriction.Configuration             = true
Query.Agent.Max.Return                       = -1
Query.Database.Type                          = db2
ARS.version                                  = 1.2.2.8
Registration.Listeners.Manager.Issue         = com.tivoli.agentmgr.registration.S
toreNewCertificate
Registration.Listeners.Manager.Request       = com.tivoli.agentmgr.registration.A
uthorizationValidator,  com.tivoli.agentmgr.registration.AuthorizationTestOnly,
com.tivoli.agentmgr.registration.AgentReregistrationTest
Registration.Listeners.Agent.Issue           = com.tivoli.agentmgr.registration.S
toreNewCertificate
Registration.Listeners.Agent.Request         = com.tivoli.agentmgr.registration.S
implePWRequestValidator,  com.tivoli.agentmgr.registration.AuthorizationTestOnly,
 com.tivoli.agentmgr.registration.AgentReregistrationTest
Config.Listener.Manager                      = com.tivoli.agentmgr.status.StoreMa
nagerStatus
Config.Listener.Agent                        = com.tivoli.agentmgr.status.StoreAg
entStatus
Health Check passed.
```

*Figure 6-8 Health Check passed*

If the Agent Manager is not operational, the **HealthCheck** command displays a summary of Java exceptions, suggests a probable cause and error, and gives the name of a file containing additional information as shown in Figure 6-9. Review this file. If you cannot resolve the problems, contact the IBM Support Center.

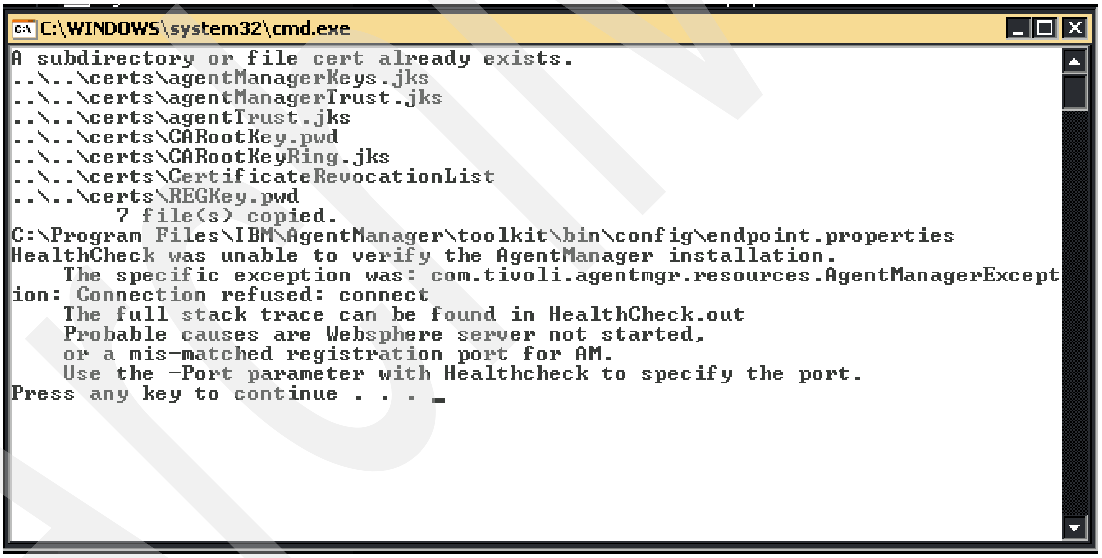

```
A subdirectory or file cert already exists.
..\..\certs\agentManagerKeys.jks
..\..\certs\agentManagerTrust.jks
..\..\certs\agentTrust.jks
..\..\certs\CARootKey.pwd
..\..\certs\CARootKeyRing.jks
..\..\certs\CertificateRevocationList
..\..\certs\REGKey.pwd
        7 file(s) copied.
C:\Program Files\IBM\AgentManager\toolkit\bin\config\endpoint.properties
HealthCheck was unable to verify the AgentManager installation.
    The specific exception was: com.tivoli.agentmgr.resources.AgentManagerExcept
ion: Connection refused: connect
    The full stack trace can be found in HealthCheck.out
    Probable causes are Websphere server not started,
    or a mis-matched registration port for AM.
    Use the -Port parameter with Healthcheck to specify the port.
Press any key to continue . . . _
```

*Figure 6-9 Example of an Agent Manager problem*

## 6.2.3 Log Collector

*Log Collector* gathers logs and other information that are needed to debug the Agent Manager. After running the HealthCheck command and there is a failure exception error, run the **LogCollector** command to gather the logs.

Log Collector creates the LogCollector.zip file, which is located in the root directory of where the Agent Manager is installed. You must run this command from the machine on which the Agent Manager is installed. The Log Collector readme file has instructions about running the commands, and it also lists known issues and limitations.

`LogCollector` is installed as part of the Agent Manager toolkit.

The files are:

- For Windows:

  toolkit/bin/**LogCollector.bat**
- For Linux and UNIX:

  toolkit/bin/**LogCollector.sh**

### Output of the LogCollector tool

The output is:

- LogCollector opens the <*Agent Manager*>/logs/amInstall.log file to look for needed install values.
- LogCollector opens a Java ZipFileStream for collecting its log files.
- System environment variables, ipConfig output “df –k” output and directory structure are written to LogCollector_notes.txt.
- Agent Manager /logs files are added to the .zip file.
- Agent Manager /logs/jacl files are added to the .zip file.
- WebSphere/LWI log files are added to the .zip file.
- WebSphere server1 log files are added to the .zip file (for Agent Manager running on WebSphere Application Server).
- If Agent Manager was added into another application server, those WebSphere log files are added to the .zip file.
- Agent Manager Database log files are added to the .zip file.
- LogCollector reads the <*Agent Manager*>/logs/amInstall.log file to determine the WebSphere/LWI location, WebSphere node, WebSphere instance, database type, and database location.

### How to run LogCollector

On Windows:

1. Start a command prompt window (CMD).
2. Make sure that your PATH environment variable points to the directory containing the Java command.
3. From the AgentManager\Toolkit\bin directory, type `LogCollector` as shown in Figure 6-10 on page 233.

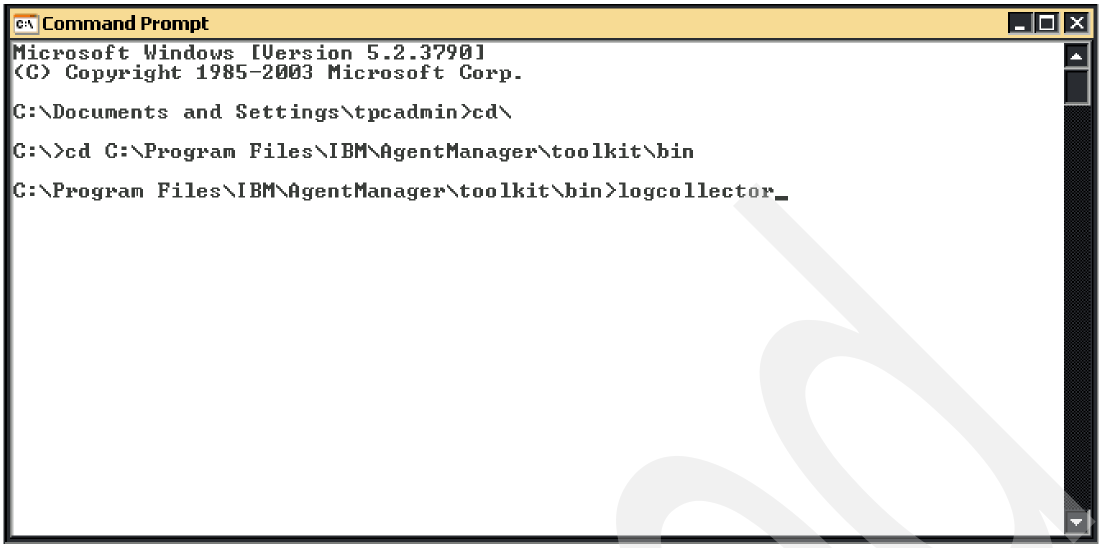

```
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

C:\Documents and Settings\tpcadmin>cd\

C:\>cd C:\Program Files\IBM\AgentManager\toolkit\bin

C:\Program Files\IBM\AgentManager\toolkit\bin>logcollector_
```

*Figure 6-10 Run LogCollector.bat*

> **Note:** Failure to retrieve all of the values from the amInstall.log causes LogCollector to exit.

On UNIX and Linux:

1. Start a terminal session.
2. Make sure that your PATH environment variable points to the directory containing the java command.
3. From the AgentManager/toolkit/bin directory, type **/LogCollector.sh**

You can send this file to the IBM support center to help you resolve problems on Agent Manager. The .zip file is located under C:\Program Files\IBM\AgentManager\ (for Windows) and /opt/IBM/AgentManager (for UNIX or Linux).

You can send this file to the IBM support center to help you resolve problems.

### 6.2.4 Deregistration tools

The Agent Manager handles the registration of resource managers and agents, such as issuing certificates and keys and performing authentication. One Agent Manager instance can manage multiple resource managers and agents. The Agent Manager can be on the same machine as the TPC server or on a separate machine.

Agents are not automatically removed from the registry when they become inactive or are uninstalled. To identify and remove obsolete agents from the registry, use the agent registration tools that are available in the toolkit subdirectory of the Agent Manager installation directory (\AgentManager\toolkit).

The registry is a database named IBMCDB that contains the current configuration of all known agents and resource managers. The registry contains the identity, certificates, and communication information for each resource manager.

And, the registry contains the following information about agents:

- Identification of every known agent and its computer system
- The certificate issued to each agent
- Basic configuration information about each agent, including information about the type and version of the hardware and operating system
- The configuration and errors of each agent (updated by the agent at a configurable interval)
- Current communication parameters for the agent, including the IP address, the port or ports for which the agent is configured, and the supported protocol
- The installation directory of each agent

The information in the registry is updated by asynchronous events, such as the registration of agents, and by updates from the agent. The agent provides a configuration update when it starts, when a bundle is installed or uninstalled, and at a configurable interval (by default, daily).

In order to manage your TPC agent environment, you need to maintain the *Agent Manager* Database (IBMCDB). *Deregister* is a collection of Java utilities and scripts that assist in the deregistration of agents.

**Note:** You must run the **Deregister** commands from the machine on which the Agent Manager is installed.

The Deregistration toolkit consists of three commands, which are **RetrieveAgents**, **LogicallyDeleteAgents,** and **PurgeAgents**. You need to issue these commands with parameters, which you can save into a configuration file to simplify the command issued. You can store configuration settings for any of your Agent Manager servers in a simple text file (db_info.cfg).

In this command **db_info.cfg / -dbconfig**, the **-dbconfig** parameter specifies a file that contains the database connection information. You use the parameter only if you do not want to type the password information on the command line.

**Note:** Instead of saving the database password in normal text that anyone can read, you have access to the server in the db_info.cfg file. You can use the **-dbPassword** parameter to specify the password on the command line when you run this script.

Before using the **-dbconfig** parameter, you must update the dbconfig file with the database connection information. The db_info.cfg file is located in the AgentManager\toolkit\bin directory, which can be modified to suit your environment. Example 6-1 shows the contents of the db_info.cfg file.

*Example 6-1 db_info.cfg file contents*

```
dbType=[db2|cloudscape|oracle]
dbName=[database name, fully qualified path if cloudscape]
dbUsername=[database user name]
dbPassword=[database password, or provide -dbPassword on invocation]
dbHostname=[host name of target database machine]
dbPort=[port on which target database is listening]
dbDriver=[type2|type4]
```

After you enter the database variable information, you can use this file as shown in Example 6-2.

*Example 6-2 File db_info.cfg with variables entered*

```
dbType=db2
dbName=IBMCDB
dbUsername= db2tpc
dbPassword= db2tpc
dbHostname= gallium.itsosj.sanjose.ibm.com
dbPort=50000
dbDriver=type2
```

> **Note:** The default dbport (`50000`) is created at installation of DB2. To verify the dbDriver type, view the <*install*>\IBM\AgentManager\install\AMInstall.properties file.

## 6.2.5 How to run RetrieveAgents, LogicallyDeleteAgents, and PurgeAgents

On Windows:

1. Start a command window (CMD).
2. Make sure that your path points to the <*install*>\AgentManager\Toolkit\bin directory, type: **`RetrieveAgents`**, **`LogicallyDeleteAgents`**, or **`PurgeAgents`** (see readme for more options)

On UNIX and Linux:

1. Start a terminal session.
2. Make sure that your path points to the <*install*>\AgentManager\Toolkit\bin directory, type: **`RetrieveAgents`**, **`LogicallyDeleteAgents`**, or **`PurgeAgents`** (see readme for more options)

### RetrieveAgents

This script lists the agents in the system. The agents displayed are the list of Common Agents that are known to Agent Manager. Review the readme file in the toolkit directory for arguments and parameters.

The commands provide the following attributes for each agent:

- Agent name
- Host name
- IP address
- Port
- Agent installation directory
- Managed element ID
- Operating system ID

The file for each platform is:

- Windows:

  toolkit\bin\**RetrieveAgents.bat**

- Linux and UNIX:

  toolkit\bin\**RetrieveAgents.sh**

Table 6-1 on page 236 shows the basic parameters.

*Table 6-1 RetrieveAgents command parameter or arguments*

| Parameter | Function and variables |
|---|---|
| -dbpassword | Enter the database password or use the db_info.cfg config file.<br>-dbpassword tpcadmin |
| -dir | Lists all the agents that the Common Agent has installed in a particular directory.<br>Specify the directory in Web site form, for example: file:///C:/Program Files/IBM/TPC/ca |
| -os | List all the agents by a specific operating system ID.<br>Specify the operating system ID: -os 62eec52116b011d8bb860006293337a5 |
| -port | List all agents by the listening port.<br>-port 9510 (default port 9510) |
| -exp | Lists the expiration status:<br>-exp y = only expired agents<br>-exp n = only non-expired agents<br>-exp a = All agents (default) |
| -out | Redirects the agent list to a specified file.<br>-out c:\AgentList.txt |

Issue the **RetrieveAgents** command with many parameters to narrow the search and the results.

Example 6-11 on page 263 is an example of viewing all of the agents registered to this Agent Manager server. We currently have two agents managed by Agent Manager.

## Syntax

The syntax for this command is **retrieveagents -dbpassword db2tpc**

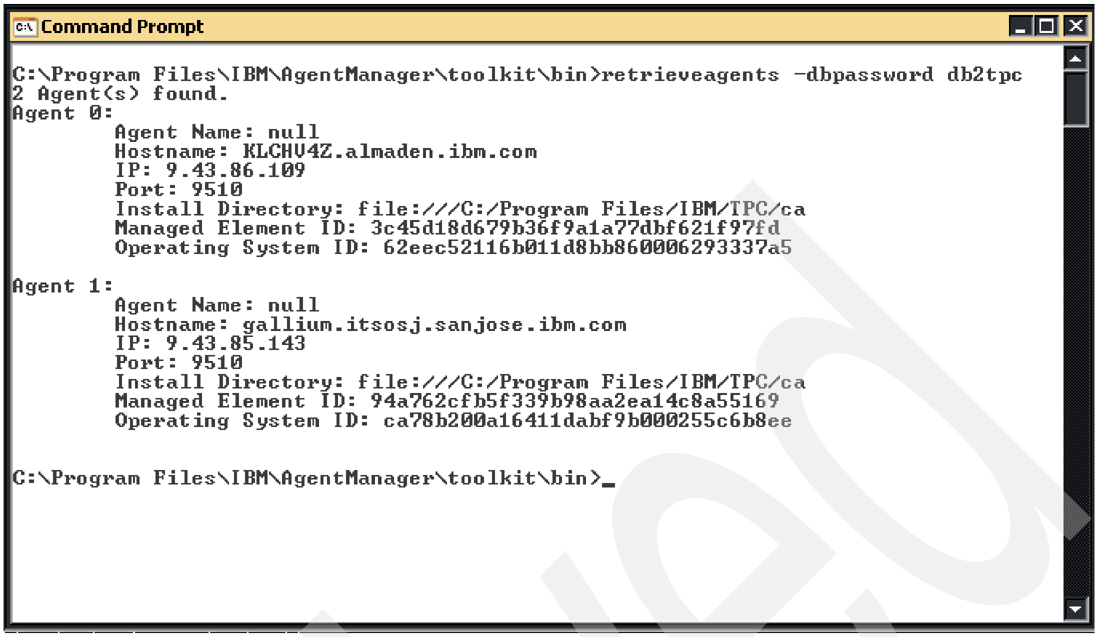

```
C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbpassword db2tpc
2 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: KLCHV4Z.almaden.ibm.com
        IP: 9.43.86.109
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 3c45d18d679b36f9a1a77dbf621f97fd
        Operating System ID: 62eec52116b011d8bb860006293337a5

Agent 1:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-11 RetrieveAgents command output*

Similarly, you can issue this command by using the db_info.cfg file as shown in Figure 6-12.

## Syntax

The syntax is **`retrieveagents -dbconfig db_info.cfg`**

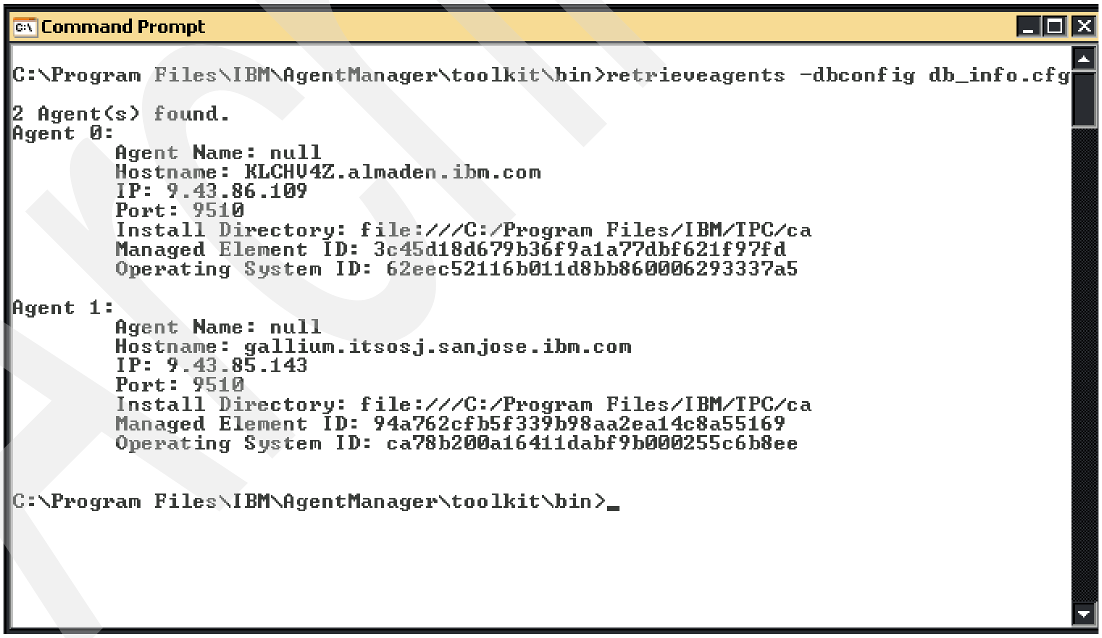

```
C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbconfig db_info.cfg

2 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: KLCHV4Z.almaden.ibm.com
        IP: 9.43.86.109
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 3c45d18d679b36f9a1a77dbf621f97fd
        Operating System ID: 62eec52116b011d8bb860006293337a5

Agent 1:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-12 RetrieveAgents command using db_info.cfg*

Issue the **`RetrieveAgents`** command by using multiple parameters as shown in Figure 6-13 on page 238.

### Syntax

The sample syntax is **`retrieveagents -dbconfig db_info.cfg -port 9510 -exp n -os ca78b200a16411dabf9b000255c6b8ee`**

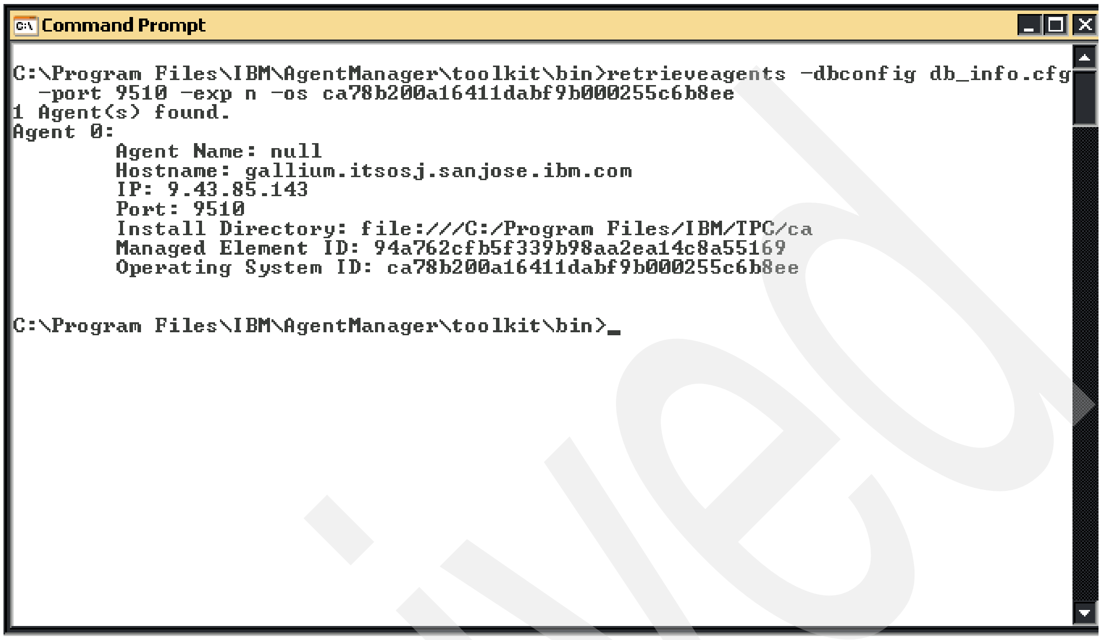

```
C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbconfig db_info.cfg
  -port 9510 -exp n -os ca78b200a16411dabf9b000255c6b8ee
1 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-13 RetrieveAgents command using multiple parameters*

## 6.2.6 LogicallyDeleteAgents

This script sets the delete timestamp to the current timestamp, which logically deletes the agent or agents. Although you have deleted the agent, you can still recover the records for that particular agent, because it has not been removed from the Agent Manager database. Review the readme file in the toolkit directory for arguments and parameters.

The files are:

- For Windows:

  toolkit\bin\**LogicallyDeleteAgents.bat**
- For Linux and UNIX:

  toolkit/bin/**LogicallyDeleteAgents.sh**

Table 6-2 on page 239 shows the basic parameters.

Table 6-2 LogicallyDeleteAgents parameters or arguments

| Parameters | Functions and variables |
|---|---|
| -dbpassword | Enter the database password or use the db_info.cfg config file.<br>-dbpassword tpcadmin |
| -me | Expires agents specified by the managed element ID. Specify one or more IDs in a comma-separated list.<br>-me 94a762cfb5f339b98aa2ea14c8a55169 |
| -os | Expires agents specified by operating system ID. All agents with this OS ID will be expired.<br>-os 62eec52116b011d8bb860006293337a5 |

Example 6-14 on page 273 is an example of expiring an agent with a specific managed element ID.

### Syntax

The syntax is:

```
LogicallyDeleteAgents -dbpassword db2tpc -me3c45d18d679b36f9a1a77dbf621f97fd
```

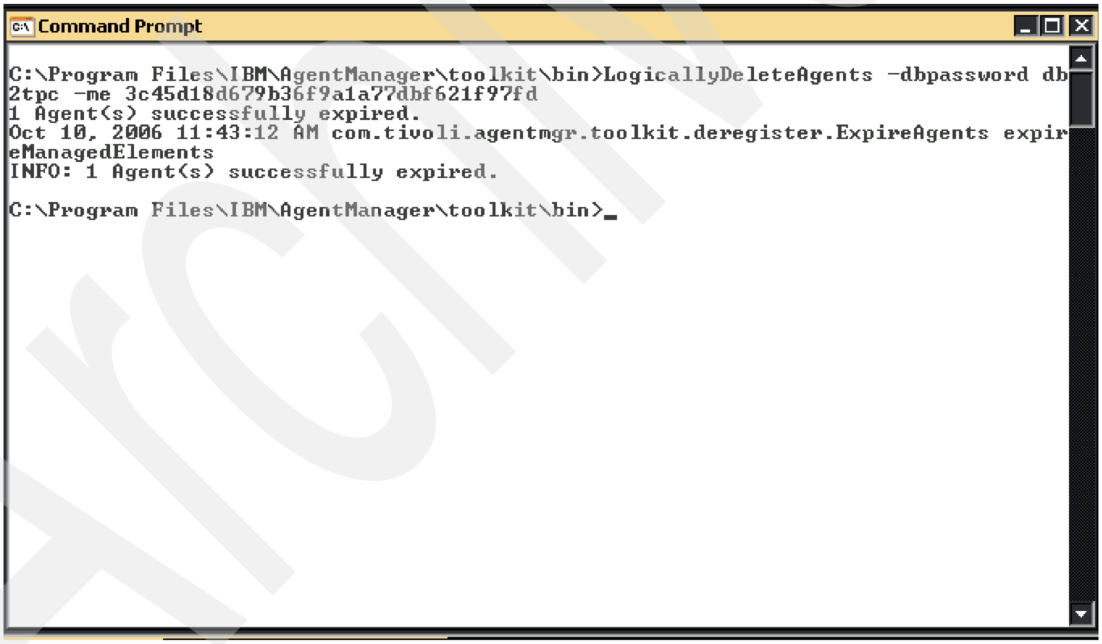

```
C:\Program Files\IBM\AgentManager\toolkit\bin>LogicallyDeleteAgents -dbpassword db
2tpc -me 3c45d18d679b36f9a1a77dbf621f97fd
1 Agent(s) successfully expired.
Oct 10, 2006 11:43:12 AM com.tivoli.agentmgr.toolkit.deregister.ExpireAgents expir
eManagedElements
INFO: 1 Agent(s) successfully expired.

C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

Figure 6-14 LogicallyDeleteAgents command with specific managed element

If you list all active agents using the *RetrieveAgents* command, you notice that one agent has been expired as shown in Figure 6-15 on page 240.

### Syntax

The syntax is **`retrieveagents -dbpassword db2tpc -exp n`**

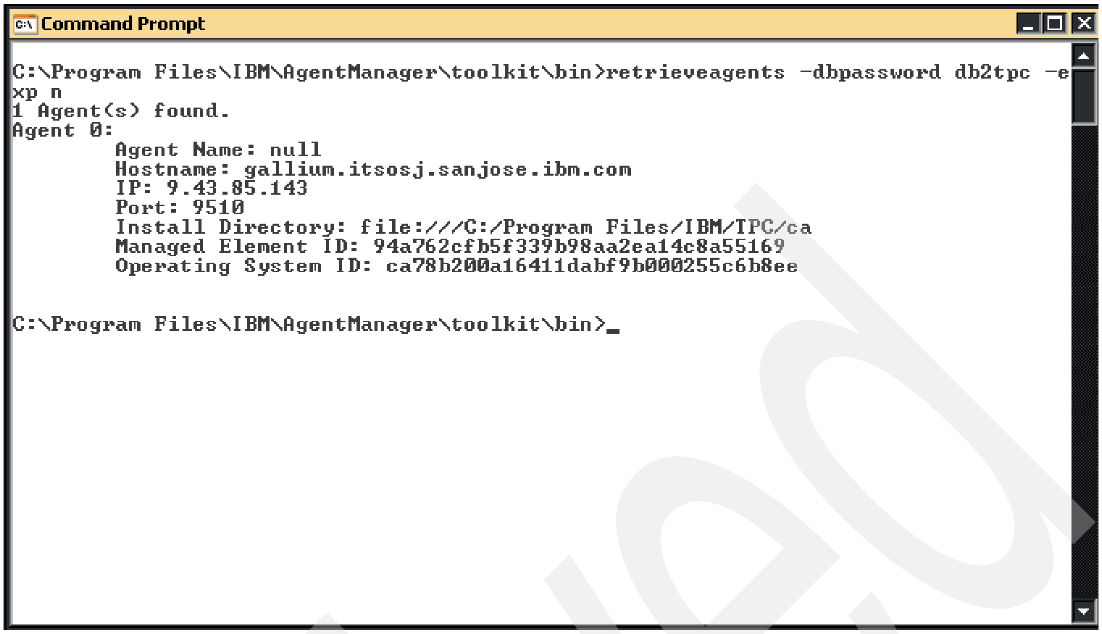

```
C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbpassword db2tpc -e
xp n
1 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-15 RetrieveAgents command with one non-expired agent*

If you list all the Agents using the `RetrieveAgents` command, you notice that we can still view both agents but one agent is expired. When you delete the agent logically, the certificate for that particular agent is revoked, which means that the certificate is added to the Certification Revocation List (CRL). After the certificate is added to the CRL, this certificate is no longer valid.

> **Note: -me** and **-os** are mutually exclusive; you cannot use them cumulatively.

### 6.2.7 PurgeAgents

Use the `PurgeAgents` command to systematically and permanently remove old and unneeded data. The term *purge* is stronger than *delete*. You often can restore deleted objects by undeleting them, but purged objects are gone forever.

When the system database gets too large, it can slow resources, corrupt databases, and use up necessary disk space. It is good to purge the system of all call records when they are no longer necessary. This script deletes agents from the database tables. Review the readme file in the toolkit directory for arguments and parameters.

The files are:

- Windows:

  toolkit\bin\**PurgeAgents.bat**

- Linux and UNIX:

  toolkit/bin/**PurgeAgents.sh**

Table 6-3 on page 241 shows the basic parameters.

*Table 6-3 PurgeAgent parameters or arguments*

| Parameters | Functions and variables |
|---|---|
| -dbpassword | Enter the database password or use the db_info.cfg config file.<br>-dbpassword tpcadmin |
| -me | Delete agents specified by the managed element ID. Specify one or more IDs in a comma-separated list.<br>-me 94a762cfb5f339b98aa2ea14c8a55169 |
| -os | Delete agents specified by operating system ID. All agents with this OS ID will be expired.<br>-os 62eec52116b011d8bb860006293337a5 |

Figure 6-16 is an example of deleting an agent with a specific managed element ID.

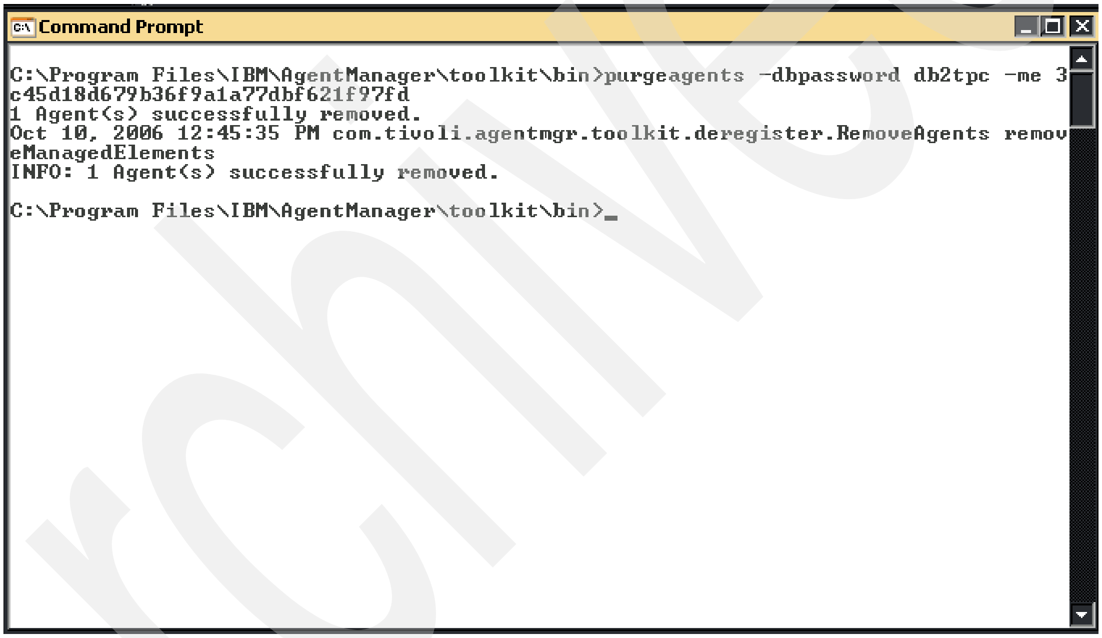

```
C:\Program Files\IBM\AgentManager\toolkit\bin>purgeagents -dbpassword db2tpc -me 3
c45d18d679b36f9a1a77dbf621f97fd
1 Agent(s) successfully removed.
Oct 10, 2006 12:45:35 PM com.tivoli.agentmgr.toolkit.deregister.RemoveAgents remov
eManagedElements
INFO: 1 Agent(s) successfully removed.

C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-16 PurgeAgents command using Managed Element ID*

If you list all agents using the **`RetrieveAgents`** command, you notice that there is only one agent managed by Agent Manager (see Figure 6-17 on page 242).

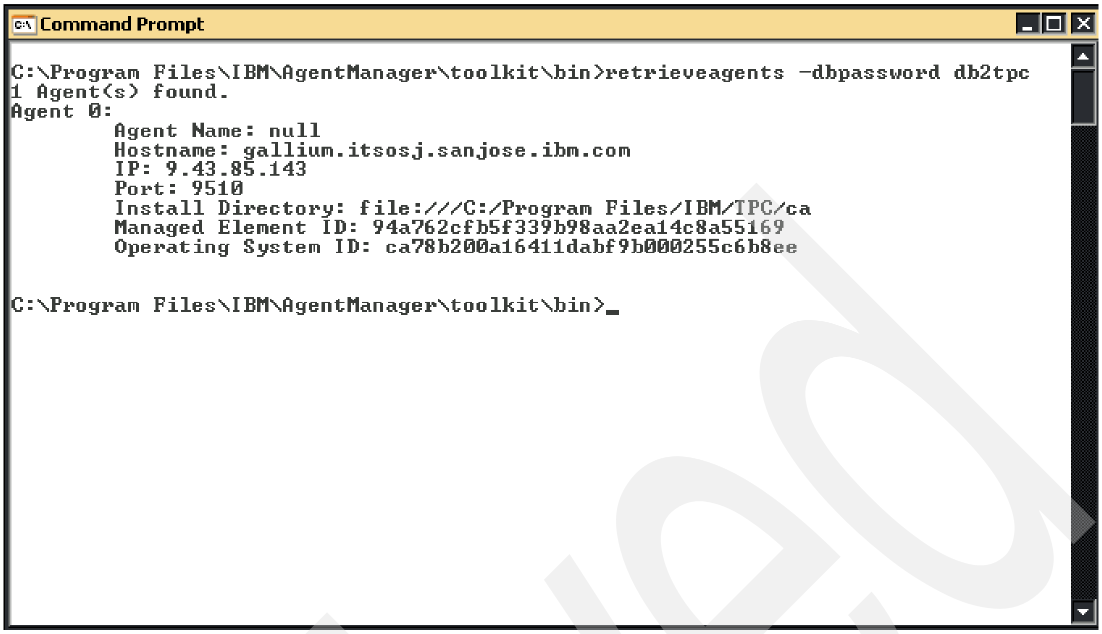

```
C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbpassword db2tpc
1 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-17 RetrieveAgents command verifying all agents managed*

> **Restriction:** You cannot purge an agent if it is not expired (LogicallyDeleted). The parameters **-me** and **-os** are mutually exclusive; you cannot use them cumulatively.

### 6.2.8 Reactivate agents

TPC allows you to reactivate the agents that you have purged in your Agent Manager database. You can only reactivate the agent when the Agent has been successfully “Purged” and not “LogicallyDeleted.”

The process of reactivating the agent requires you to delete the certificate files and restart your IBM Tivoli Common Agent services, so the certificates can be downloaded from Agent Manager again and can be reinstated with normal operation. You do not need to reinstall the agents. These certificates are instrumental in maintaining the connection between the TPC server and the agents.

These steps guide you through the process reactivate a purged agent:

1. Figure 6-18 on page 243 shows that we have two non-expired agents that are currently reported by the TPC server. We used the command:

   ```
   retrieveagents -dbpassword db2tpc -exp n
   ```

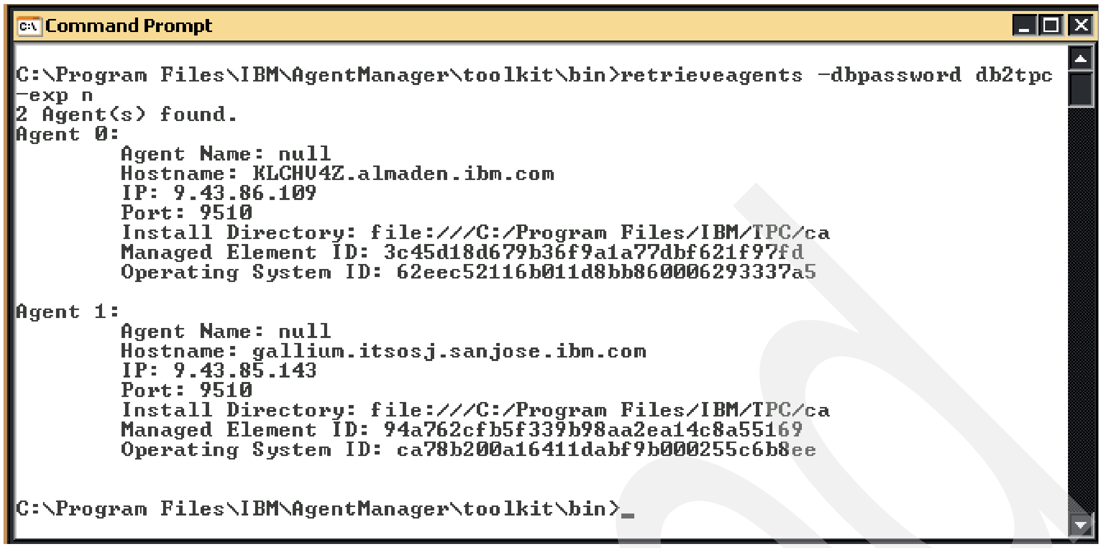

```
C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbpassword db2tpc
-exp n
2 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: KLCHV4Z.almaden.ibm.com
        IP: 9.43.86.109
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 3c45d18d679b36f9a1a77dbf621f97fd
        Operating System ID: 62eec52116b011d8bb860006293337a5

Agent 1:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

Figure 6-18   List of non-expired agents

2. We now “LogicallyDeleteAgents” and then “PurgeAgent,” so `Agent0` (`KLCHV4Z.almaden.ibm.com`) will be expired and purged from the Agent Manager Database as shown in Figure 6-19.

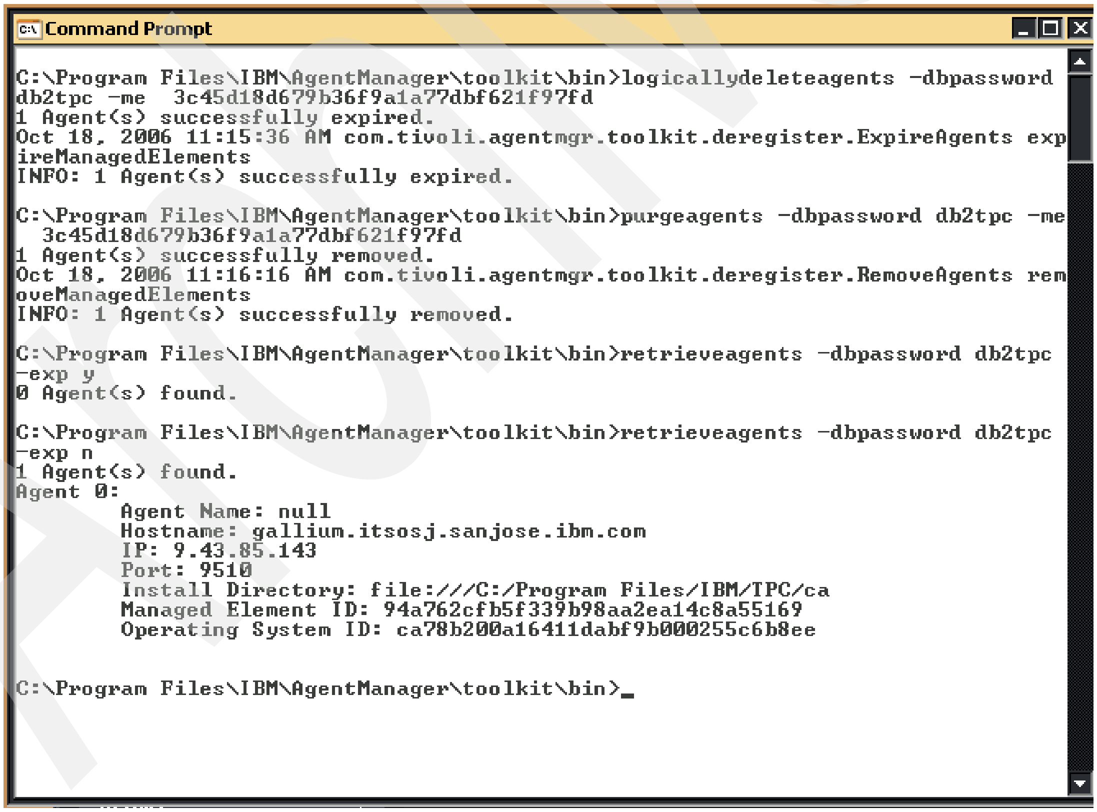

```
C:\Program Files\IBM\AgentManager\toolkit\bin>logicallydeleteagents -dbpassword
db2tpc -me  3c45d18d679b36f9a1a77dbf621f97fd
1 Agent(s) successfully expired.
Oct 18, 2006 11:15:36 AM com.tivoli.agentmgr.toolkit.deregister.ExpireAgents exp
ireManagedElements
INFO: 1 Agent(s) successfully expired.

C:\Program Files\IBM\AgentManager\toolkit\bin>purgeagents -dbpassword db2tpc -me
  3c45d18d679b36f9a1a77dbf621f97fd
1 Agent(s) successfully removed.
Oct 18, 2006 11:16:16 AM com.tivoli.agentmgr.toolkit.deregister.RemoveAgents rem
oveManagedElements
INFO: 1 Agent(s) successfully removed.

C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbpassword db2tpc
-exp y
0 Agent(s) found.

C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbpassword db2tpc
-exp n
1 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

Figure 6-19   Agent KLCHV4Z.almaden.ibm.com purged from the database

3. We will need to remotely log on to the agent `KLCHV4Z.almaden.ibm.com` server to delete the contents of the cert directory. The cert directory is located in the directory shown in Figure 6-20:
   <*install dir*> \Program Files\IBM\TPC\ca

   Do not delete the cert directory, only the contents.

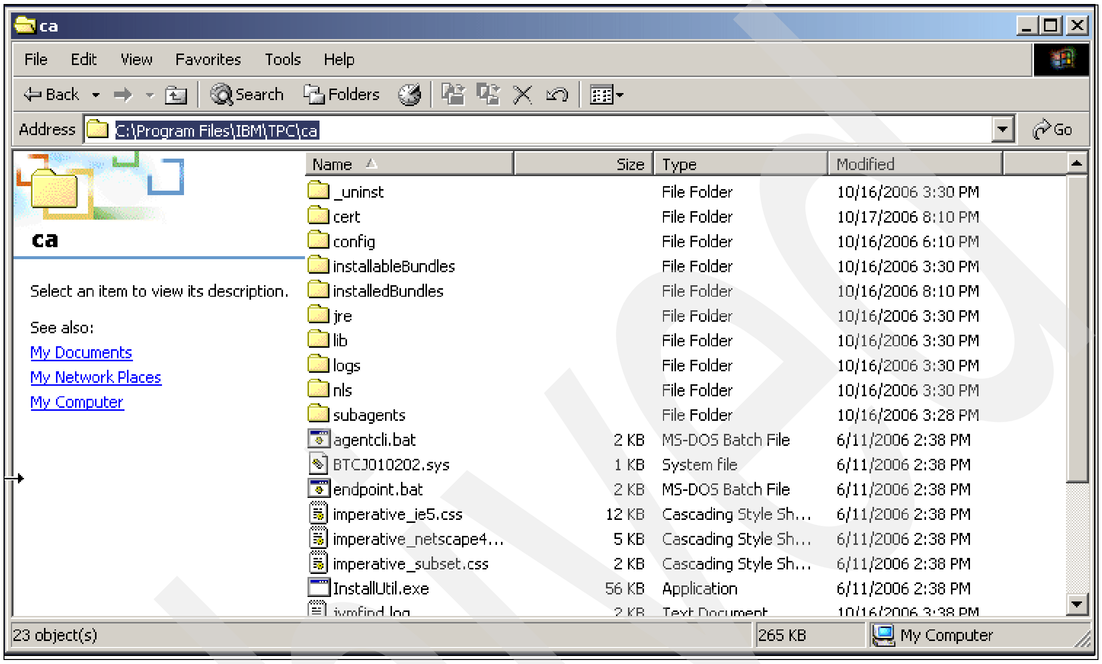

Address: C:\Program Files\IBM\TPC\ca

| Name | Size | Type | Modified |
|---|---|---|---|
| _uninst | | File Folder | 10/16/2006 3:30 PM |
| cert | | File Folder | 10/17/2006 8:10 PM |
| config | | File Folder | 10/16/2006 6:10 PM |
| installableBundles | | File Folder | 10/16/2006 3:30 PM |
| installedBundles | | File Folder | 10/16/2006 8:10 PM |
| jre | | File Folder | 10/16/2006 3:30 PM |
| lib | | File Folder | 10/16/2006 3:30 PM |
| logs | | File Folder | 10/16/2006 3:30 PM |
| nls | | File Folder | 10/16/2006 3:30 PM |
| subagents | | File Folder | 10/16/2006 3:28 PM |
| agentcli.bat | 2 KB | MS-DOS Batch File | 6/11/2006 2:38 PM |
| BTCJ010202.sys | 1 KB | System file | 6/11/2006 2:38 PM |
| endpoint.bat | 2 KB | MS-DOS Batch File | 6/11/2006 2:38 PM |
| imperative_ie5.css | 12 KB | Cascading Style Sh... | 6/11/2006 2:38 PM |
| imperative_netscape4... | 5 KB | Cascading Style Sh... | 6/11/2006 2:38 PM |
| imperative_subset.css | 2 KB | Cascading Style Sh... | 6/11/2006 2:38 PM |
| InstallUtil.exe | 56 KB | Application | 6/11/2006 2:38 PM |
| jvmfind.log | 2 KB | Text Document | 10/16/2006 3:38 PM |

23 object(s) | 265 KB

*Figure 6-20 Cert directory*

> **Tip:** Ensure that you make a backup of the cert directory prior to deleting the contents of the directory.

4. After you delete the contents of the cert directory that contains these files, you need to restart the IBM Tivoli Common Agent services (see Figure 6-21 on page 245):
   – agentKeys.jks
   – agentTrust.jks
   – CertificateRevocationList
   – pwd

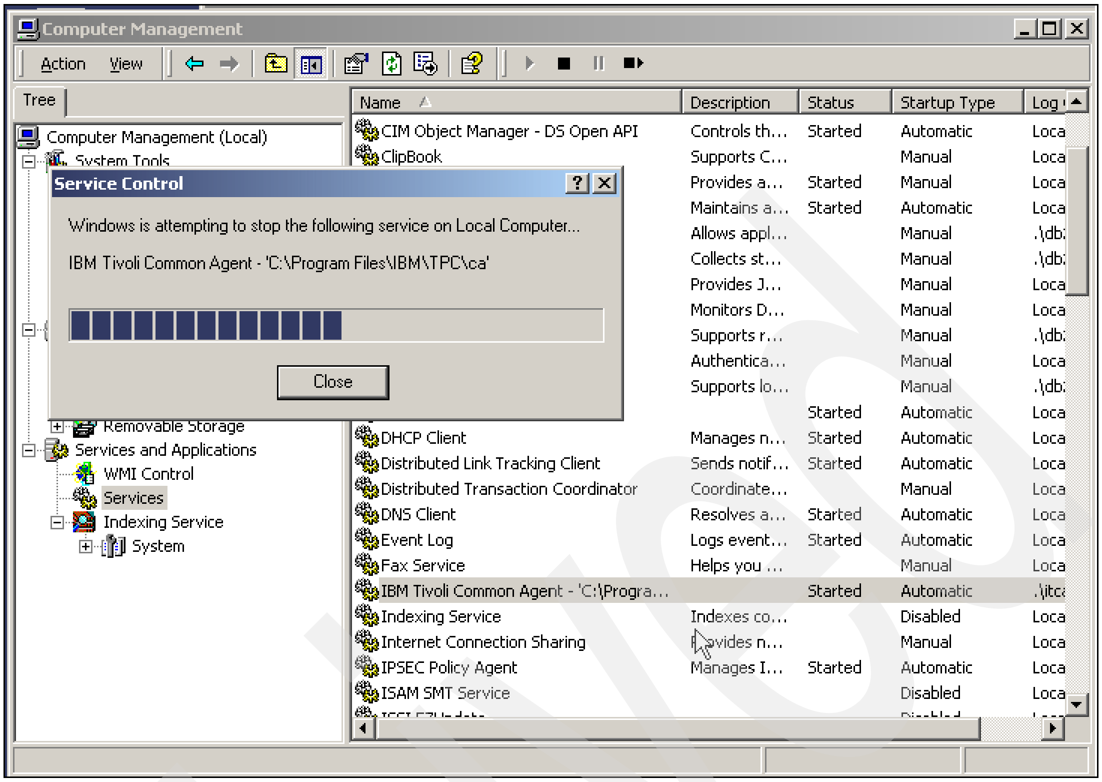

Service Control

Windows is attempting to stop the following service on Local Computer...

IBM Tivoli Common Agent - 'C:\Program Files\IBM\TPC\ca'

| Name | Description | Status | Startup Type | Log |
|---|---|---|---|---|
| CIM Object Manager - DS Open API | Controls th... | Started | Automatic | Loca |
| ClipBook | Supports C... | | Manual | Loca |
| | Provides a... | Started | Manual | Loca |
| | Maintains a... | Started | Automatic | Loca |
| | Allows appl... | | Manual | .\db: |
| | Collects st... | | Manual | .\db: |
| | Provides J... | | Manual | Loca |
| | Monitors D... | | Manual | Loca |
| | Supports r... | | Manual | .\db: |
| | Authentica... | | Manual | Loca |
| | Supports lo... | | Manual | .\db: |
| | | Started | Automatic | Loca |
| DHCP Client | Manages n... | Started | Automatic | Loca |
| Distributed Link Tracking Client | Sends notif... | Started | Automatic | Loca |
| Distributed Transaction Coordinator | Coordinate... | | Manual | Loca |
| DNS Client | Resolves a... | Started | Automatic | Loca |
| Event Log | Logs event... | Started | Automatic | Loca |
| Fax Service | Helps you ... | | Manual | Loca |
| IBM Tivoli Common Agent - 'C:\Progra... | | Started | Automatic | .\itca |
| Indexing Service | Indexes co... | | Disabled | Loca |
| Internet Connection Sharing | Provides n... | | Manual | Loca |
| IPSEC Policy Agent | Manages I... | Started | Automatic | Loca |
| ISAM SMT Service | | | Disabled | Loca |

*Figure 6-21 Restart IBM Tivoli Common Agent*

5. The files deleted in the cert directory will be recreated after the restart of the IBM Tivoli Common Agent Service (see Figure 6-22).

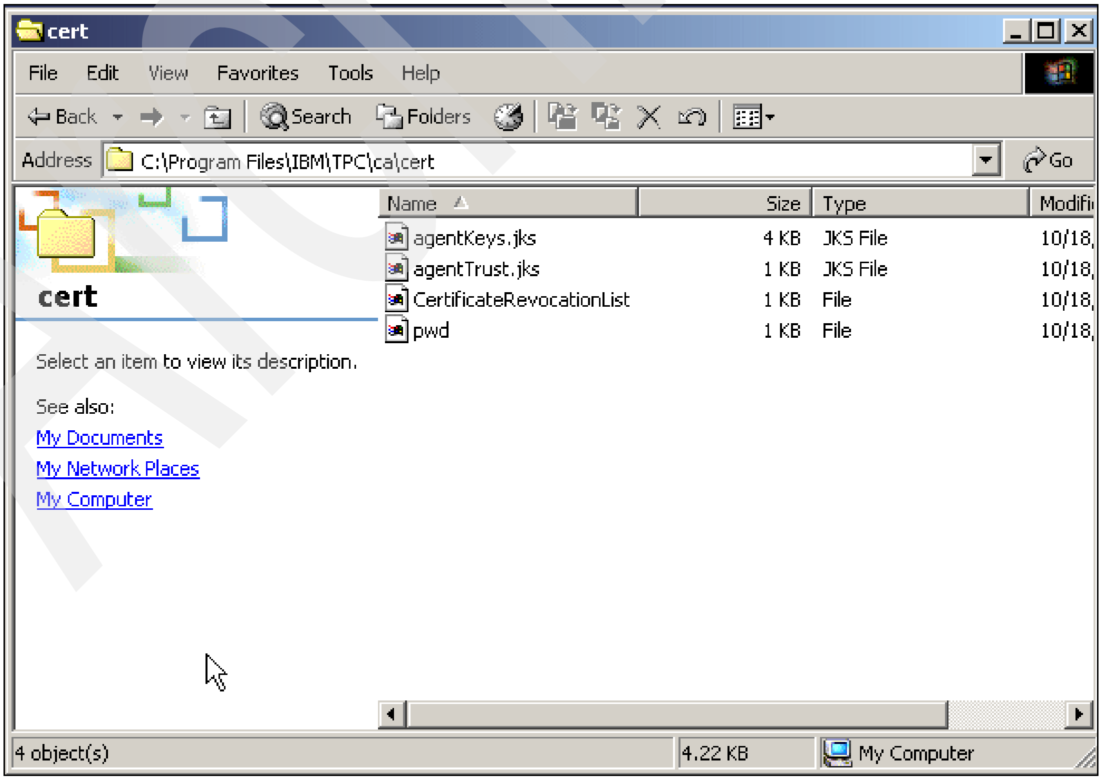

Address C:\Program Files\IBM\TPC\ca\cert

| Name | Size | Type | Modifi |
|---|---|---|---|
| agentKeys.jks | 4 KB | JKS File | 10/18, |
| agentTrust.jks | 1 KB | JKS File | 10/18, |
| CertificateRevocationList | 1 KB | File | 10/18, |
| pwd | 1 KB | File | 10/18, |

*Figure 6-22 Contents of the cert directory recreated*

6. The Agent will now reactivate itself with the Agent Manager Database and normal operations will resume. If we list the Agents again using the **`RetrieveAgents`** command, we notice the agent has been reactivated as shown in Figure 6-23.

   **Syntax**

   The syntax is **`retrieveagents -dbpassword db2tpc`**

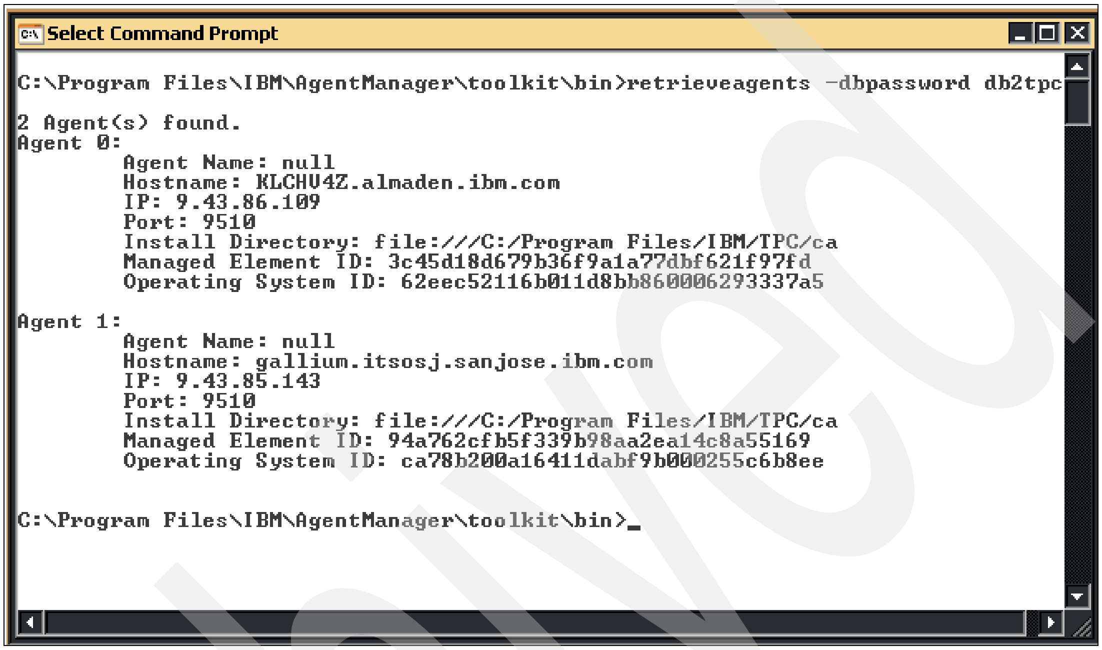

```
C:\Program Files\IBM\AgentManager\toolkit\bin>retrieveagents -dbpassword db2tpc

2 Agent(s) found.
Agent 0:
        Agent Name: null
        Hostname: KLCHV4Z.almaden.ibm.com
        IP: 9.43.86.109
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 3c45d18d679b36f9a1a77dbf621f97fd
        Operating System ID: 62eec52116b011d8bb860006293337a5

Agent 1:
        Agent Name: null
        Hostname: gallium.itsosj.sanjose.ibm.com
        IP: 9.43.85.143
        Port: 9510
        Install Directory: file:///C:/Program Files/IBM/TPC/ca
        Managed Element ID: 94a762cfb5f339b98aa2ea14c8a55169
        Operating System ID: ca78b200a16411dabf9b000255c6b8ee


C:\Program Files\IBM\AgentManager\toolkit\bin>_
```

*Figure 6-23 Agent reactivated*

### Toolkit example

This a good example of when these tools are very important. The Common Agent is running on a workstation. The workstation is stolen from your server room, and an unauthorized person might use the certificates to connect to your TPC server. In this situation, delete the agent as fast as possible to revoke the stolen certificates. These tools are designed to manage your Agent Manager database to help you to achieve the optimal performance from your application.

## 6.3 Collectlogs tool

The Collectlogs tool allows you to collect all information required for analysis by IBM support. The tool only collects information that is currently being monitored by the CIMOM server. The tool does not collect any configuration information of the TPC server and the underlying infrastructure, for example, LUN, array, MDisk, or I/O group information.

### 6.3.1 DS Open API CIM Agent

If you have problems with DS Open API CIM Agent, the **collectlogs** tool helps collect information for all installed IBM DS subsystems from a centralized location. This tool detects the configuration of the system on which it is running and collects the appropriate information. The information will be placed in a .zip file. This tool is installed upon installation of the DS Open API CIM Agent code. Support will request the tool when it needs to be run for analysis. The script generates a collectedLogs_*YYYY-MM-DD_HH-MM-SS*.zip file. *YYYY-MM-DD_HH-MM-SS* represents the timestamp where the logs were collected.

## Output of the DS Open API Collectlogs tool

The output of this tool is:

- Java home, version, and classpath
- JVM implementation name and version
- CIMOM installation and uninstallation logs
- All config files
- Diagnostic information regarding the system and its services
- Internal service logs
- Trace logs
- Provider logs
- CIMOM properties and configuration files

The files are:

- For Windows:

  C:\Program Files\IBM\CIMAgent\

  Run the following command: **`run collectLogs.bat`**

- For Linux or UNIX:

  /opt/IBM/cimagent/

  Run the following command: **`run collectLogs.sh`**

## How to run

The steps are:

1. Navigate to the directory where the collectlog.bat file is located: C:\Program Files\IBM\CIMAgent\
2. Run **`Collectlogs.bat`** (see Figure 6-24).

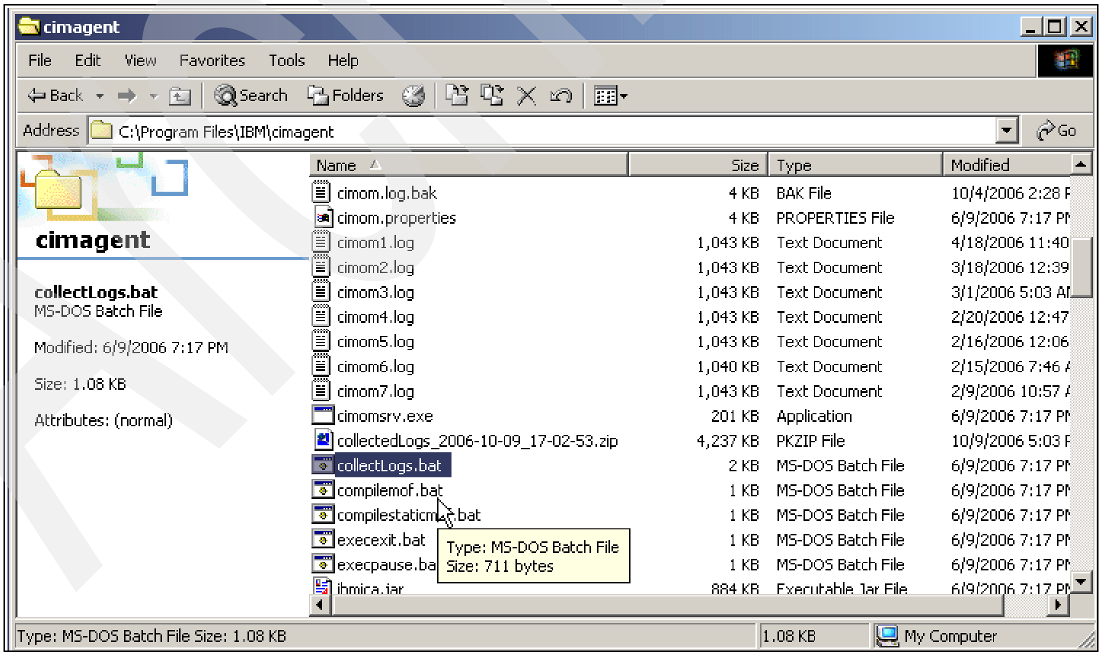

Address: C:\Program Files\IBM\cimagent

cimagent

collectLogs.bat
MS-DOS Batch File

Modified: 6/9/2006 7:17 PM

Size: 1.08 KB

Attributes: (normal)

| Name | Size | Type | Modified |
|---|---|---|---|
| cimom.log.bak | 4 KB | BAK File | 10/4/2006 2:28 F |
| cimom.properties | 4 KB | PROPERTIES File | 6/9/2006 7:17 PI |
| cimom1.log | 1,043 KB | Text Document | 4/18/2006 11:40 |
| cimom2.log | 1,043 KB | Text Document | 3/18/2006 12:39 |
| cimom3.log | 1,043 KB | Text Document | 3/1/2006 5:03 AI |
| cimom4.log | 1,043 KB | Text Document | 2/20/2006 12:47 |
| cimom5.log | 1,043 KB | Text Document | 2/16/2006 12:06 |
| cimom6.log | 1,040 KB | Text Document | 2/15/2006 7:46 / |
| cimom7.log | 1,043 KB | Text Document | 2/9/2006 10:57 / |
| cimomsrv.exe | 201 KB | Application | 6/9/2006 7:17 PI |
| collectedLogs_2006-10-09_17-02-53.zip | 4,237 KB | PKZIP File | 10/9/2006 5:03 F |
| collectLogs.bat | 2 KB | MS-DOS Batch File | 6/9/2006 7:17 PI |
| compilemof.bat | 1 KB | MS-DOS Batch File | 6/9/2006 7:17 PI |
| compilestaticm bat | 1 KB | MS-DOS Batch File | 6/9/2006 7:17 PI |
| execexit.bat | 1 KB | MS-DOS Batch File | 6/9/2006 7:17 PI |
| execpause.ba | 1 KB | MS-DOS Batch File | 6/9/2006 7:17 PI |
| ibmcia.jar | 884 KB | Executable Jar File | 6/9/2006 7:17 PI |

Type: MS-DOS Batch File
Size: 711 bytes

Type: MS-DOS Batch File Size: 1.08 KB | 1.08 KB | My Computer

*Figure 6-24 DS Open API CollectLogs.bat*

3. A command window runs in the foreground verifying that the information that has been collected for analysis and saves it in the collectedLogs_*YYYY-MM-DD_HH-MM-SS*.zip file (see Figure 6-25).

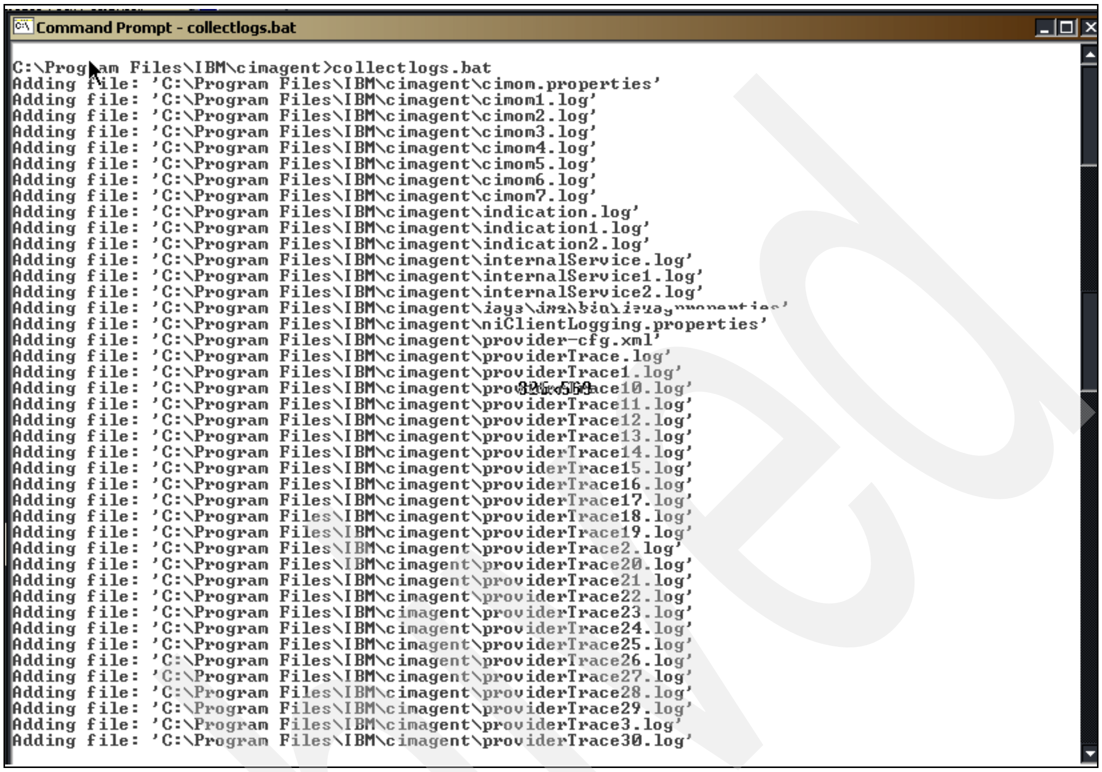

```
C:\Program Files\IBM\cimagent>collectlogs.bat
Adding file: 'C:\Program Files\IBM\cimagent\cimom.properties'
Adding file: 'C:\Program Files\IBM\cimagent\cimom1.log'
Adding file: 'C:\Program Files\IBM\cimagent\cimom2.log'
Adding file: 'C:\Program Files\IBM\cimagent\cimom3.log'
Adding file: 'C:\Program Files\IBM\cimagent\cimom4.log'
Adding file: 'C:\Program Files\IBM\cimagent\cimom5.log'
Adding file: 'C:\Program Files\IBM\cimagent\cimom6.log'
Adding file: 'C:\Program Files\IBM\cimagent\cimom7.log'
Adding file: 'C:\Program Files\IBM\cimagent\indication.log'
Adding file: 'C:\Program Files\IBM\cimagent\indication1.log'
Adding file: 'C:\Program Files\IBM\cimagent\indication2.log'
Adding file: 'C:\Program Files\IBM\cimagent\internalService.log'
Adding file: 'C:\Program Files\IBM\cimagent\internalService1.log'
Adding file: 'C:\Program Files\IBM\cimagent\internalService2.log'
Adding file: 'C:\Program Files\IBM\cimagent\java\jre\bin\java.properties'
Adding file: 'C:\Program Files\IBM\cimagent\niClientLogging.properties'
Adding file: 'C:\Program Files\IBM\cimagent\provider-cfg.xml'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace1.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace10.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace11.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace12.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace13.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace14.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace15.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace16.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace17.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace18.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace19.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace2.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace20.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace21.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace22.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace23.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace24.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace25.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace26.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace27.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace28.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace29.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace3.log'
Adding file: 'C:\Program Files\IBM\cimagent\providerTrace30.log'
```

*Figure 6-25 DS Open API command window verifying the information collected*

4. The collected information is placed in a collectedLogs_*YYYY-MM-DD_HH-MM-SS*.zip file in the directory C:\Program Files\IBM\CIMAgent\.

You can send these files to the IBM support center to help you resolve problems on DS Open API CIM Agent.

## 6.3.2 SAN Volume Controller

Similarly, if you are having problems collecting performance information with SVC CIM Agent, you can run the similar script to collect log and configuration information. This tool is run on the SVC Hardware Management Console (HMC). The CIM Agent is an integrated part of the HMC software.

### Output of the SAN Volume Controller (SVC) Collectlogs tool

This tool collects the following information:

- Server exception logs
- Provider logs where the performance information from the SVC is captured
- CIMOM logs where the TPC server and the CIM Agent communicate
- Installation logs of the SAN Volume Controller (SVC) Master Console
- SVC service log
- The required configuration files

The file is in Windows: C:\Program Files\IBM\svcconsole\support\

## How to run

The steps are:

1. Navigate to the directory where the collectlogs.bat file is located: C:\Program Files\IBM\svcconsole\support\
2. Run **`Collectlogs.bat`** (see Figure 6-26).

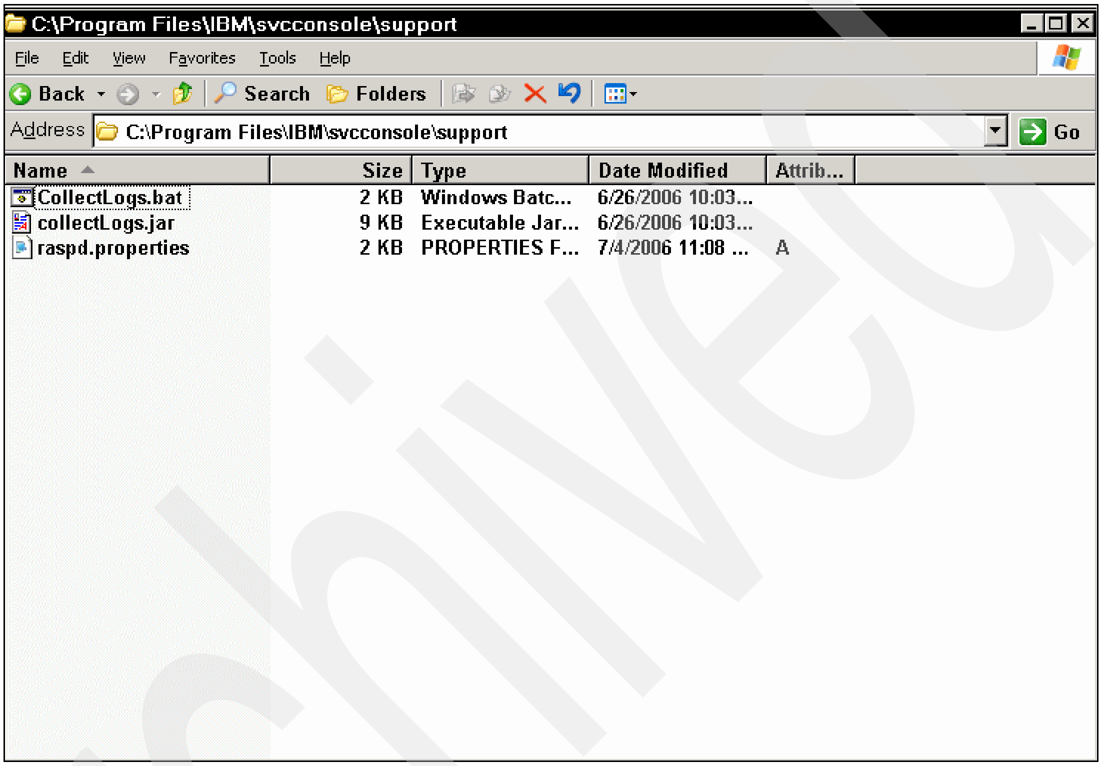
C:\Program Files\IBM\svcconsole\support
File
Edit
View
Favorites
Tools
Help
Back
Search
Folders
Address
C:\Program Files\IBM\svcconsole\support
Go

| Name | Size | Type | Date Modified | Attrib... |
|---|---|---|---|---|
| CollectLogs.bat | 2 KB | Windows Batc... | 6/26/2006 10:03... | |
| collectLogs.jar | 9 KB | Executable Jar... | 6/26/2006 10:03... | |
| raspd.properties | 2 KB | PROPERTIES F... | 7/4/2006 11:08 ... | A |

*Figure 6-26 SVC CollectLogs.bat*

3. A command window runs in the foreground verifying the information that has been collected for analysis as shown in Figure 6-27 on page 250.

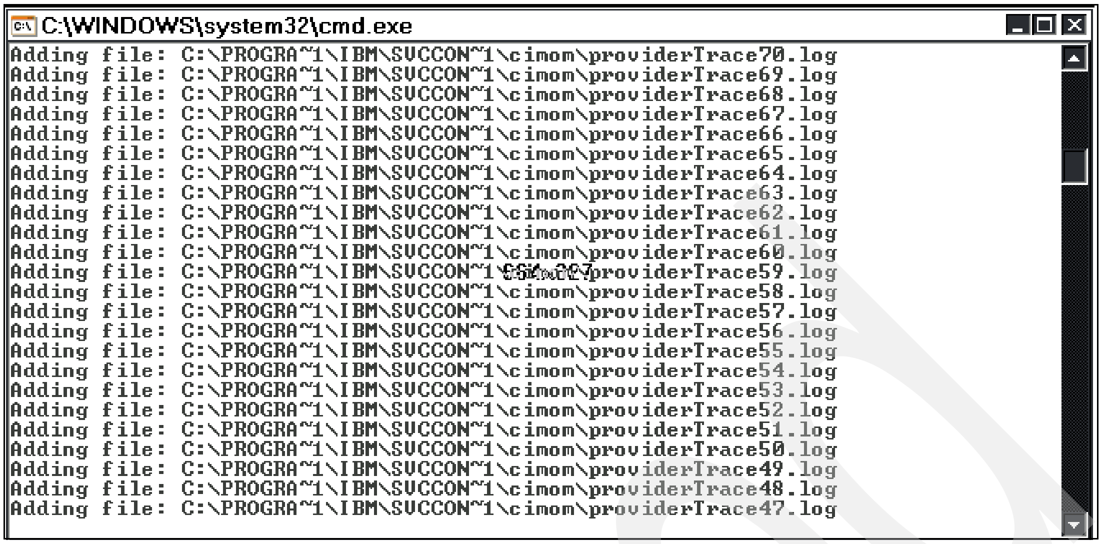

```
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace70.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace69.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace68.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace67.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace66.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace65.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace64.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace63.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace62.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace61.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace60.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace59.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace58.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace57.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace56.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace55.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace54.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace53.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace52.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace51.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace50.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace49.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace48.log
Adding file: C:\PROGRA~1\IBM\SVCCON~1\cimom\providerTrace47.log
```

*Figure 6-27 Command window verifying the collected information*

4. The information collected is placed in a collectedLogs.zip file in the following directory: C:\Program Files\IBM\svcconsole\support\

You can send these files to the IBM support center to help you resolve problems with the SAN Volume Controller CIM Agent.

## 6.4 TPCTOOL

Many TPC users want the ability to produce performance reports from the TPC database with the ability to produce multiple metric graphs as provided by other IBM Storage products. IBM TPC developed TPCTOOL.

TPCTOOL is command line interface (CLI) program, which interacts with the TPC Device server. Commands are entered as lines of text, that is, sequences of types of characters, from a keyboard and output can be received as text.

Characteristics of TPCTOOL are:

- The tool provides queries, management, and reporting capabilities.
- You cannot initiate discoveries, probes, and performance collection from the tool.
- With the release of TPC V3.1, Storage Resource Management Command Processor (SCRMCP), perfcli (Performance Management), and AgentCLI have been integrated into the TPCTOOL.
- You can install TPCTOOL anywhere. The CLI code can be deployed on your personal computer that has access to the network where the TPC server is located.
- The tool connects through TCP/HTTP/SOAP to the Web Services API.
- You can use the tool for storage provisioning and management.
- Standalone Java Client.
- Authentication:
  - User authentication, which requires a user ID and password authenticated in the Device server's authentication domain. Role-based authorization is enforced on a per command basis.
  - Super user or host-based authentication is for the user ID *tpc_superuser*. The password for this user ID bypasses role-based authorization. This is the required authentication method for AIX-based and UNIX-based Device servers.

You install the **tpctool** in the following default directories:

- Windows:

  C:\Program Files\IBM\TPC\cli
- UNIX or Linux:

  /<usr or opt>/IBM/TPC/cli

You must set your path or CHDIR to <*install*>/cli (see Figure 6-28).

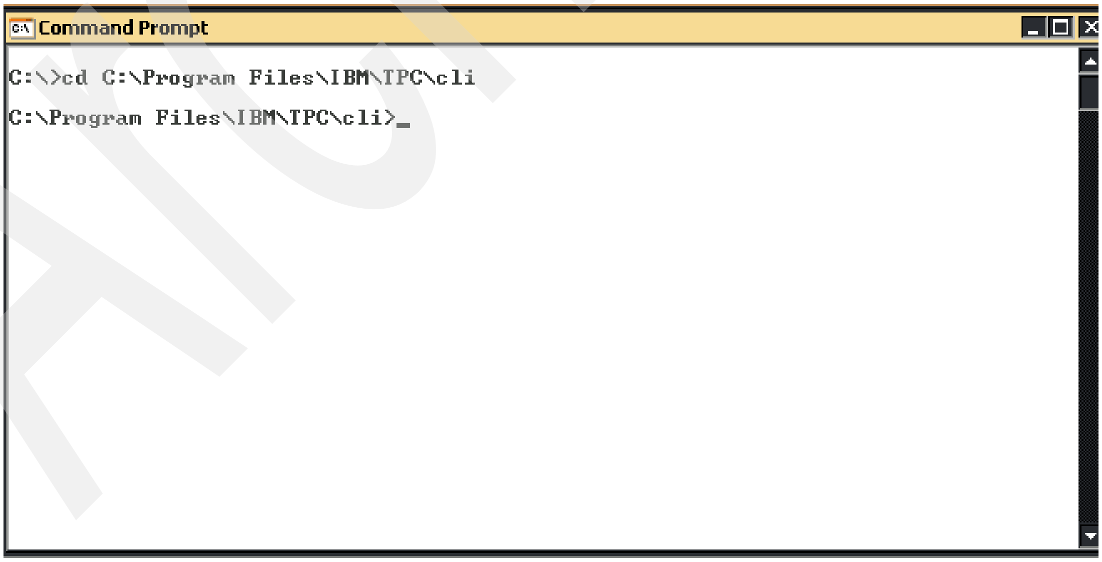

```
C:\>cd C:\Program Files\IBM\TPC\cli

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-28 Set path to tpctool*

## 6.4.1 Basic syntax

This section provides an example of how to display syntax for a TPCTOOL command.

The TPCTOOL command is the basis of the TotalStorage Productivity Center CLI. Consider a CLI as consisting of syntax and semantics. The semantic defines the type of operations that are possible, and on what sort of data these operations can be performed. The CLI displays a prompt, the user enters a command and terminates the command. The command is executed and the result is text output. The command either can be issued on its own, which starts an interactive session, or it must be the first element in a command.

> **Tip:** To display keywords and arguments for a command, enter a question mark (?) or type help at the configuration prompt or after entering part of a command followed by a space.

The basic syntax for TPCTOOL is:

**tpctool <*command*> -url <*host:port*> -user <*uid*> -pwd <*pwd*> <*parms*>**

The variables are:

- <*command*> cli command such as lsdev, encrypt, lsvol, and so forth
- <*host:port*> is a Device server host: port, for example, `localhost:9550`
- <*uid*> is either a user known to the Device server node or the ‘tpc_superuser’
- <*pwd*> is the password for the user (uid)
- <*parms*> are any parameters, options, or arguments to the command

TPCTOOL has a help function to assist you:

**`-help | h | -?`** Lists help for the command.

## 6.4.2 Command modes

Because the CLI is divided into many modes, the commands available to you at any given time depend on the mode that you are in at that moment. Entering a question mark (?) or help at the CLI prompt allows you to obtain a list of commands available for each command mode.

### Single-shot mode

Use the TPCTOOL CLI single-shot command mode if you want to issue a single occasional command.

You must supply the login information and issue the command that you want to process at the same time. Perform the following steps to use the single-shot mode:

1. Start a command window (CMD).
2. You must set your path or CHDIR to <*install*>/cli.
3. From the <*install*>/cli directory, type your command at the shell prompt:

```
shell> tpctool lsdev -user tpcadmin -pwd tpcadmin -url 9.43.85.143:9550
```

### Interactive mode

Use the TPCTOOL CLI interactive command mode when you have multiple transactions to process that cannot be incorporated into a script. The interactive command mode provides a history function that makes repeating or checking prior command usage easy to do.

You can enter the interactive mode by entering the **tpctool** command with no command line options. Perform the following steps to use the interactive mode:

1. Start a command window (CMD).
2. You must set your path or CHDIR to <*install*>/cli.
3. From the <*install*>/cli directory, type **tpctool** at the shell prompt, and you are now within the interactive session.
4. At the prompt, you can enter any valid TPCTOOL CLI command.

```
shell> tpctool
tpctool> lsdev -user tpcadmin -pwd tpcadmin -url 9.43.85.143:9550
```

### Multiple/Script command mode

You can create a file that contains multiple TPCTOOL CLI commands. Login commands can be included in the command file.

Use the TPCTOOL CLI Multiple/Script command mode if you want to issue a sequence of CLI commands. Administrators can use this mode to create automated processes, for example, establishing volume performance reports for SVC.

Consider the following when using the TPCTOOL CLI Multiple/Script command mode:

- The TPCTOOL CLI script can contain only TPCTOOL CLI commands. Use of shell commands results in a process failure. The syntax is:

  ```
  shell> tpctool -script
  ```

  - You can add comments to the scripts. Comments must be prefixed by the number sign (#), for example, `# This script contains a list of metrics available for DS8000 subsystem volume performance.`

## 6.4.3 Output syntax to a file

For programs that display a lot of text, such as TPCTOOL, consider redirecting text, which is usually displayed, to a file. Displaying a lot of text slows down execution; scrolling text in a terminal window on a workstation can cause an I/O bottleneck and use more CPU time. The contents of the syntax can be redirected to a file for later use.

The command in Example 6-3 shows how to run the tool more efficiently by redirecting output to a file and then displaying the program output.

*Example 6-3 Redirecting TPCTOOL output*

```
tpctool lsdev -user ***** -pwd ***** -url localhost:9550
> C:\reports\Output.txt
```

You can import the file into a spreadsheet to create custom reports.

## 6.4.4 Return codes

CLI and any scheduler usually exit with return codes that accurately reflect the success or failure of the operation.

Table 6-4 on page 254 contains the codes returned by TPCTOOL.

*Table 6-4 Return codes used by TPCTOOL*

| Code | Description |
|---|---|
| 0 | The command completed successfully. |
| 1 | The command was unknown to tpctool and was not resolved as an alias. |
| 2 | A required option was not provided. |
| 3 | An option was unknown to tpctool or was not applicable to the command. |
| 4 | An option was missing a required parameter. |
| 5 | An option contained a parameter that was of an invalid format. |
| 6 | An argument was in an invalid format. |
| 7 | An extraneous argument or argument list was provided. |
| 8 | The tpctool client could not connect to the Device server. |
| 9 | The tpctool client could not log in to the Device server using the specified credentials. |
| 10 | The specified credentials are not authorized to perform the requested action. |
| 11 | A required component, for example, Disk Manager or Fabric Manager, is not installed and enabled. |
| 12 | The command might have started but the connection with the Device server was lost. The command might not be completed successfully. |
| 13 | Some operations were completed partially before the Device server returned a failure. |
| 14 | The command failed. |

## 6.4.5 How to remotely install the TPCTOOL CLI

The TPCTOOL CLI can be installed on any workstation that has direct access to the TPC server. By remotely running commands by using CLI, you do not need to log on directly to the TPC server. Scripts can be scheduled to be run when required from the remote CLI installation.

> **Restriction:** The TPCTOOL CLI must have direct access to the server; there must be no firewall constraints between them.

### How to deploy the TPCTOOL CLI

You need to have the TPC server code available to install the TPCTOOL CLI.

The steps are:

1. Browse to the software and run **`setup.exe`** as shown in Figure 6-29 on page 255, or use Disk 1 of the TPC code to install the CLI software.

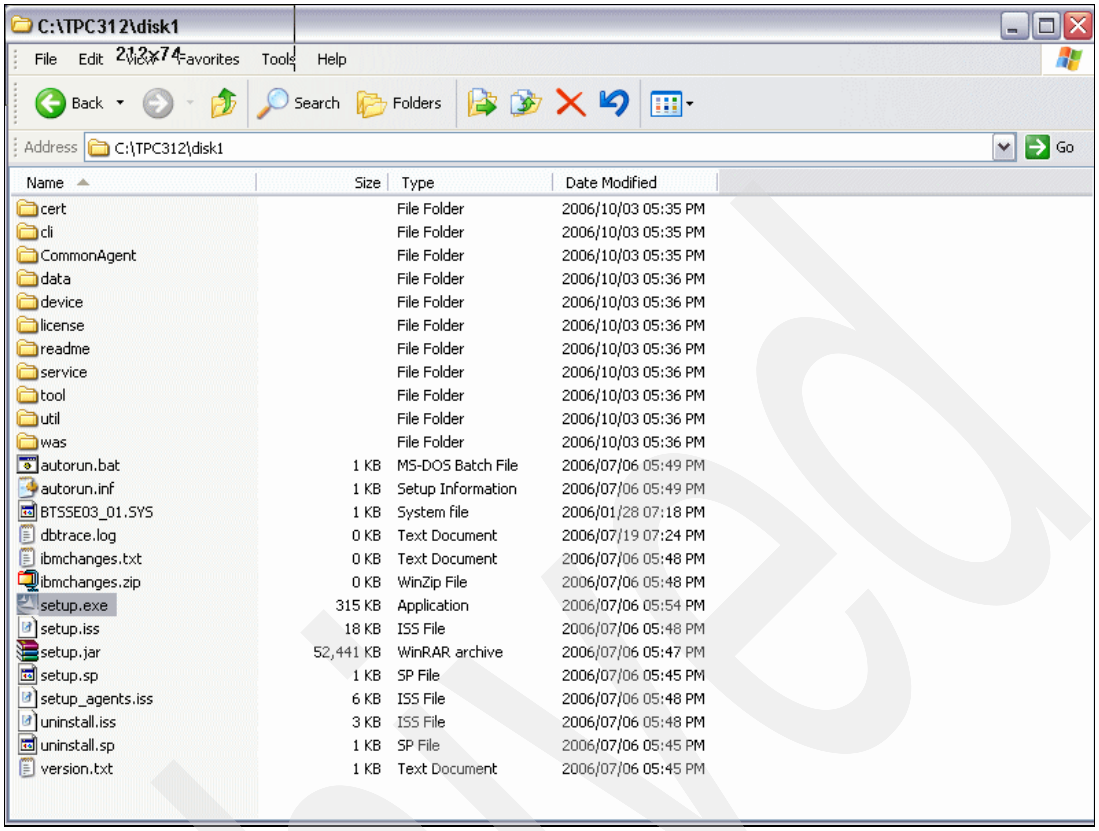

| Name | Size | Type | Date Modified |
|---|---|---|---|
| cert | | File Folder | 2006/10/03 05:35 PM |
| cli | | File Folder | 2006/10/03 05:35 PM |
| CommonAgent | | File Folder | 2006/10/03 05:35 PM |
| data | | File Folder | 2006/10/03 05:36 PM |
| device | | File Folder | 2006/10/03 05:36 PM |
| license | | File Folder | 2006/10/03 05:36 PM |
| readme | | File Folder | 2006/10/03 05:36 PM |
| service | | File Folder | 2006/10/03 05:36 PM |
| tool | | File Folder | 2006/10/03 05:36 PM |
| util | | File Folder | 2006/10/03 05:36 PM |
| was | | File Folder | 2006/10/03 05:36 PM |
| autorun.bat | 1 KB | MS-DOS Batch File | 2006/07/06 05:49 PM |
| autorun.inf | 1 KB | Setup Information | 2006/07/06 05:49 PM |
| BTSSE03_01.SYS | 1 KB | System file | 2006/01/28 07:18 PM |
| dbtrace.log | 0 KB | Text Document | 2006/07/19 07:24 PM |
| ibmchanges.txt | 0 KB | Text Document | 2006/07/06 05:48 PM |
| ibmchanges.zip | 0 KB | WinZip File | 2006/07/06 05:48 PM |
| setup.exe | 315 KB | Application | 2006/07/06 05:54 PM |
| setup.iss | 18 KB | ISS File | 2006/07/06 05:48 PM |
| setup.jar | 52,441 KB | WinRAR archive | 2006/07/06 05:47 PM |
| setup.sp | 1 KB | SP File | 2006/07/06 05:45 PM |
| setup_agents.iss | 6 KB | ISS File | 2006/07/06 05:48 PM |
| uninstall.iss | 3 KB | ISS File | 2006/07/06 05:48 PM |
| uninstall.sp | 1 KB | SP File | 2006/07/06 05:45 PM |
| version.txt | 1 KB | Text Document | 2006/07/06 05:45 PM |

*Figure 6-29 Browse to software and select setup.exe*

2. Select the language to be used during the install.
3. Accept the license agreement (see Figure 6-30 on page 256) and click **Next** to continue.

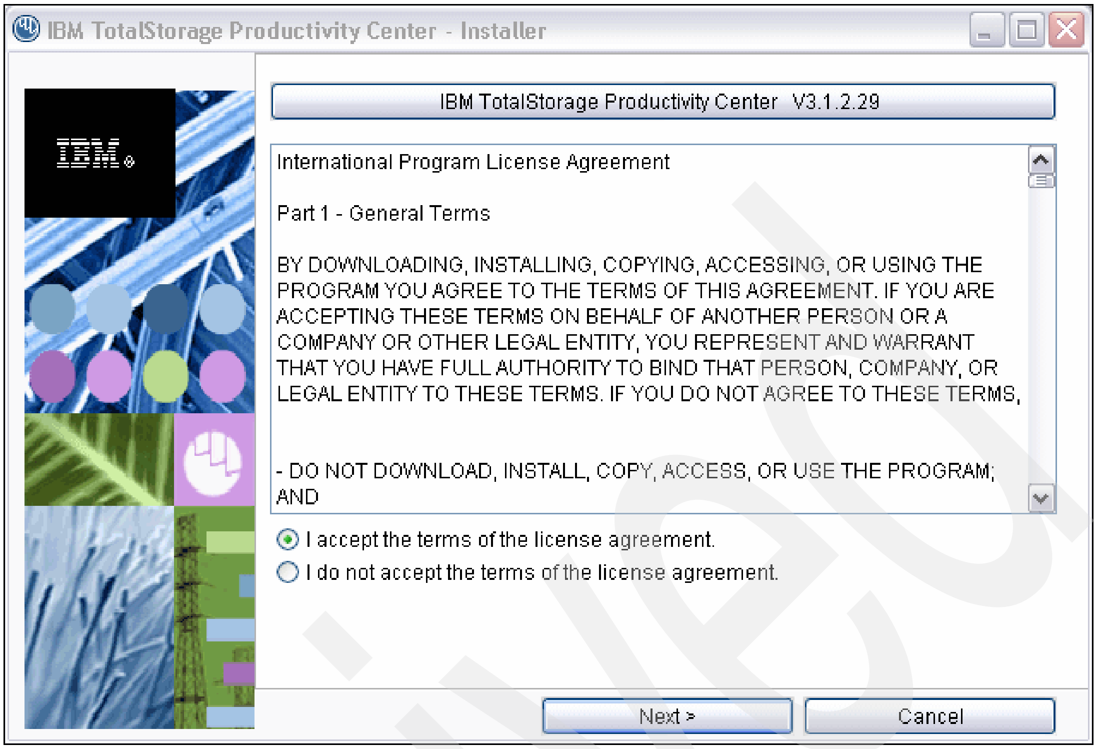
IBM TotalStorage Productivity Center - Installer
IBM TotalStorage Productivity Center V3.1.2.29

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, OR USING THE PROGRAM YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF ANOTHER PERSON OR A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON, COMPANY, OR LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

- DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, OR USE THE PROGRAM; AND

● I accept the terms of the license agreement.

○ I do not accept the terms of the license agreement.

*Figure 6-30 Accept the license agreement*

4. Select custom installation (see Figure 6-31 on page 257). A typical installation installs the TPC server code with the agents and the CLI. The default location of the installation code appears in the TPC Installation Location. If you want to change the installation path, enter the preferred location. Click **Next** to continue.

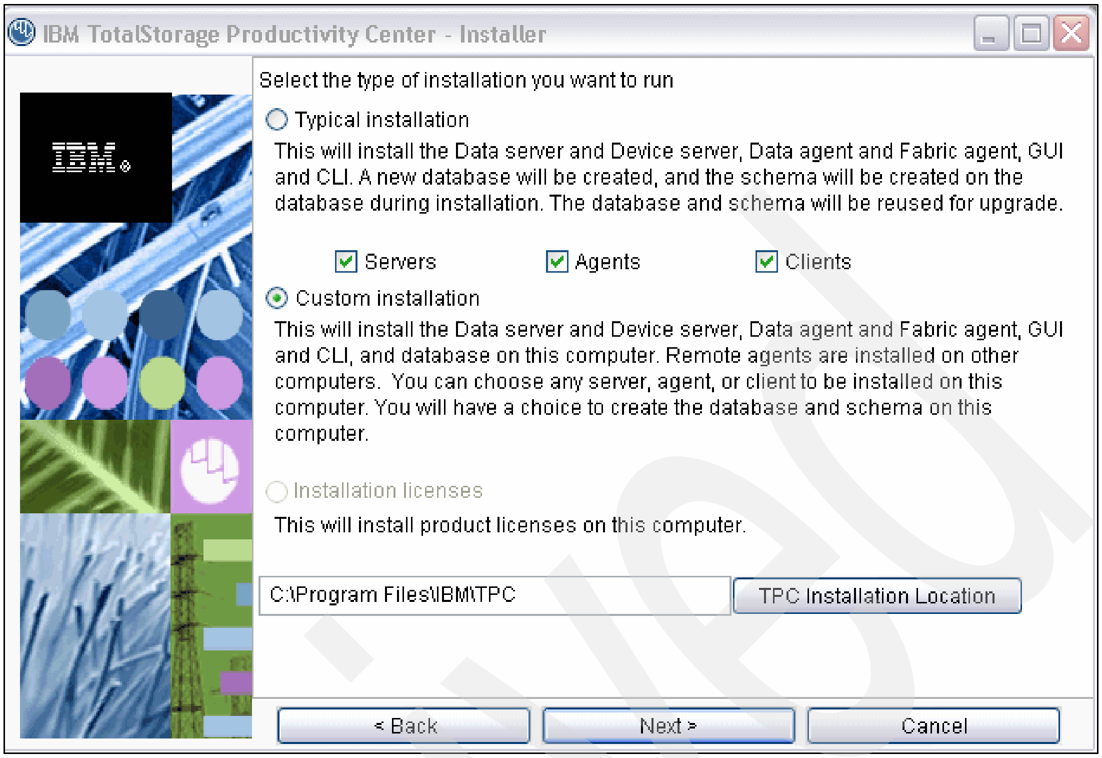

Select the type of installation you want to run

○ Typical installation

This will install the Data server and Device server, Data agent and Fabric agent, GUI and CLI. A new database will be created, and the schema will be created on the database during installation. The database and schema will be reused for upgrade.

☑ Servers ☑ Agents ☑ Clients

● Custom installation

This will install the Data server and Device server, Data agent and Fabric agent, GUI and CLI, and database on this computer. Remote agents are installed on other computers. You can choose any server, agent, or client to be installed on this computer. You will have a choice to create the database and schema on this computer.

○ Installation licenses

This will install product licenses on this computer.

C:\Program Files\IBM\TPC

*Figure 6-31 Select custom installation*

5. Select CLI (see Figure 6-32), ensure that all the other options are unchecked. Click **Next** to continue.

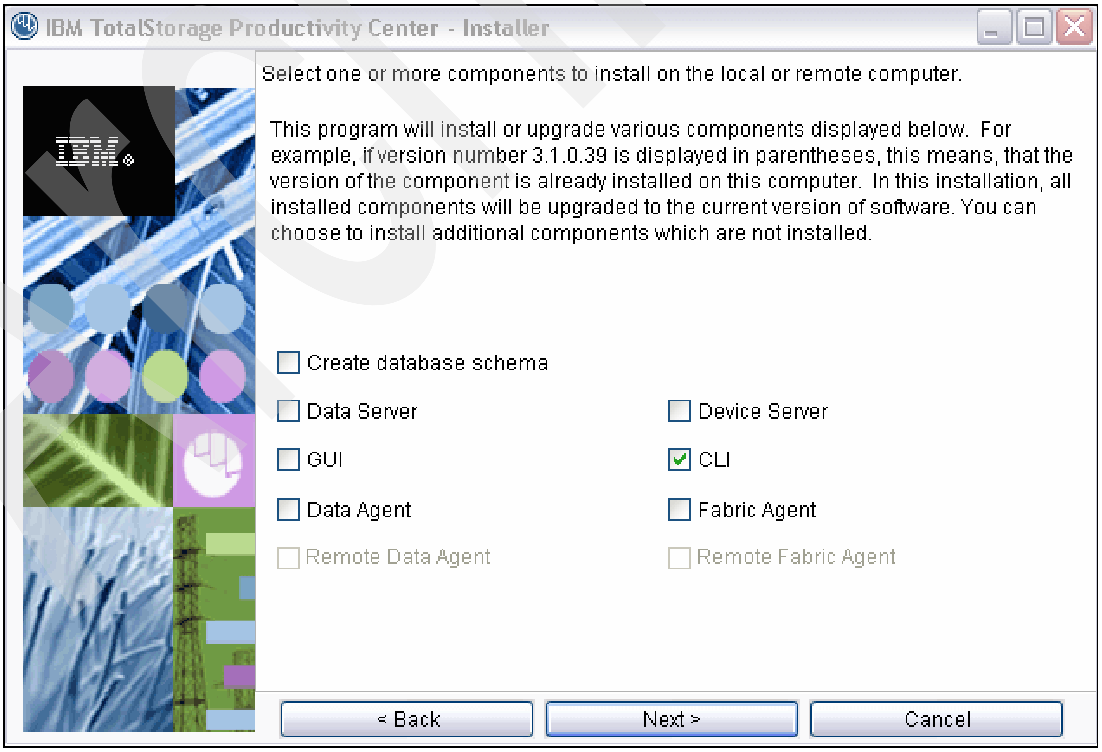

Select one or more components to install on the local or remote computer.

This program will install or upgrade various components displayed below. For example, if version number 3.1.0.39 is displayed in parentheses, this means, that the version of the component is already installed on this computer. In this installation, all installed components will be upgraded to the current version of software. You can choose to install additional components which are not installed.

☐ Create database schema

☐ Data Server ☐ Device Server

☐ GUI ☑ CLI

☐ Data Agent ☐ Fabric Agent

☐ Remote Data Agent ☐ Remote Fabric Agent

*Figure 6-32 Select CLI to install*

6. Enter the Device server (TPC server), DNS name, or IP address, and the port with which the Device server communicates. The default communication port is `9550`. You also need to enter the host authentication password, which was created upon installation of Agent Manager. The default password is `changeMe` (see Figure 6-33). Click **Next** to continue.

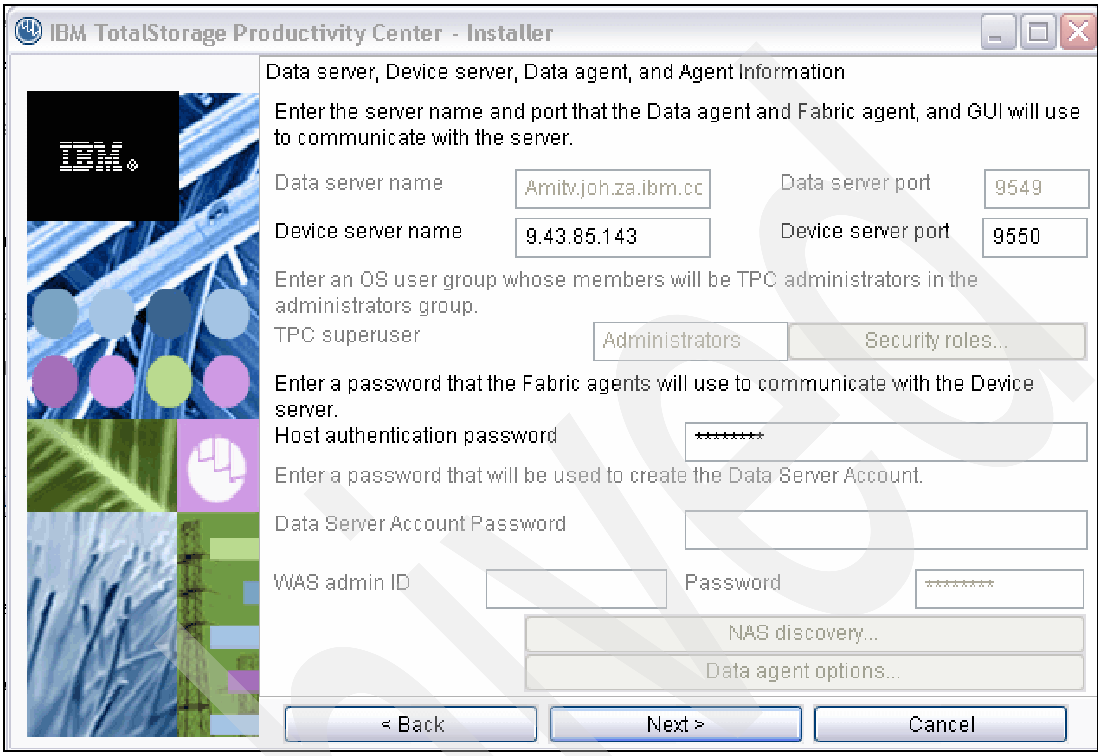

Data server, Device server, Data agent, and Agent Information

Enter the server name and port that the Data agent and Fabric agent, and GUI will use to communicate with the server.

Data server name: Amitv.joh.za.ibm.cc
Data server port: 9549
Device server name: 9.43.85.143
Device server port: 9550

Enter an OS user group whose members will be TPC administrators in the administrators group.

TPC superuser: Administrators

Enter a password that the Fabric agents will use to communicate with the Device server.

Host authentication password: *********

Enter a password that will be used to create the Data Server Account.

Data Server Account Password:

WAS admin ID:
Password: *********

*Figure 6-33 Enter Device server name, port address, and host authentication password*

7. TPC now gives you a summary of the information that you have entered in the previous windows as well as the size of the installation (see Figure 6-34 on page 259). Click **Install** to proceed.

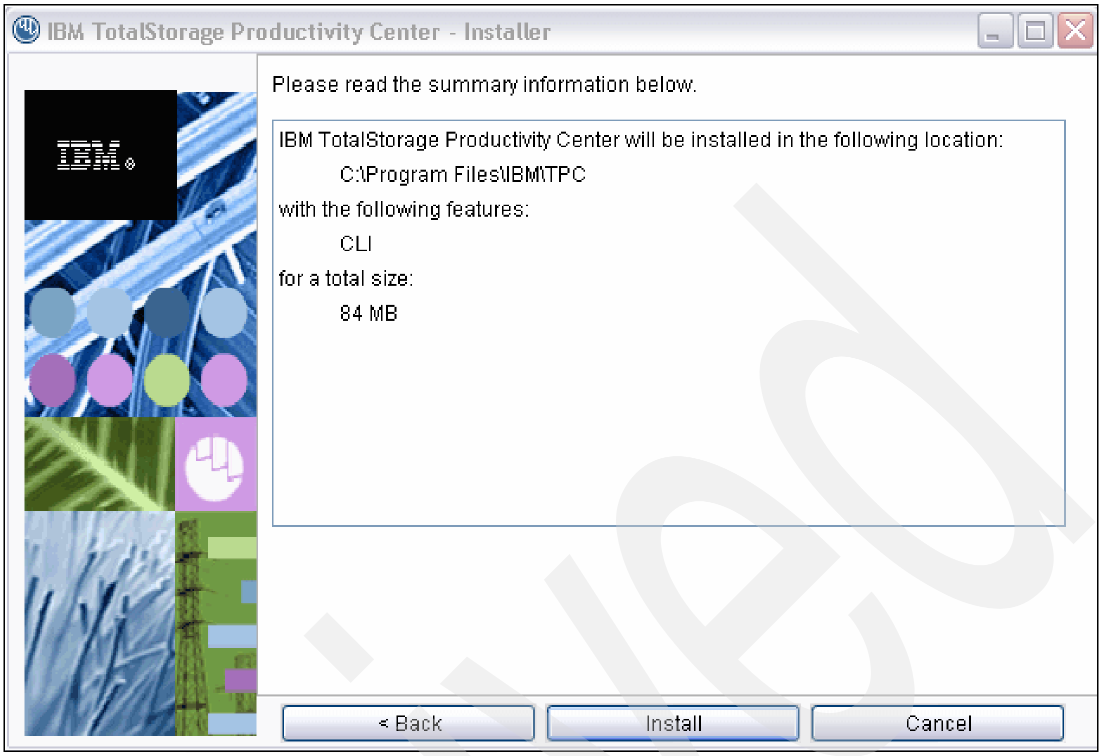

Please read the summary information below.

IBM TotalStorage Productivity Center will be installed in the following location:

C:\Program Files\IBM\TPC

with the following features:

CLI

for a total size:

84 MB

*Figure 6-34 Summary of installation*

8. The installation installs Java and the required components. A status bar indicates the progress of the installation (see Figure 6-35).

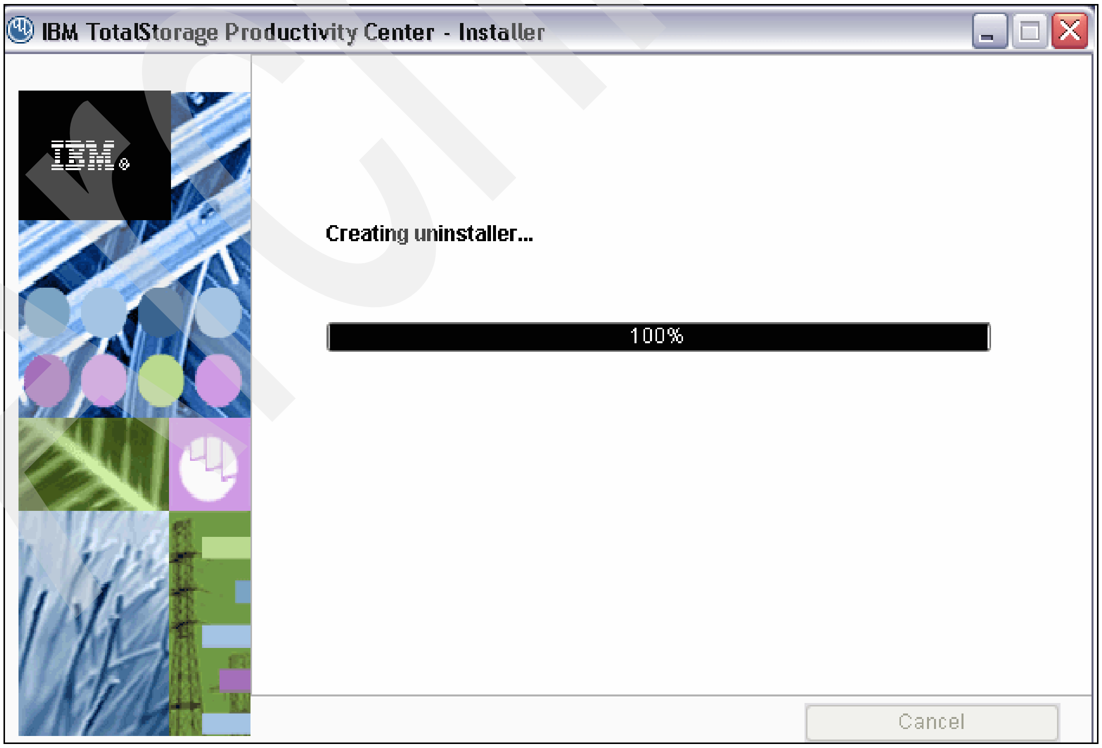

Creating uninstaller...

*Figure 6-35 Installer status*

9. After the installation has completed successfully, a summary displays as shown in Figure 6-36. Click **Finish**.

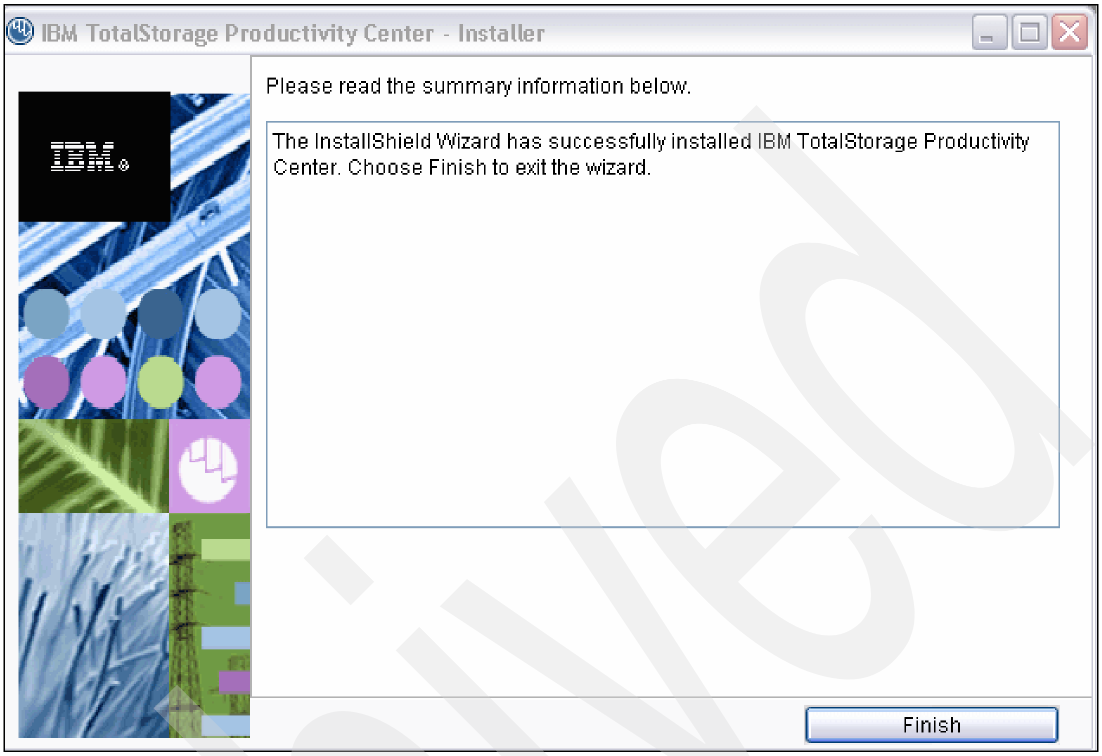

Please read the summary information below.

The InstallShield Wizard has successfully installed IBM TotalStorage Productivity Center. Choose Finish to exit the wizard.

*Figure 6-36 Installation summary*

10. You have successfully installed the TPCTOOL CLI on your workstation. You can now use the TOOL to access your TPC installation.

## 6.4.6 TPCTOOL configuration file

You use configuration files to configure the initial settings for computer operations. TPCTOOL allows you to utilize a configuration file to enter your settings to access your TPC Device server and to execute syntax:

- *Command* is a textual substitution of the command string with a defined alias.
- *Parameters* can be provided with default values that can be substituted.
- Command and Parameter can be *aliased* and stored in a configuration file on the client machine or server.
- The aliased commands are saved in the command configuration file.

The TPCCLI.CONF file is a method to define alias commands that can be executed either in a script or through the interactive CLI interface.

**Note:** The TPCCLI.CONF file is not provided by default.

Create the file and save it in the following directories:

- Windows:

  **C:\Program Files\IBM\TPC\cli\libs**

- UNIX or Linux:

  **/<usr or opt>/IBM/TPC/cli/libs**

The aliased commands are written in ASCII and are line-oriented. The lines are terminated by a new line. The TPC administrator or SAN administrator needs to create the configuration file. This configuration file needs to be maintained on a regular basis, depending on the changing environment and the function that is required.

### Build the configuration file

The steps to build the configuration file are:

1. You first need to define your basic syntax, which includes your connection options (the user ID, password, and url) to connect to your TPC Device server as shown in Example 6-4.

> **Tip:** We recommend that the password is defined in interactive mode so that it is encryted in the configuration file. See encrypt under 6.4.7, “General commands” on page 263.

*Example 6-4 Example of a configuration file 1*

```
#TPC Cli configuration property file DO NOT EDIT
# Sun Sept 17 11:00am
user=amit
password=#@$sheetal143
url=localhost:9550
```

2. After entering your connection options, you need to define the devices that you will manage and from which you will collect performance information (see Example 6-5). These are the devices that you monitor with TPC. You can obtain a list of devices by using the **`lsdev`** command.

*Example 6-5 Example of a configuration file with device listed*

```
#TPC Cli configuration property file DO NOT EDIT
# Sun Sept 17 11:05am
user=amit
password=#@$sheetal143
url=155.237.5.67:9550
SVC1=10.5.3.56:0000020060c8990+0
Subsys1:=10.5.3.88:000000200656c9990+0
```

3. You can use the configuration file to script the entire command. By merely executing the alias that we defined, we can gather the information from the TPC database (see Example 6-6).

*Example 6-6 Example of a config with a command and its parameters*

```
#TPC Cli configuration property file DO NOT EDIT
# Sun Sept 17 11:15am
Lsperf1=lsdev -user amit -pwd #@$sheetal143 -url 155.237.5.67:9550 -perf
```

4. This configuration file can be edited to suit your environment. You can add commands into the configuration file, and therefore, run the defined alias. Using the generic, disk, fabric, and reporting commands available, you can generate performance reports. In

Example 6-7, we aliased a command to generate a volume performance graph for the SVC. This data will be extracted from the TPC database.

*Example 6-7 Example of a configuration file and the get report (getrpt) command*

```
#TPC CLi configuration property file DO NOT EDIT
#Sun Sept 17 11:05am

url=155.237.5.67.77:9550
SVC1=10.5.3.56:0000020060c8990+0
SVC2=10.5.3.57:0000020065e8290+0
user=amit
password=#@$sheetal143

asvc1vdisk=getrpt -url url -user user -pwd password -ctype 12
-columns 803,806,809,812,815,818,819,820,821,825,826,827,830,831,832,833 -level
sample -subsys SVC1 -fs ;

asvc2vdisk=getrpt -url url -user user -pwd password -ctype 12
-columns 803,806,809,812,815,818,819,820,821,825,826,827,830,831,832,833 -level
sample -subsys SVC2 -fs;
```

5. To utilize the alias defined in the TPCCLI.CONF file, you can execute the alias, which is named `asvc1vdisk` in Example 6-8, from the command line through script mode.

*Example 6-8 Executing the alias command through a script from the command line*

```
Cd:C:\Program File\IBM\TPC\cli
tpctool asvc1vdisk
```

6. This command returns an output similar to Example 6-9.

*Example 6-9 Output of the executed script*

```
Timestamp;Device;Component;803;806;809;812;815;818;819;820;821;825;826;827;830;
831;832;833
=====================================================================
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x977vd16;262.3137;0.0000;262.3137;-;-;-;0
.5123;0.0000;0.5123;2.0000;0.0000;2.0000;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x916vd18;0.0011;0.0067;0.0078;-;-;-;0.000
0;0.0000;0.0000;4.0000;4.0000;4.0000;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x982vd4;0.1375;0.0887;0.2262;-;-;-;0.0003
;0.0002;0.0004;1.9758;2.0000;1.9853;-;-;-;-
```

7. Example 6-10 specifies a **-start** variable and a **-duration** variable to provide a defined period of collected data.

*Example 6-10 Output of the executed script with a timestamp*

```
tpctool asvc1vdisk -start 2006:08:01:08:00:00 -duration 27000
```

The `asvc1vdisk` is the script that we created in the TPCCIL.CONF file.

The command in Example 6-10 on page 262 returns an output similar to Example 6-11.

*Example 6-11   Output of the executed script with data and a timestamp*

```
Timestamp;Device;Component;803;806;809;812;815;818;819;820;821;825;826;827;830;
831;832;833
===========================================================================
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x977vd16;262.3137;0.0000;262.3137;-;-;-;0
.5123;0.0000;0.5123;2.0000;0.0000;2.0000;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x916vd18;0.0011;0.0067;0.0078;-;-;-;0.000
0;0.0000;0.0000;4.0000;4.0000;4.0000;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x982vd4;0.1375;0.0887;0.2262;-;-;-;0.0003
;0.0002;0.0004;1.9758;2.0000;1.9853;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x982vd8;0.1685;0.0011;0.1696;-;-;-;0.0004
;0.0000;0.0004;2.4408;2.0000;2.4379;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x977vd15;351.1807;0.0000;351.1807;-;-;-;0
.6859;0.0000;0.6859;2.0000;0.0000;2.0000;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x982vd2;0.0222;0.0089;0.0310;-;-;-;0.0000
;0.0000;0.0001;1.8500;2.0000;1.8929;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x916vd4;0.0000;0.3814;0.3814;-;-;-;0.0000
;0.0016;0.0016;0.0000;4.3605;4.3605;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;MCS_XPvd10;1.5421;12.9035;14.4457;-;-;-;0
.0253;0.1626;0.1879;16.7850;12.9024;13.3169;-;-;-;-
2006.07.31:16:13:55;SVC-2145-SVC1-IBM;x9e9vd1;0.0000;3.2683;3.2683;-;-;-;0.0000
;0.0131;0.0131;0.0000;4.1072;4.1072;-;-;-;-
```

The difference between Example 6-10 on page 262 and Example 6-11 is that in Example 6-10 on page 262, where the start and duration values are not provided, the output only returns the last data collected, compared to in Example 6-11, where the start and duration values were defined, all of the data for the time period was provided by the output.

## 6.4.7 General commands

In this section, we describe a few general commands that you can execute by using TPCTOOL.

### -help

The following command lists all the commands that you can execute with TPCTOOL.

**`-help | h | -?`**

Entering a question mark (**?**) or **help** or **h** at the CLI prompt displays a list of commands available for each command mode. You can also get a list of keywords and arguments associated with any command by using the context-sensitive help feature.

### -ver

To list the version of TPC, use the **-ver** command as shown in Figure 6-37 on page 264.

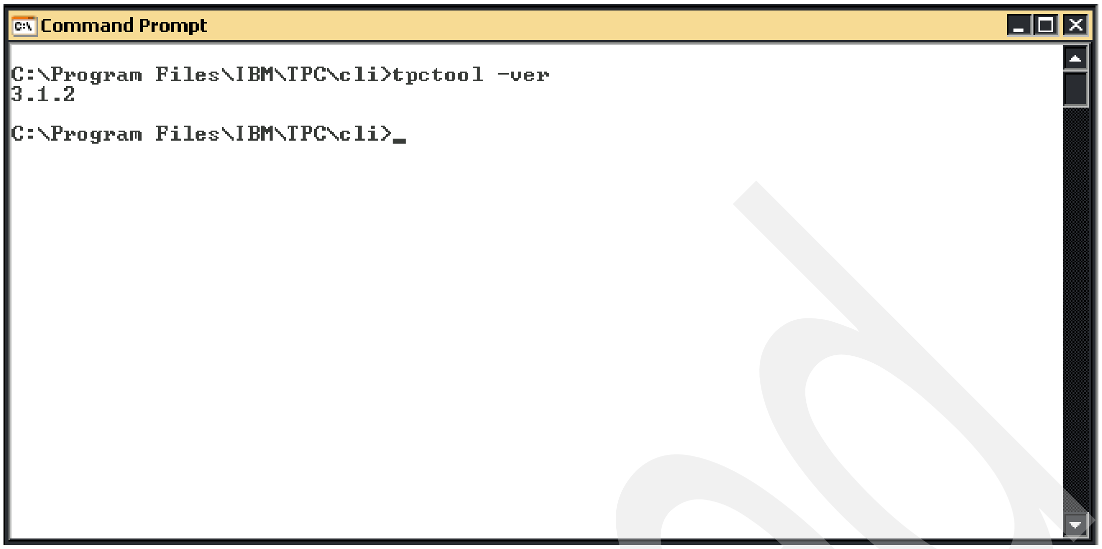

```
C:\Program Files\IBM\TPC\cli>tpctool -ver
3.1.2

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-37 tpctool ver command output*

## -lsdev

The **lsdev** command lists information about the storage subsystem, fabrics, and switches that are defined in TPC. You must be a disk administrator (TPC role-based administration) to issue this command. Role-based administration allows a single user sign-on.

The command is:

```
lsdev -user myuser -pwd mypass -url myurl
```

Figure 6-38 shows the command.

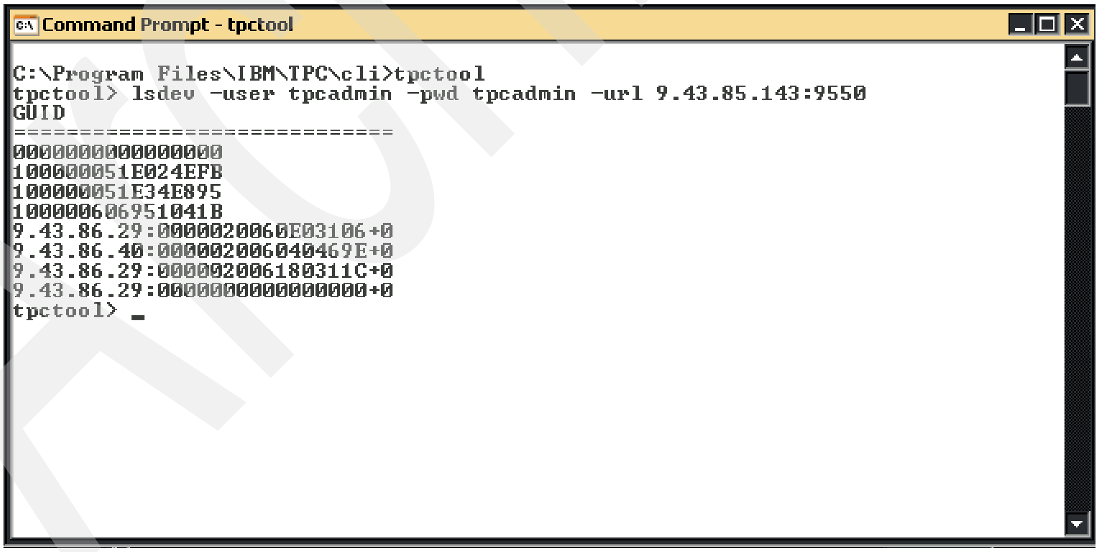

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lsdev -user tpcadmin -pwd tpcadmin -url 9.43.85.143:9550
GUID
=============================
0000000000000000
100000051E024EFB
100000051E34E895
100000606951041B
9.43.86.29:0000020060E03106+0
9.43.86.40:000002006040469E+0
9.43.86.29:000002006180311C+0
9.43.86.29:0000000000000000+0
tpctool> _
```

*Figure 6-38 tpctool lsdev command output*

## -encrypt

Use the **encrypt** command to generate an encrypted password for use in the configuration file. This command takes text from standard input and generates a 7-bit ASCII-equivalent character. Use encryption to ensure secrecy, particularly to verify the integrity and authenticity of the TPC environment. Figure 6-39 on page 265 shows the command.

The command is:

```
tpctool encrypt passwordfortpcserver
```

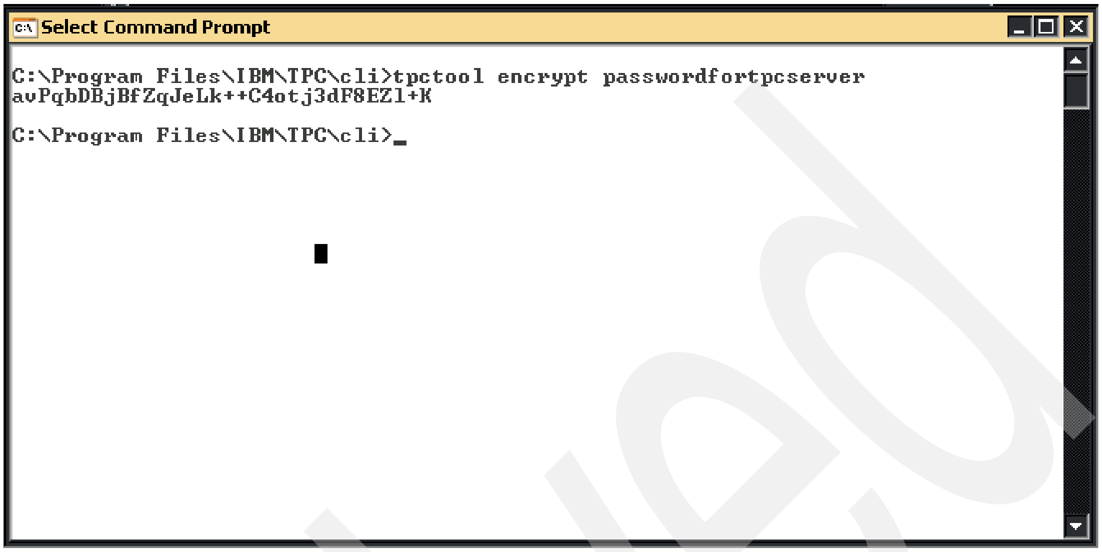

```
C:\Program Files\IBM\TPC\cli>tpctool encrypt passwordfortpcserver
avPqbDBjBfZqJeLk++C4otj3dF8EZl+K

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-39 tpctool encrypt command output*

**Restriction:** You can only use the encrypt function with the TPCTOOL in IBM TotalStorage Productivity Center Version 3.3.

For more information about TPCTOOL, refer to *IBM Totalstorage Productivity Center Command-Line Interface Reference,* GC32-1777.

## 6.4.8 Disk commands

In this section, we describe a few Disk commands that you can execute by using the TPCTOOL. You must be a disk administrator (TPC role-based authority) to issue these commands.

### -lsvol

You use the **lsvol** command to list all the volumes on a subsystem, a specific volume or volumes, or volumes on a specific array. You must use the device GUID variable to output a list of volumes.

The command is:

```
lsvol -user myuser -pwd mypass -url myurl -dev 600A0B80000CBDAD00000000440F3BA5+15
```

Figure 6-40 on page 266 shows the command.

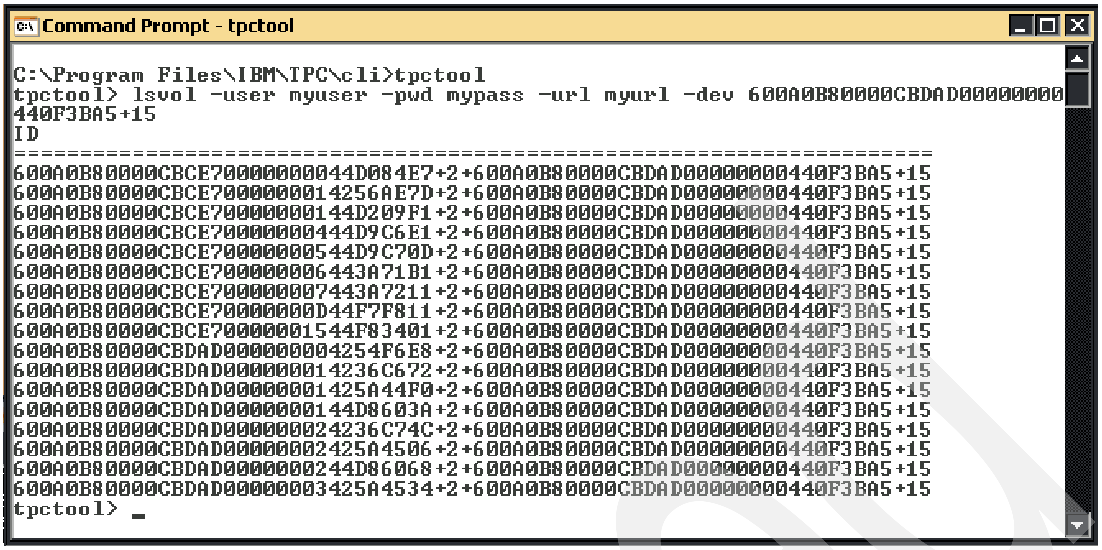

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lsvol -user myuser -pwd mypass -url myurl -dev 600A0B80000CBDAD00000000
440F3BA5+15
ID
======================================================================
600A0B80000CBCE70000000044D084E7+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE7000000014256AE7D+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE70000000144D209F1+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE70000000444D9C6E1+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE70000000544D9C70D+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE700000006443A71B1+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE700000007443A7211+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE70000000D44F7F811+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBCE70000001544F83401+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD000000004254F6E8+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD000000014236C672+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD00000001425A44F0+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD0000000144D8603A+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD000000024236C74C+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD00000002425A4506+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD0000000244D86068+2+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD00000003425A4534+2+600A0B80000CBDAD00000000440F3BA5+15
tpctool> _
```

*Figure 6-40 tpctool lsvol command output*

**Tip:** Use the `lsdev` command to determine a list of devices.

## -lsarray

You use the `lsarray` command to list information on an array per device.

The command is:

```
lsarray -user myuser -pwd mypass -url myurl -dev
600A0B80000CBDAD00000000440F3BA5+15
```

Figure 6-41 shows the command.

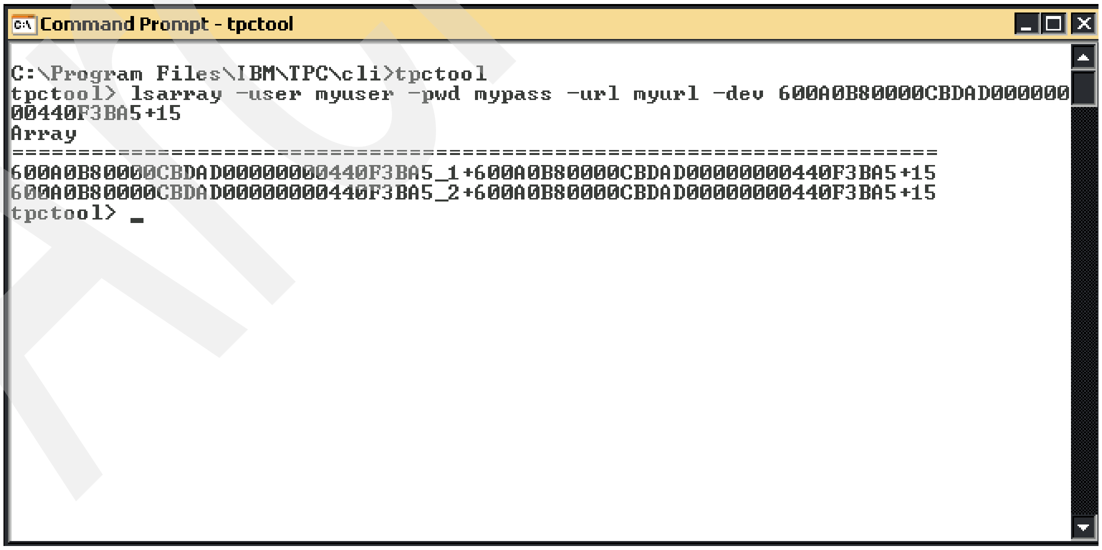

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lsarray -user myuser -pwd mypass -url myurl -dev 600A0B80000CBDAD000000
00440F3BA5+15
Array
======================================================================
600A0B80000CBDAD00000000440F3BA5_1+600A0B80000CBDAD00000000440F3BA5+15
600A0B80000CBDAD00000000440F3BA5_2+600A0B80000CBDAD00000000440F3BA5+15
tpctool> _
```

*Figure 6-41 tpctool lsarray command output*

**Tip:** Use the lsarray command to determine information about the arrays on specific subsystems.

## -mkvol

You use the **mkvol** command to create volumes. It is currently not supported on the SVC.

The command is:

```
mkvol -user myuser -pwd mypass -url myurl -array
600A0B80001744310000000041EFF538_11+600A0B80001744310000000041EFF538+15 -size 1
-count 1
```

Figure 6-42 shows the command.

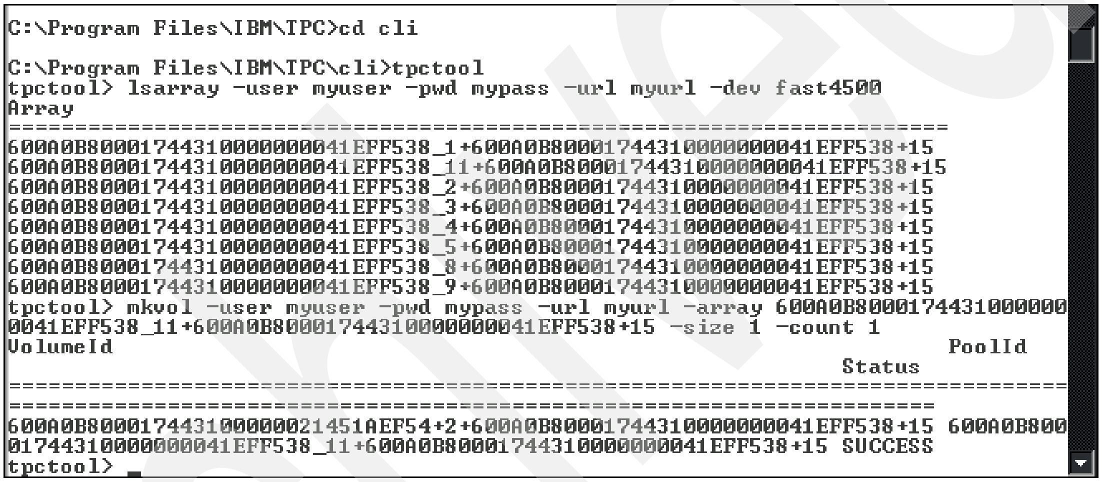

```
C:\Program Files\IBM\TPC>cd cli

C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lsarray -user myuser -pwd mypass -url myurl -dev fast4500
Array
=====================================================================
600A0B80001744310000000041EFF538_1+600A0B80001744310000000041EFF538+15
600A0B80001744310000000041EFF538_11+600A0B80001744310000000041EFF538+15
600A0B80001744310000000041EFF538_2+600A0B80001744310000000041EFF538+15
600A0B80001744310000000041EFF538_3+600A0B80001744310000000041EFF538+15
600A0B80001744310000000041EFF538_4+600A0B80001744310000000041EFF538+15
600A0B80001744310000000041EFF538_5+600A0B80001744310000000041EFF538+15
600A0B80001744310000000041EFF538_8+600A0B80001744310000000041EFF538+15
600A0B80001744310000000041EFF538_9+600A0B80001744310000000041EFF538+15
tpctool> mkvol -user myuser -pwd mypass -url myurl -array 600A0B8000174431000000
0041EFF538_11+600A0B80001744310000000041EFF538+15 -size 1 -count 1
VolumeId                                                               PoolId
                                                              Status
================================================================================
=====================================================================
600A0B80001744310000021451AEF54+2+600A0B80001744310000000041EFF538+15 600A0B800
0174431000000041EFF538_11+600A0B80001744310000000041EFF538+15 SUCCESS
tpctool> _
```

Figure 6-42 tpctool mkvol command output

## -rmvol

You use the **rmvol** command to remove volumes. It is currently not supported on the SVC.

The command is:

```
rmvol  -user myuser -pwd mypass -url myurl
600A0B80001744310000021451AEF54+2+600A0B80001744310000000041EFF538+15
```

Figure 6-43 on page 268 shows the command.

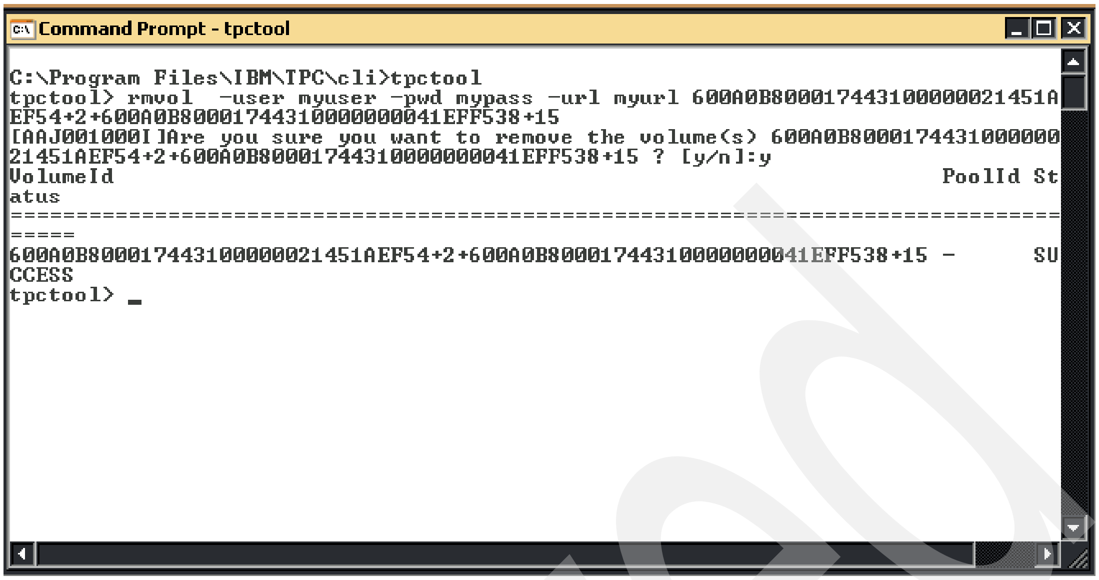

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> rmvol  -user myuser -pwd mypass -url myurl 600A0B80001744310000021451A
EF54+2+600A0B800017443100000000041EFF538+15
[AAJ001000I]Are you sure you want to remove the volume(s) 600A0B80001744310000000
21451AEF54+2+600A0B800017443100000000041EFF538+15 ? [y/n]:y
VolumeId                                                               PoolId St
atus
================================================================================
=====
600A0B80001744310000021451AEF54+2+600A0B800017443100000000041EFF538+15 -      SU
CCESS
tpctool> _
```

*Figure 6-43 tpctool rmvol command output*

**Tip:** Do not use the **-f** parameter, because you disable the confirmation command before the volumes are deleted.

## 6.4.9 Fabric commands

In this section, we discuss a few Fabric commands that you can execute by using TPCTOOL. You must be a Fabric administrator (TPC role-based authority) to issue these commands.

### -start

Use the **`start`** command to start a transaction. You must run certain commands with the start command to enable you access to the fabric.

Sample syntax is:

```
start -user myuser -pwd mypass -url myurl -fabric 100000051E34E895
```

### -lszs

You use the **`lszs`** command to list information about a zone set.

Sample syntax is:

```
lszs -user myuser -pwd mypass -url myurl -fabric myfabric
```

Figure 6-44 on page 269 shows the command.

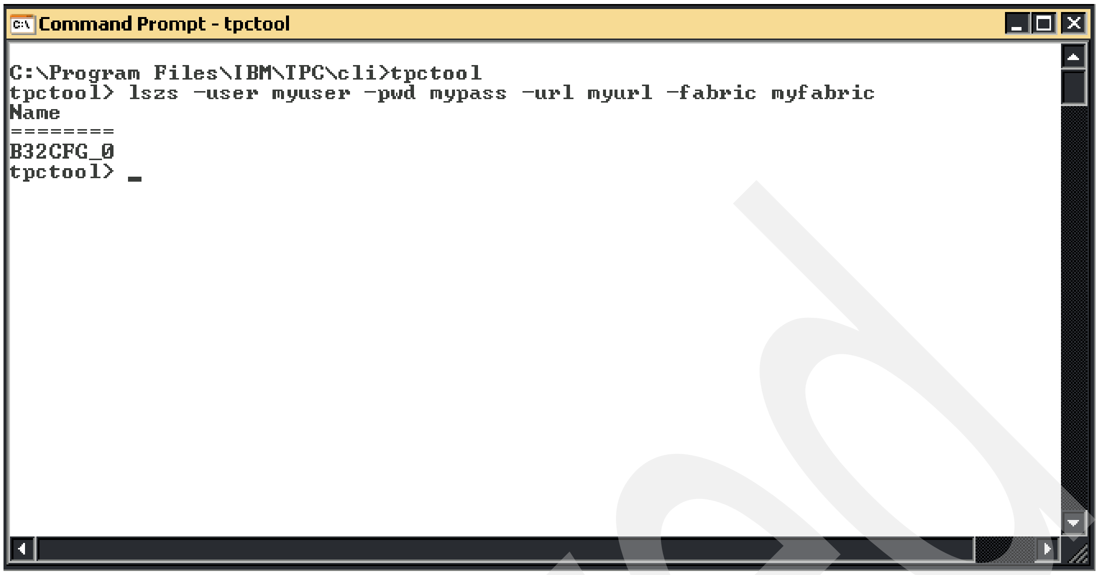

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lszs -user myuser -pwd mypass -url myurl -fabric myfabric
Name
========
B32CFG_0
tpctool> _
```

*Figure 6-44 tpctool lszs command output*

**Tip:** Use the **-active** parameter to display only the active zones. The default is to display both active and inactive zone sets.

### -lszone

You use the `lszone` command to list all the zones in a zone set.

The sample syntax is:

```
lszone -user myuser -pwd mypass -url myurl -fabric myfabric B32CFG_0
```

Figure 6-45 shows the command.

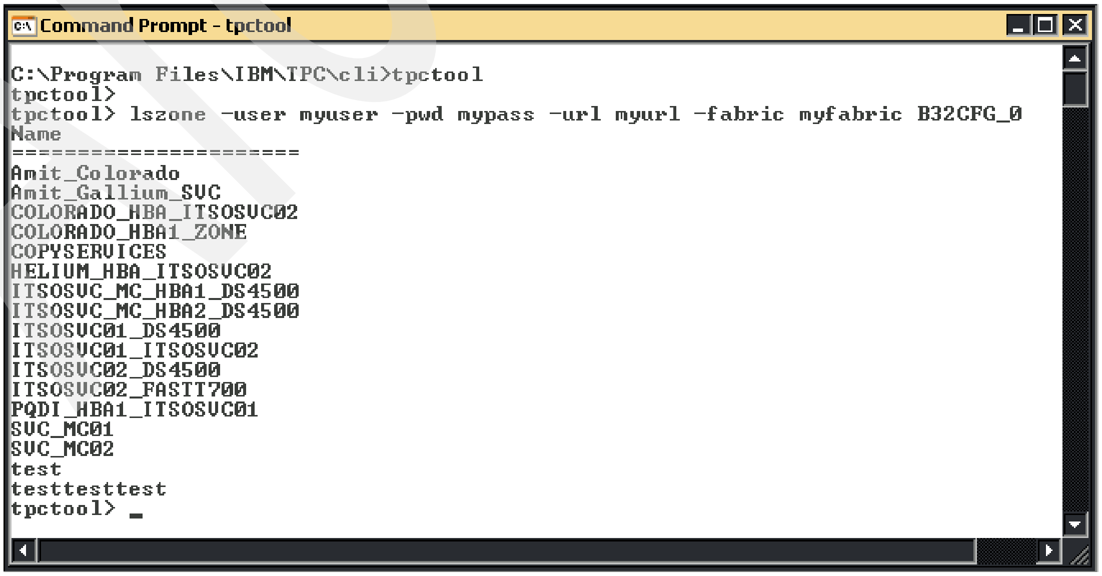

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool>
tpctool> lszone -user myuser -pwd mypass -url myurl -fabric myfabric B32CFG_0
Name
======================
Amit_Colorado
Amit_Gallium_SVC
COLORADO_HBA_ITSOSVC02
COLORADO_HBA1_ZONE
COPYSERVICES
HELIUM_HBA_ITSOSVC02
ITSOSVC_MC_HBA1_DS4500
ITSOSVC_MC_HBA2_DS4500
ITSOSVC01_DS4500
ITSOSVC01_ITSOSVC02
ITSOSVC02_DS4500
ITSOSVC02_FASTT700
PQDI_HBA1_ITSOSVC01
SVC_MC01
SVC_MC02
test
testtesttest
tpctool> _
```

*Figure 6-45 tpctool lszone command output*

### -ckzone

You use the **ckzone** command to verify that the fabric contains a zone. Figure 6-46 shows the command.

The sample syntax is:

```
ckzone -user myuser -pwd mypass -url myurl -fabric myfabric ITSOSVC_MC_HBA1_DS4500
```

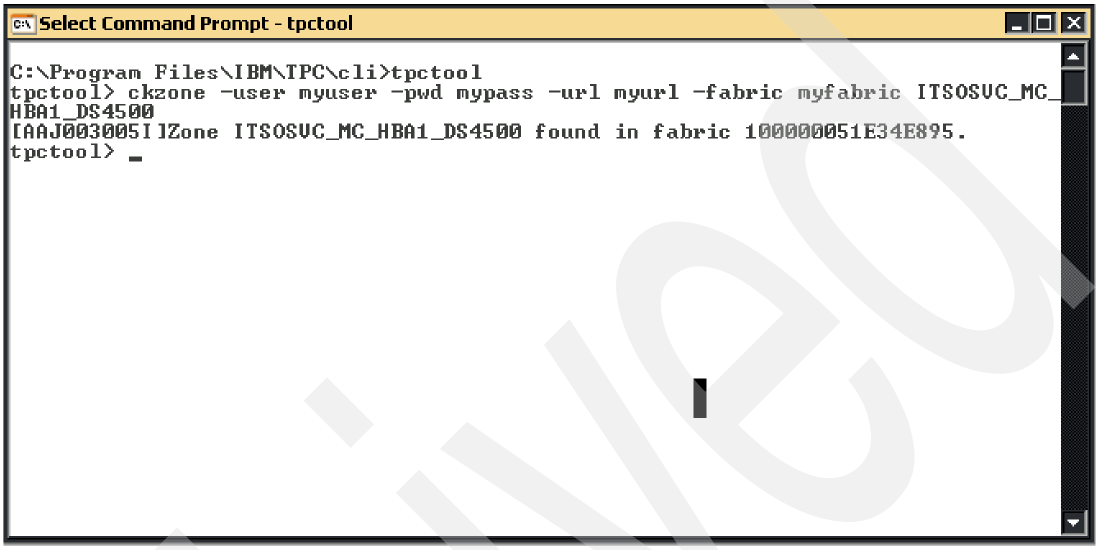

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> ckzone -user myuser -pwd mypass -url myurl -fabric myfabric ITSOSVC_MC_
HBA1_DS4500
[AAJ003005I]Zone ITSOSVC_MC_HBA1_DS4500 found in fabric 100000051E34E895.
tpctool> _
```

*Figure 6-46 tpctool ckzone command output*

## 6.4.10 Reporting commands

In this section, we describe all of the reporting commands that you can execute by using TPCTOOL. Sometimes, you only need to report the output for one or more data expressions. To produce a simple report, you can use the **getrpt** command. This command automatically loops over the dimensions of the data and formats the output in a default layout.

### -lstype

You use the **lstype** command to list the component types that are available on a fabric and a storage subsystem. Note the association between the **Name** and **Type** fields. This association is used in TPCTOOL reports.

Sample syntax is:

```
tpctool lstype
```

Figure 6-47 on page 271 shows the command.

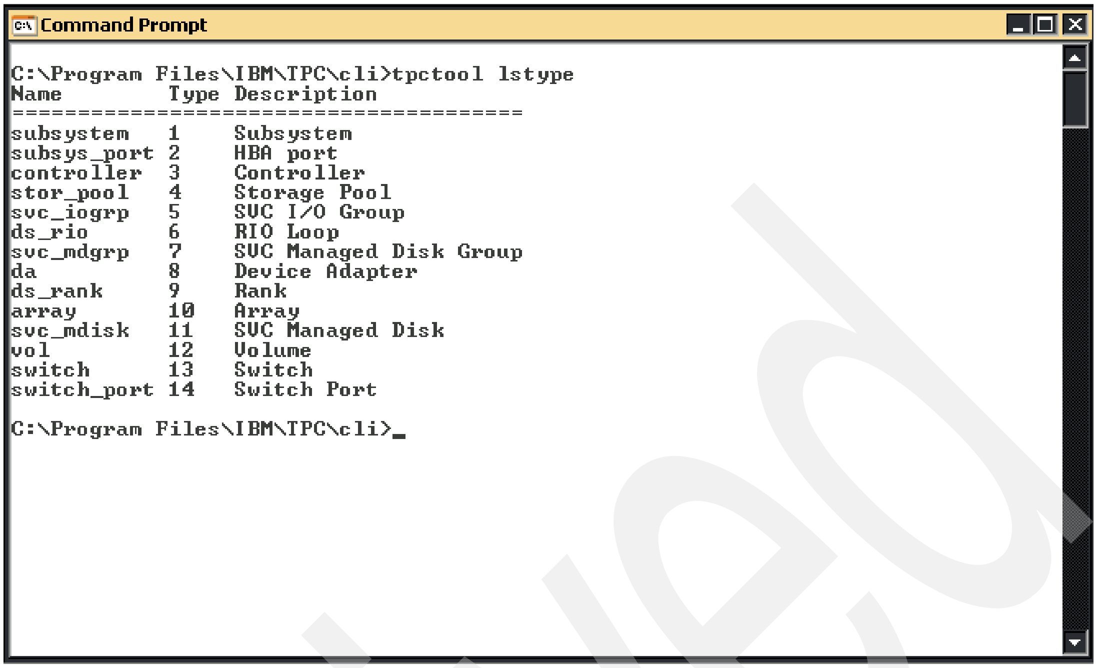

```
C:\Program Files\IBM\TPC\cli>tpctool lstype
Name        Type Description
========================================
subsystem   1    Subsystem
subsys_port 2    HBA port
controller  3    Controller
stor_pool   4    Storage Pool
svc_iogrp   5    SVC I/O Group
ds_rio      6    RIO Loop
svc_mdgrp   7    SVC Managed Disk Group
da          8    Device Adapter
ds_rank     9    Rank
array       10   Array
svc_mdisk   11   SVC Managed Disk
vol         12   Volume
switch      13   Switch
switch_port 14   Switch Port

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-47 tpctool lstype command output*

When using the available metrics for each **lstype**, you can use the name or the type, for example, you can use `subsystem` or `1` as a ctype variable. Example 6-12 shows a ctype using the metric name `subsystem`. Example 6-13 uses the type number `1` as a metric.

*Example 6-12 Using the lstype name as a ctype metric*

```
tpctool lsmetrics -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype 
subsystem -level sample
```

*Example 6-13 Using the lstype type as a ctype metric*

```
tpctool lsmetrics -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype 1 
-level sample
```

Using either option is entirely up to the TPC administrator. The results will be the same, which are shown in Figure 6-48 on page 272.

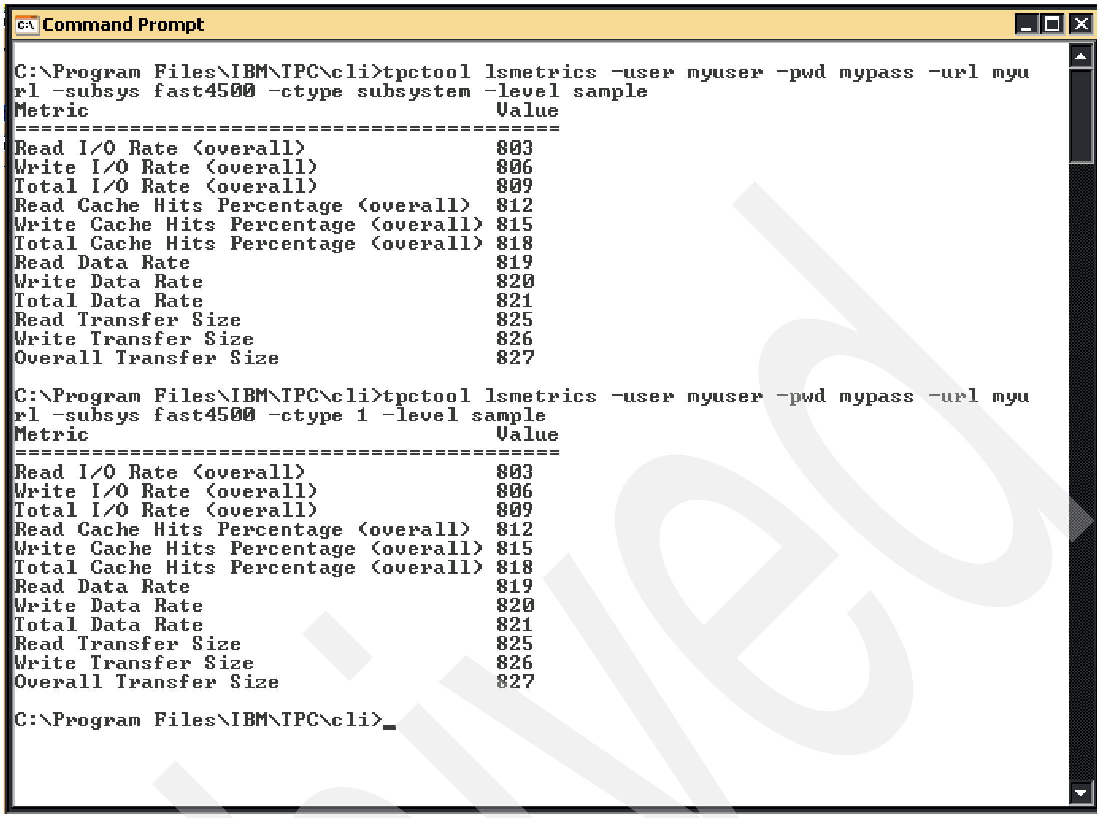

```
C:\Program Files\IBM\TPC\cli>tpctool lsmetrics -user myuser -pwd mypass -url myu
rl -subsys fast4500 -ctype subsystem -level sample
Metric                                  Value
==========================================
Read I/O Rate (overall)                 803
Write I/O Rate (overall)                806
Total I/O Rate (overall)                809
Read Cache Hits Percentage (overall)    812
Write Cache Hits Percentage (overall)   815
Total Cache Hits Percentage (overall)   818
Read Data Rate                          819
Write Data Rate                         820
Total Data Rate                         821
Read Transfer Size                      825
Write Transfer Size                     826
Overall Transfer Size                   827

C:\Program Files\IBM\TPC\cli>tpctool lsmetrics -user myuser -pwd mypass -url myu
rl -subsys fast4500 -ctype 1 -level sample
Metric                                  Value
==========================================
Read I/O Rate (overall)                 803
Write I/O Rate (overall)                806
Total I/O Rate (overall)                809
Read Cache Hits Percentage (overall)    812
Write Cache Hits Percentage (overall)   815
Total Cache Hits Percentage (overall)   818
Read Data Rate                          819
Write Data Rate                         820
Total Data Rate                         821
Read Transfer Size                      825
Write Transfer Size                     826
Overall Transfer Size                   827

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-48 tpctool lsmetrics using -cytpe subsystem or 1 results*

## -lsmetrics

You use the `lsmetrics` command to display a list of reports and metrics for devices and components. You must have disk and fabric operator authority to use this command.

Sample syntax is:

```
tpctool lsmetrics -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype
subsystem -level sample
```

Figure 6-49 on page 273 shows the command.

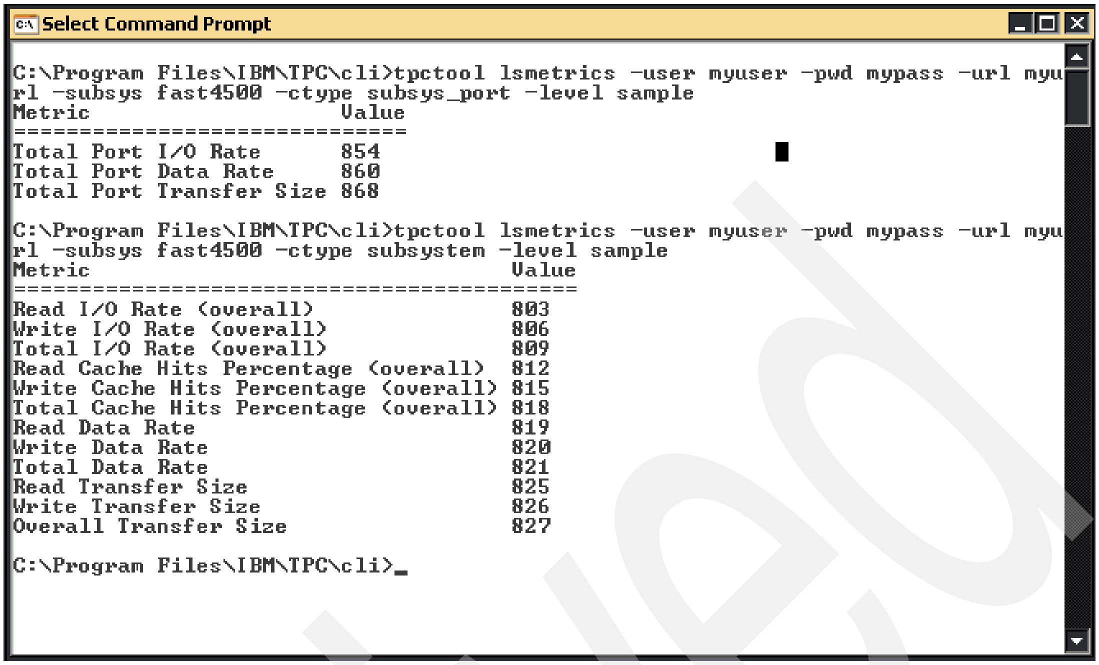

```
C:\Program Files\IBM\TPC\cli>tpctool lsmetrics -user myuser -pwd mypass -url myu
rl -subsys fast4500 -ctype subsys_port -level sample
Metric                   Value
==============================
Total Port I/O Rate      854
Total Port Data Rate     860
Total Port Transfer Size 868

C:\Program Files\IBM\TPC\cli>tpctool lsmetrics -user myuser -pwd mypass -url myu
rl -subsys fast4500 -ctype subsystem -level sample
Metric                                Value
===========================================
Read I/O Rate (overall)               803
Write I/O Rate (overall)              806
Total I/O Rate (overall)              809
Read Cache Hits Percentage (overall)  812
Write Cache Hits Percentage (overall) 815
Total Cache Hits Percentage (overall) 818
Read Data Rate                        819
Write Data Rate                       820
Total Data Rate                       821
Read Transfer Size                    825
Write Transfer Size                   826
Overall Transfer Size                 827

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-49 tpctool lsmetrics command output*

You use these metrics in your **`getrpt`** command to interpret your request and determine what data to extract from your TPC database (see Example 6-14). In this example, you only extract the Read I/O (803), Write I/O (806), and Total I/O Rates (809) from an IBM subsystem.

*Example 6-14 tpctool getrpt command with specific metrics for subsystem (ctype)*

```
getrpt -user myuser -pwd mypass -url myurl -subsys fast4500 -level sample
-ctype subsystem -columns 803,806,809
```

**Tip:** Use the **`lstype`** command for a list of ctype variables.

**Note:** Metrics might differ between device type and version level, for example, SVC 2.1 and SVC 4.1.

## -lstime

You use the **`lstime`** command to display a list of time ranges for which performance data is available. You must have disk and fabric operator authority to use this command as shown in Figure 6-50 on page 274.

Sample syntax is:

**`lstime -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype subsystem -level sample`**

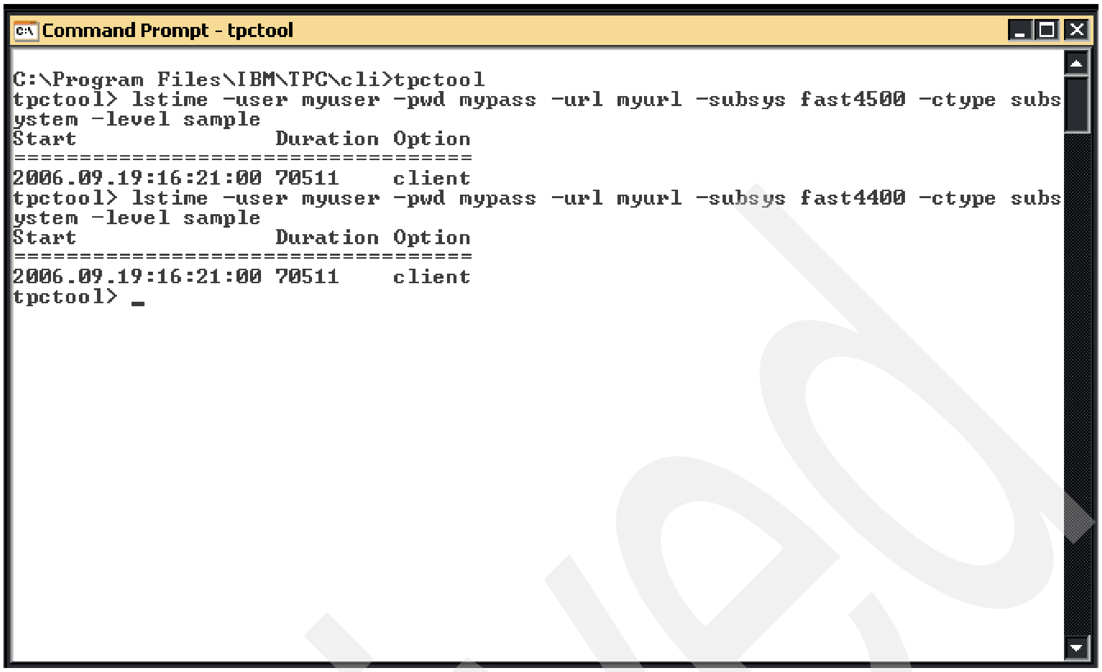

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lstime -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype subs
ystem -level sample
Start               Duration Option
===================================
2006.09.19:16:21:00 70511    client
tpctool> lstime -user myuser -pwd mypass -url myurl -subsys fast4400 -ctype subs
ystem -level sample
Start               Duration Option
===================================
2006.09.19:16:21:00 70511    client
tpctool> _
```

Figure 6-50 tpctool lstime command output

## -lscomp

You use the `lscomp` command to list all of the components for which performance data has been collected by TPCTOOL in the TPC database. Figure 6-51 shows the command.

Sample syntax is:

```
lscomp -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype subsystem -level
sample -start 2006.09.19:16:00:00 -duration 80000
```

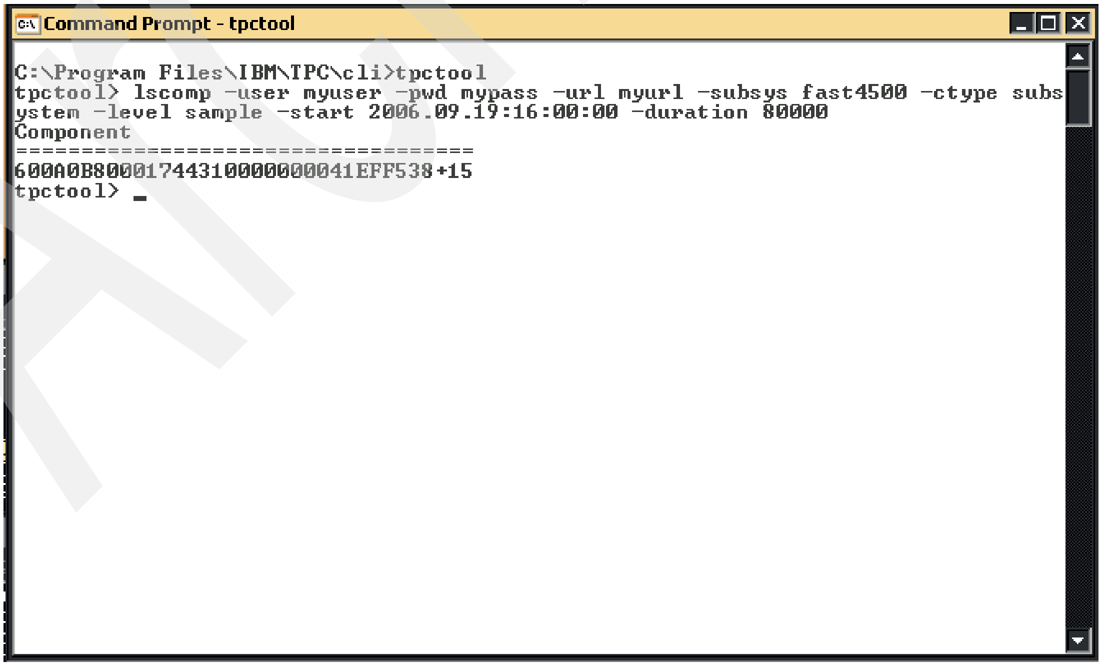

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lscomp -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype subs
ystem -level sample -start 2006.09.19:16:00:00 -duration 80000
Component
===================================
600A0B80001744310000000041EFF538+15
tpctool> _
```

Figure 6-51 tpctool lscomp command output

## -lscounter

Use the **lscounter** command to print a list of reports and metrics for the associated devices and components. You must have disk and fabric operator authority (TPC role-based authority) to use this command. Figure 6-52 shows the output of the command.

Sample syntax is:

```
lscounters -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype vol -level
sample
```

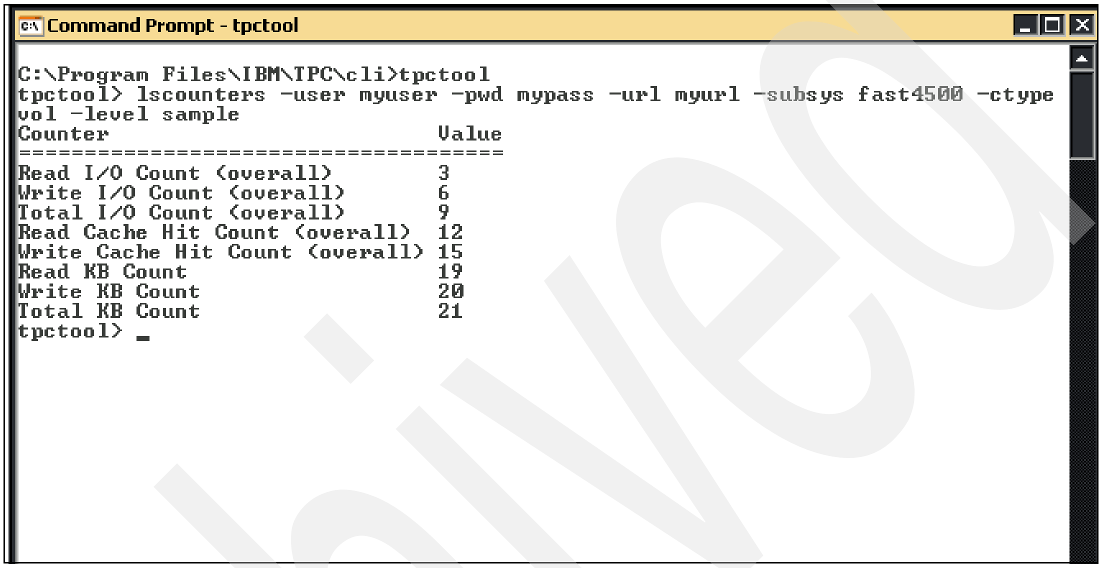

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> lscounters -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype
vol -level sample
Counter                         Value
=====================================
Read I/O Count (overall)        3
Write I/O Count (overall)       6
Total I/O Count (overall)       9
Read Cache Hit Count (overall)  12
Write Cache Hit Count (overall) 15
Read KB Count                   19
Write KB Count                  20
Total KB Count                  21
tpctool> _
```

*Figure 6-52 tpctool lscounters command output*

These values can also be determined from measurements taken from an existing system. For example, the basic metrics for I/O performance are throughput and response time.

## -getrpt

You use the **getrpt** command to display a performance metrics report. You must have disk and fabric operator authority (TPC role-based authority) to use this command. Figure 6-53 on page 276 shows the command.

Sample syntax is:

```
getrpt -user myuser -pwd mypass -url myurl -subsys fast4400 -level sample -ctype
subsystem -columns 803,806 -start 2006.09.19:16:21:00 -duration 80000 -fs ;
```

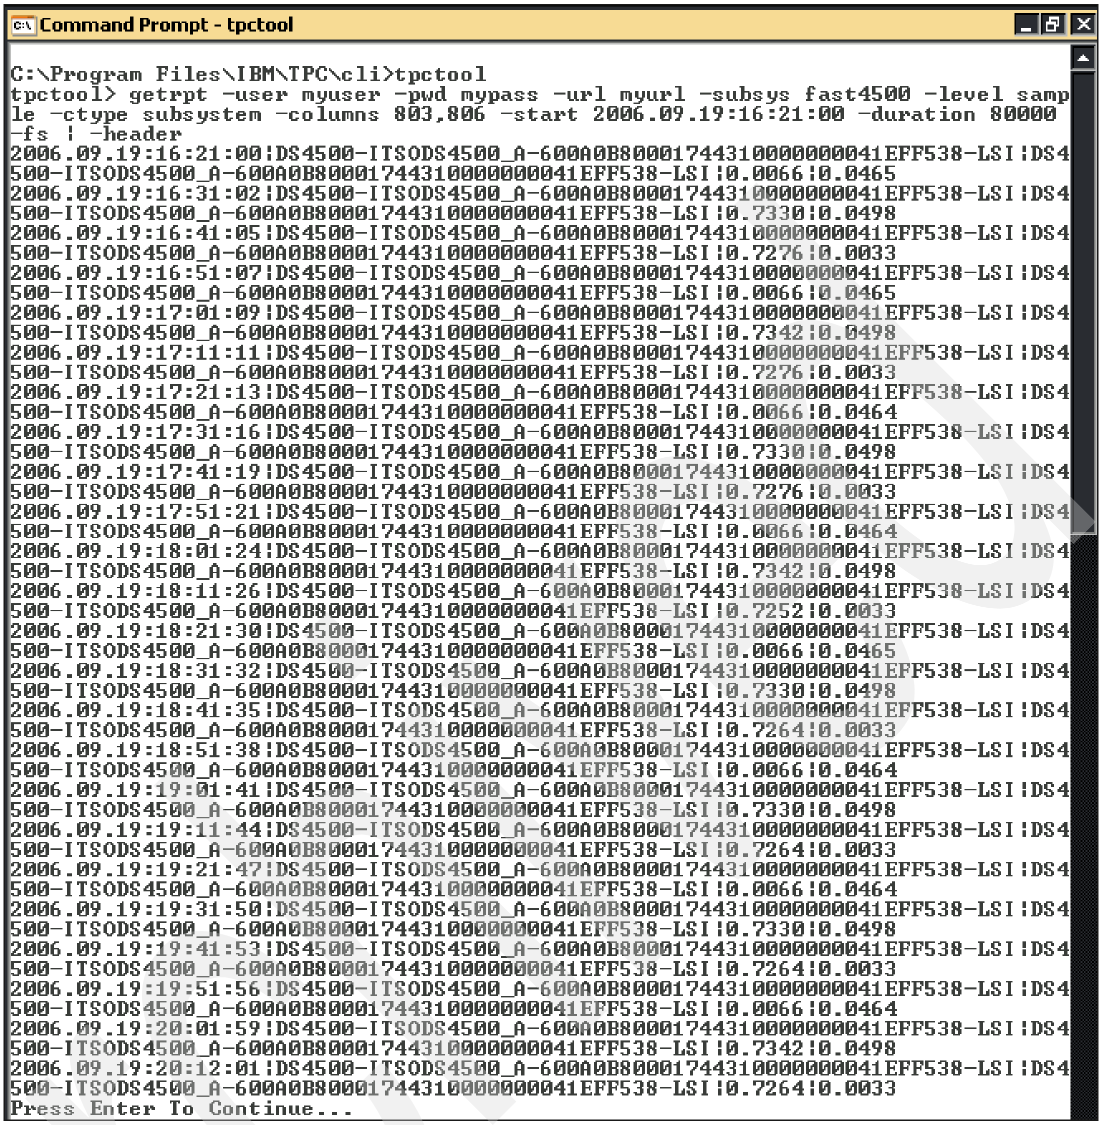

```
C:\Program Files\IBM\TPC\cli>tpctool
tpctool> getrpt -user myuser -pwd mypass -url myurl -subsys fast4500 -level samp
le -ctype subsystem -columns 803,806 -start 2006.09.19:16:21:00 -duration 80000
-fs ¦ -header
2006.09.19:16:21:00¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0465
2006.09.19:16:31:02¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7330¦0.0498
2006.09.19:16:41:05¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7276¦0.0033
2006.09.19:16:51:07¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0465
2006.09.19:17:01:09¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7342¦0.0498
2006.09.19:17:11:11¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7276¦0.0033
2006.09.19:17:21:13¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0464
2006.09.19:17:31:16¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7330¦0.0498
2006.09.19:17:41:19¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7276¦0.0033
2006.09.19:17:51:21¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0464
2006.09.19:18:01:24¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7342¦0.0498
2006.09.19:18:11:26¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7252¦0.0033
2006.09.19:18:21:30¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0465
2006.09.19:18:31:32¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7330¦0.0498
2006.09.19:18:41:35¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7264¦0.0033
2006.09.19:18:51:38¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0464
2006.09.19:19:01:41¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7330¦0.0498
2006.09.19:19:11:44¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7264¦0.0033
2006.09.19:19:21:47¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0464
2006.09.19:19:31:50¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7330¦0.0498
2006.09.19:19:41:53¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7264¦0.0033
2006.09.19:19:51:56¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.0066¦0.0464
2006.09.19:20:01:59¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7342¦0.0498
2006.09.19:20:12:01¦DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦DS4
500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI¦0.7264¦0.0033
Press Enter To Continue...
```

*Figure 6-53 tpctool getrpt command output*

In Figure 6-53, we extract subsystem performance data for the FAStT 4400 device for columns 803 (Read I/O Rate) and 806 (Write I/O Rate).

**Tips:**

Take advantage of these tips:

- Use **`lsmetrics`** or **`lscounters`** to obtain a list of columns that will appear in the report.
- You can use the **`-fs`** command with a character that separates the fields in the output.
- You can use the **`-header`** command to suppress (remove) the column headings. By default, the reports include the column headings.

By combining these various types of syntax in a command, you can create custom reports with special formatting. For example, you can create a report with a different variable in each column or with a different variable in each row.

## 6.4.11 Custom reports

To create a customized report, you combine the reporting commands, functions, and options. Although you use a combination of various syntax in commands, you can experiment with most of the syntax individually.

Here are examples of how you can benefit from the customization of reports:

- The facility allows you to produce reports whenever you need them.
- You choose the plans and dates on the report for which you want data displayed.
- You have access to many preformatted reports with the ability to create your own custom reports.
- You can save a template of your report for future use.
- You can sort and display the report any way that you see fit.
- You can schedule reports to run automatically.
- You receive an e-mail when your reports are ready.

## 6.4.12 TPCTOOL improvements

### CLI changes in TPC V3.1.2

In this section, we describe the changes that have been made to the CLI from TPC V3.1.1 to TPC V3.1.2:

- The `lsdev` command was modified to add the **Label** column. The user is able to display the user-provided name.
- The `lsarray` command was modified to add the **Label** column. The user will be able to display the name of the storage pool.
- The `lsvol` command was modified to add a **Label** column and a **Format** column to identify the volume and the size in GBs.
- The `lstype` command was modified to add a **Description** column to display the name and the format of the storage volume.
- The `unassignvol` was modified to show the worldwide port name (WWPN).
- The tpctool return code when using `cygwin` was fixed.
- The `getrpt` command was modified in order to use the continuation token and to retrieve all the data chunks and to identify components.
- The `lstime` command was modified to display the **Duration** column and the **Options** column.
- The `ver` command was modified to return the client version instead of help.
- A subsystem ID was added in the report commands.
- The CLI reports work correctly in distinct time zones.

### CLI changes in TPC Version 3.1.3

In this section, we discuss the changes that have been made to the CLI from TPC V3.1.2 to TPC V3.1.3:

- The ability to display SVC node performance information has been added. This is accomplished by using the new component type **svc_node**.
- The strings that contain the separator character are quoted in order to correctly generate a comma delimited output format.
- If you do not specify the quotation mark character (‘), these strings are enclosed by double quotation marks (“ ”).
- The **-grouping** option has been added for the **`lsarray`**, **`lsvol`**, and **`getrpt`** commands. The digit grouping is enabled (for example, numeric values such as `12000` are displayed as 12,000). The grouping character (for example, a comma for US numbers) is determined by the location currently in effect. By default, the digit grouping is disabled.
- The **-colname** option has been added for the **`getrpt`** command to display a column name instead of a number.

## 6.4.13 Producing custom performance reports with TPCTOOL

In this section, we use these CLI commands to generate custom performance reports that you can use in a client environment.

The TPC Graphical User Interface (GUI) only allows you to present a report with a single performance metric (for example, `Read I/O`, `Read Transfer Rate`, or `Read Cache Hit`). In order to utilize the data gathered by TPC to its full potential and to fully understand the performance of the device, all of the grouped metrics need to be plotted on a graph from the host to the device (for example, `Read I/O`, `Write I/O`, and `Total I/O`) and not a single metric, which is the case with the TPC GUI.

The CLI/TPCTOOL method has the capability to present multiple metrics. In this method, you use the interactive user interface and scripted files to extract the data from the DB2 database across multiple metrics. This method requires customization in terms of importing data into EXCEL or a similar spreadsheet and plotting graphs against this extracted data. You can save the customized scripts and graphs as templates for later use to avoid the efforts of recreating the graphs and scripts.

### Subsystem performance report

We will produce a paradigm for performance analysis of an IBM FAStT Storage Subsystem. The paradigm will represent a method to analyze the data collected by TPC for a client’s environment. We need to utilize the tpccli.conf file to script the commands and parameters.

The steps are:

1. First, set the basic connection options within the tpccli.conf file, such as user name, passwords, and URLs, as shown in Example 6-15.

*Example 6-15 tpccli.conf with connection options*

```
#TPC Cli configuration property file DO NOT EDIT
# Sun Sept 21 11:05am
myuser=amit
mypass=#@$sheetal143
myurl=9.43.85.143:9550
```

2. Then, you need to add the subsystems for which you want to extract TPC data. Remember to use the **lsdev** command to verify which devices are available as shown in Example 6-16.

*Example 6-16 tpctool listing all available devices by using command lsdev*

```
lsdev -user myuser -pwd mypass -url myurl

Output

GUID
===================================
100000051E34E895
9.43.86.40:000002006040469E+0
9.43.86.29:000002006180311C+0
600A0B80001744310000000041EFF538+15
600A0B80000CBDAD00000000440F3BA5+15
```

3. Copy the entire device string, as displayed in the output in Example 6-16, into the tpccli.conf file. It will now look like Example 6-17.

*Example 6-17 tpccli.conf with devices*

```
#TPC Cli configuration property file DO NOT EDIT
# Sun Sept 21 11:05am
myuser=amit
mypass=#@$sheetal143
myurl=9.43.85.143:9550
fast4400=600A0B80000CBDAD00000000440F3BA5+15
fast4500=600A0B80001744310000000041EFF538+15
```

4. From now on, you will need to view the list of components that are recognized by TPC. Use the **lstype** command to verify this as shown in Example 6-18.

*Example 6-18 tpctool listing all available reports by using the lstype command*

```
tpctool lstype

Output

Name        Type Description
=======================================
subsystem   1    Subsystem
subsys_port 2    HBA port
controller  3    Controller
stor_pool   4    Storage Pool
svc_iogrp   5    SVC I/O Group
ds_rio      6    RIO Loop
svc_mdgrp   7    SVC Managed Disk Group
da          8    Device Adapter
ds_rank     9    Rank
array       10   Array
svc_mdisk   11   SVC Managed Disk
vol         12   Volume
switch      13   Switch
switch_port 14   Switch Port
```

5. The output will look like Figure 6-54.

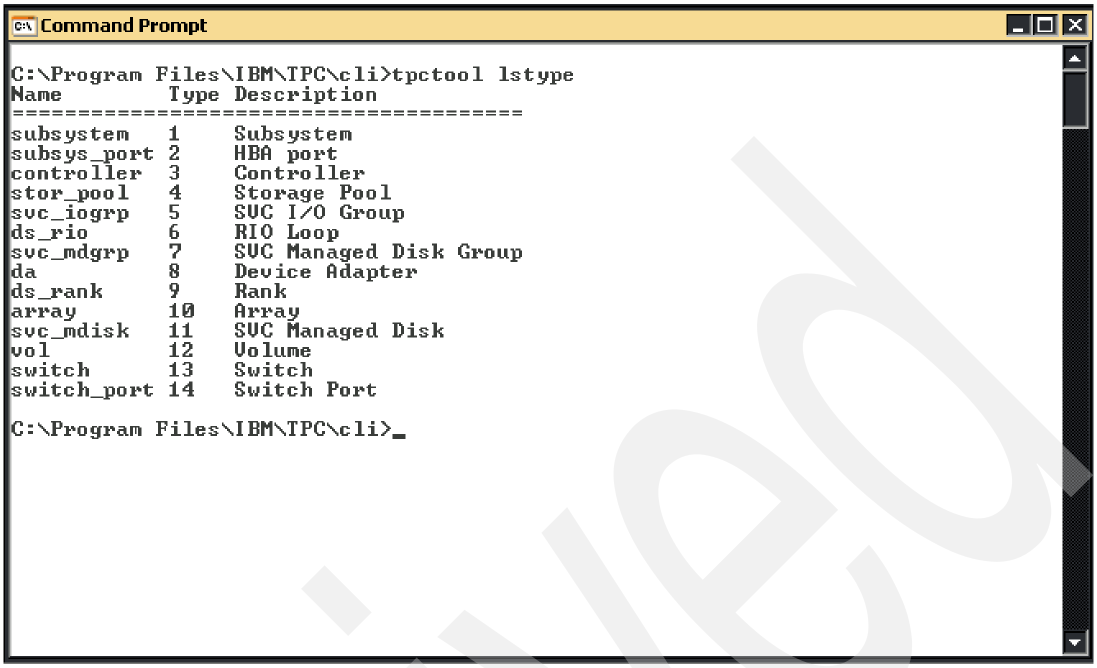

```
C:\Program Files\IBM\TPC\cli>tpctool lstype
Name        Type Description
=======================================
subsystem   1    Subsystem
subsys_port 2    HBA port
controller  3    Controller
stor_pool   4    Storage Pool
svc_iogrp   5    SVC I/O Group
ds_rio      6    RIO Loop
svc_mdgrp   7    SVC Managed Disk Group
da          8    Device Adapter
ds_rank     9    Rank
array       10   Array
svc_mdisk   11   SVC Managed Disk
vol         12   Volume
switch      13   Switch
switch_port 14   Switch Port

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-54 tpctool lstype command output list*

6. From the list, you can decide the type of component on which you want to report. For example, we use **ctype** subsystem (1), because we are reporting on an IBM subsystem (`FAStT`). We now need to determine the metrics. Use the **`lsmetrics`** command to display a list of subsystem metrics as shown in Example 6-19.

*Example 6-19 tpctool listing all available metrics for ctype subsystem (1) by using lsmetrics*

```
tpctool lsmetrics -user myuser -pwd mypass -url myurl -subsys fast4500 -ctype
subsystem -level sample
```

7. The output looks like Figure 6-55.

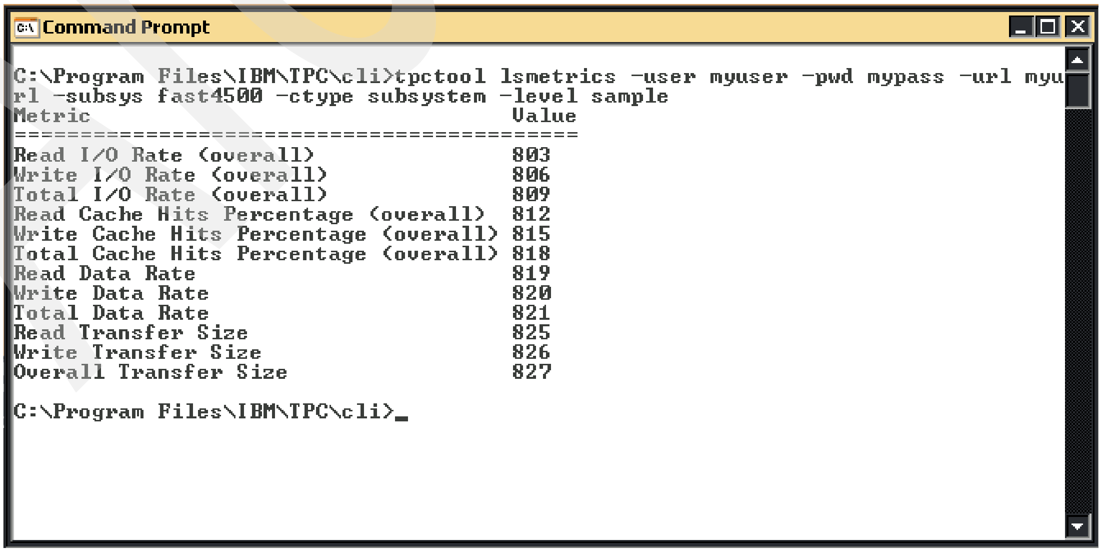

```
C:\Program Files\IBM\TPC\cli>tpctool lsmetrics -user myuser -pwd mypass -url myu
rl -subsys fast4500 -ctype subsystem -level sample
Metric                                Value
=========================================
Read I/O Rate (overall)               803
Write I/O Rate (overall)              806
Total I/O Rate (overall)              809
Read Cache Hits Percentage (overall)  812
Write Cache Hits Percentage (overall) 815
Total Cache Hits Percentage (overall) 818
Read Data Rate                        819
Write Data Rate                       820
Total Data Rate                       821
Read Transfer Size                    825
Write Transfer Size                   826
Overall Transfer Size                 827

C:\Program Files\IBM\TPC\cli>_
```

*Figure 6-55 tpctool lsmetrics subsystem command output*

8. After gathering the components and the metrics, we need to script a **getrpt** command to extract the metrics from the TPC database as shown in Example 6-20. In this example, we use all the metrics that are currently supported by component type `subsystem`.

*Example 6-20 getrpt command using all the subsystem metrics*

```
getrpt -user myuser -pwd mypass -url myurl -subsys fast4500 -level sample
-ctype subsystem -columns 803,806,809,812,815,818,819,820,821,825,826,827-fs
```

9. As we have described, if we do not stipulate a date and timestamp with a duration, this command will only return the last data collected. Example 6-21 shows the command with a start time.

*Example 6-21 getrpt command using all the subsystem metrics with timestamp and duration*

```
getrpt -user myuser -pwd mypass -url myurl -subsys fast4500 -level sample
-ctype subsystem -columns 803,806,809,812,815,818,819,820,821,825,826,827
-start 2006.09.19:16:21:00 -duration 80000 -fs
```

10. This **getrpt** command can now be added into our tpccli.conf file (see Example 6-22).

*Example 6-22 tpccli.conf with getrpt command*

```
#TPC Cli configuration property file DO NOT EDIT
# Sun Sept 21 09:05pm
myuser=amit
mypass=#@$sheetal143
myurl=9.43.85.143:9550

fast4400=600A0B80000CBDAD00000000440F3BA5+15
fast4500=600A0B800017443100000000041EFF538+15

rptfast4400=getrpt -user myuser -pwd mypass -url myurl -subsys fast4400 -level
sample -ctype subsystem -columns
803,806,809,812,815,818,819,820,821,825,826,827 -start 2006.09.19:16:21:00
-duration 80000 -fs ;

rptfast4500=getrpt -user myuser -pwd mypass -url myurl -subsys fast4500 -level
sample -ctype subsystem -columns
803,806,809,812,815,818,819,820,821,825,826,827 -start 2006.09.19:16:21:00
-duration 80000 -fs ;
```

11. We run the script `rptfast4500` (the alias given to the **getrpt** command), which was created in Example 6-22. The output is similar to Example 6-23.

*Example 6-23 tpctool getrpt output*

```
Timestamp;Device;Component;803;806;809;812;815;818;819;820;821;825;826;827
==========================================================================
2006.09.20:14:26:49;DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI;0.0066;0.0464;0.0531;10,
000;0;1,250;0.0000;0.0000;0.0000;0.5000;0.9286;0.8750
2006.09.19:16:21:00;DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI;0.0066;0.0465;0.0532;10,
000;0;1,250;0.0000;0.0000;0.0000;0.5000;0.9286;0.8750
2006.09.19:16:31:02;DS4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B80001744310000000041EFF538-LSI;0.7330;0.0498;0.7828;10,
000;0;9,364;0.1415;0.0000;0.1416;197.7036;0.8667;185.1928
```

```
2006.09.19:16:41:05;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7276;0.0033;0.7309;10,
000;0;9,954;0.1418;0.0000;0.1418;199.5046;1.0000;198.6023
2006.09.19:16:51:07;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.0066;0.0465;0.0532;10,
000;0;1,250;0.0000;0.0000;0.0000;0.5000;0.9286;0.8750
2006.09.19:17:01:09;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7342;0.0498;0.7841;10,
000;0;9,364;0.1418;0.0000;0.1418;197.7036;0.8667;185.1928
2006.09.19:17:11:11;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7276;0.0033;0.7309;10,
000;0;9,954;0.1418;0.0000;0.1418;199.5046;1.0000;198.6023
2006.09.19:17:21:13;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.0066;0.0464;0.0531;10,
000;0;1,250;0.0000;0.0000;0.0000;0.5000;0.9286;0.8750
2006.09.19:17:31:16;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7330;0.0498;0.7828;10,
000;0;9,364;0.1415;0.0000;0.1416;197.7036;0.8667;185.1928
2006.09.19:17:41:19;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7276;0.0033;0.7309;10,
000;0;9,954;0.1418;0.0000;0.1418;199.5046;1.0000;198.6023
2006.09.19:17:51:21;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.0066;0.0464;0.0531;10,
000;0;1,250;0.0000;0.0000;0.0000;0.5000;0.9286;0.8750
2006.09.19:18:01:24;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7342;0.0498;0.7841;10,
000;0;9,364;0.1418;0.0000;0.1418;197.7036;0.8667;185.1928
2006.09.19:18:11:26;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7252;0.0033;0.7285;10,
000;0;9,954;0.1413;0.0000;0.1413;199.5046;1.0000;198.6023
2006.09.19:18:21:30;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.0066;0.0465;0.0532;10,
000;0;1,250;0.0000;0.0000;0.0000;0.5000;0.9286;0.8750
2006.09.19:18:31:32;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7330;0.0498;0.7828;10,
000;0;9,364;0.1415;0.0000;0.1416;197.7036;0.8667;185.1928
2006.09.19:18:41:35;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;0.7264;0.0033;0.7297;10,
000;0;9,954;0.1415;0.0000;0.1415;199.5046;1.0000;198.6023
2006.09.19:18:51:38;DS4500-ITSODS4500_A-600A0B8000174431000000041EFF538-LSI;DS
```

12. The extracted data that was returned can be redirected to a file and imported into a spreadsheet using the command in Example 6-24.

*Example 6-24 tpctool getrpt command redirected data to an output file*

```
tpctool rptfast4500 > c:\performace\rptfast4500.out
```

## Importing data into Excel

We will show you in this section how to import the output data from the **getrpt** (rptfast4500) script that we executed for the subsystem performance report into Excel or a similar spreadsheet and create a template for later use. This imported data will now be in an easier to read format for analysis and creating graphs (reports).

You might need to export or import data regularly. In this case, the data has be exported to the rptfast4500.out file and then read by the application (Microsoft Excel). Alternatively, you can copy data on an as needed basis.

We built this example using Microsoft Excel Office 2000. The first task is to import the data into Excel:

1. Open a new Excel document as shown in Figure 6-56.

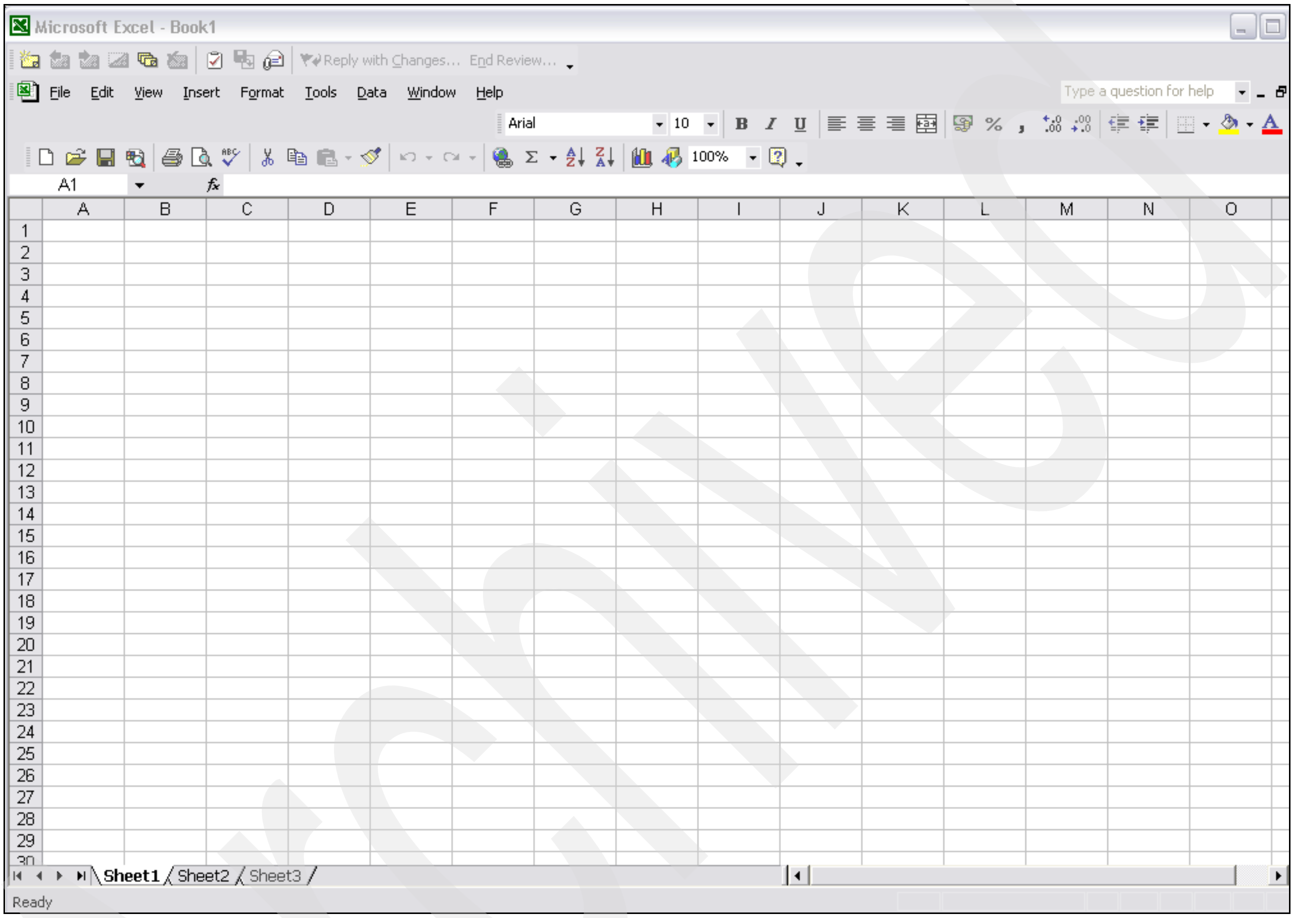
Microsoft Excel - Book1
Reply with Changes... End Review...
File Edit View Insert Format Tools Data Window Help
Type a question for help
Arial
10
100%
A1
Sheet1 Sheet2 Sheet3
Ready

*Figure 6-56 A new Excel document*

2. Select **Data** → **Import External Data** → **Import Data** as shown in Figure 6-57 on page 284.

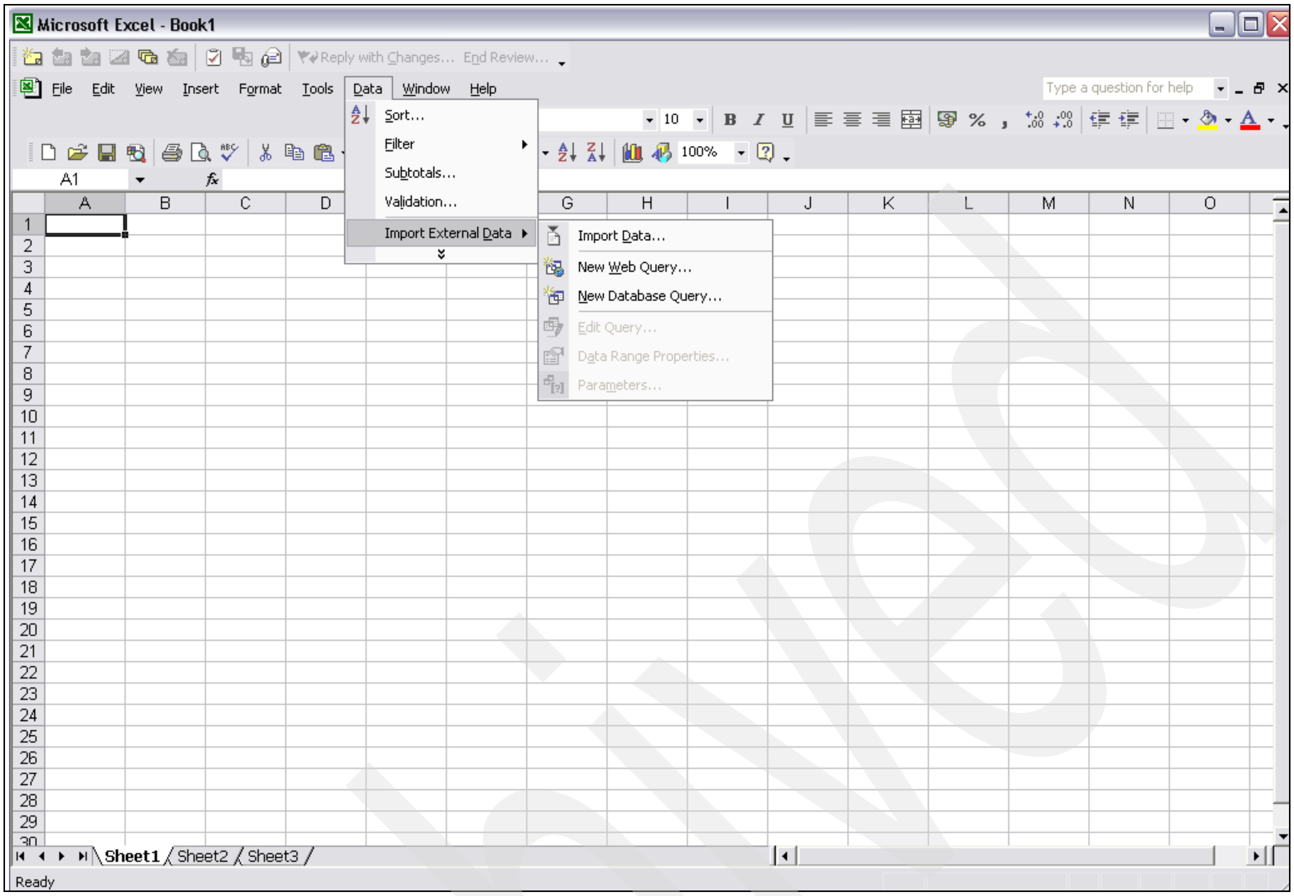

*Figure 6-57 Select Data → Import External Data → Import Data*

3. Locate the data file in the directory where you store the TPC CLI output file (**rptfast4500.out**). After you select the data file, click **Open** to start the import process as shown in Figure 6-58 on page 285.

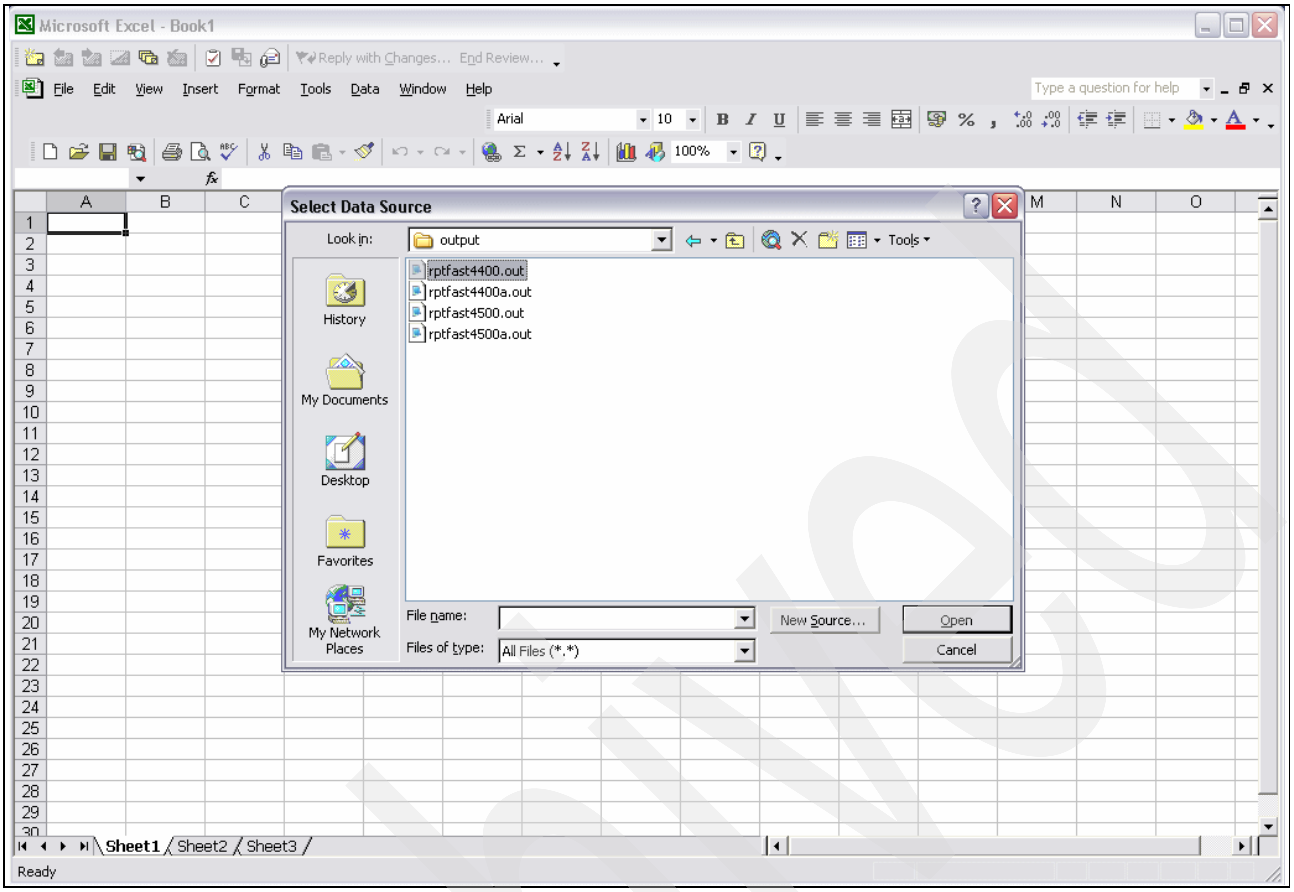

*Figure 6-58 Select the data source to begin the import process*

4. This starts the text import wizard. See Figure 6-59 on page 286. Select **Delimited** (default). Click **Next** to continue.

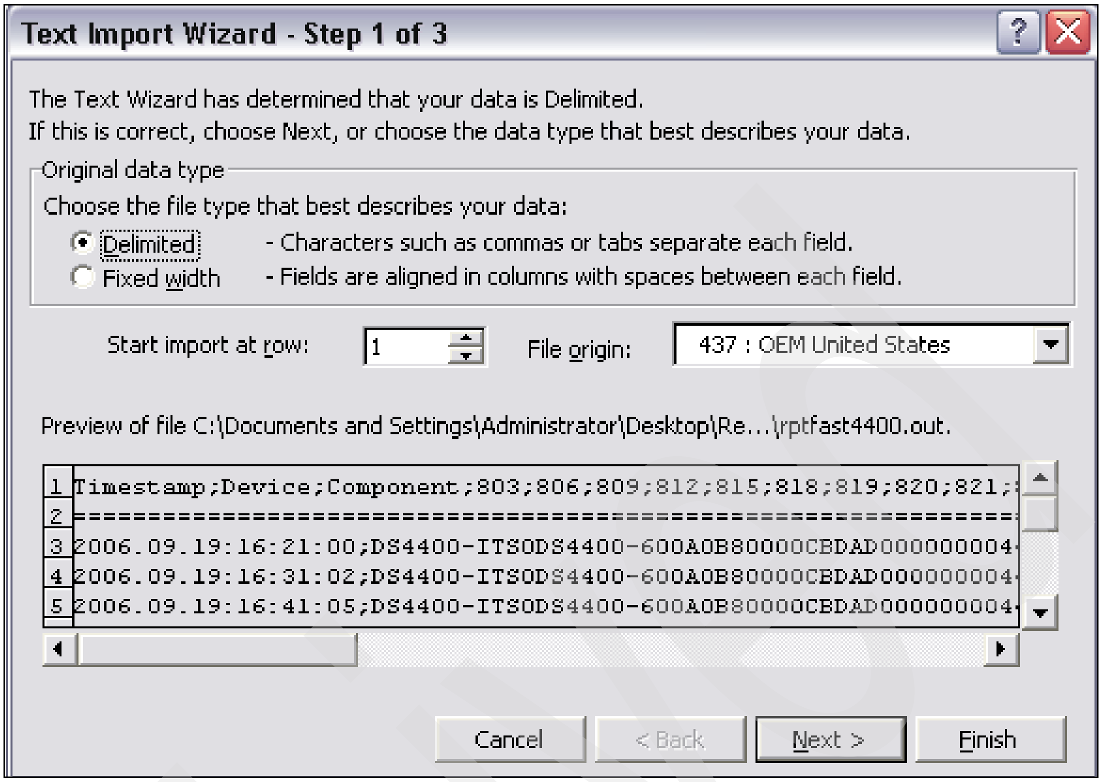

Text Import Wizard - Step 1 of 3

The Text Wizard has determined that your data is Delimited.
If this is correct, choose Next, or choose the data type that best describes your data.

Original data type

Choose the file type that best describes your data:

● Delimited - Characters such as commas or tabs separate each field.
○ Fixed width - Fields are aligned in columns with spaces between each field.

Start import at row: 1 File origin: 437 : OEM United States

Preview of file C:\Documents and Settings\Administrator\Desktop\Re...\rptfast4400.out.

```
1 Timestamp;Device;Component;803;806;809;812;815;818;819;820;821;
2 ==============================================================
3 2006.09.19:16:21:00;DS4400-ITSODS4400-600A0B80000CBDAD000000004
4 2006.09.19:16:31:02;DS4400-ITSODS4400-600A0B80000CBDAD000000004
5 2006.09.19:16:41:05;DS4400-ITSODS4400-600A0B80000CBDAD000000004
```

*Figure 6-59 Select Delimited (default) and click Next*

5. Select **Tab** (default) and enter the delimiter you select in your script. For our example, we selected the semicolon character (;) by checking the box to the left of **Semicolon**. See Example 6-25 and Figure 6-60 on page 287.

*Example 6-25 The delimiter chosen in our script -fs*

```
getrpt -user myuser -pwd mypass -url myurl -subsys fast4500 -level sample
-ctype subsystem -columns 803,806,809,812,815,818,819,820,821,825,826,827
-start 2006.09.19:16:21:00 -duration 80000 -fs ;
```

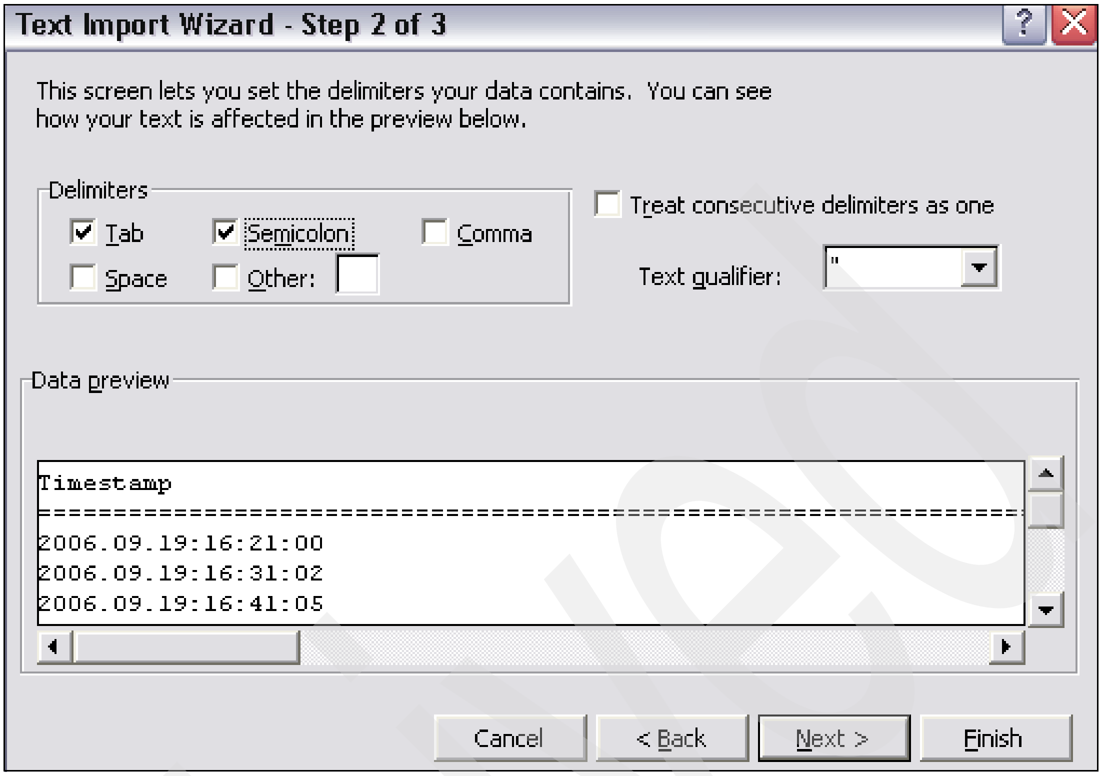

## Text Import Wizard - Step 2 of 3

This screen lets you set the delimiters your data contains. You can see how your text is affected in the preview below.

Delimiters

☑ Tab ☑ Semicolon □ Comma

□ Space □ Other: □

□ Treat consecutive delimiters as one

Text qualifier: "

Data preview

```
Timestamp
=====================================================================
2006.09.19:16:21:00
2006.09.19:16:31:02
2006.09.19:16:41:05
```

Figure 6-60 Select delimiter Semicolon

6. After you select the delimiter, you click **Next** and then click **Finish** (see Figure 6-61 on page 288).

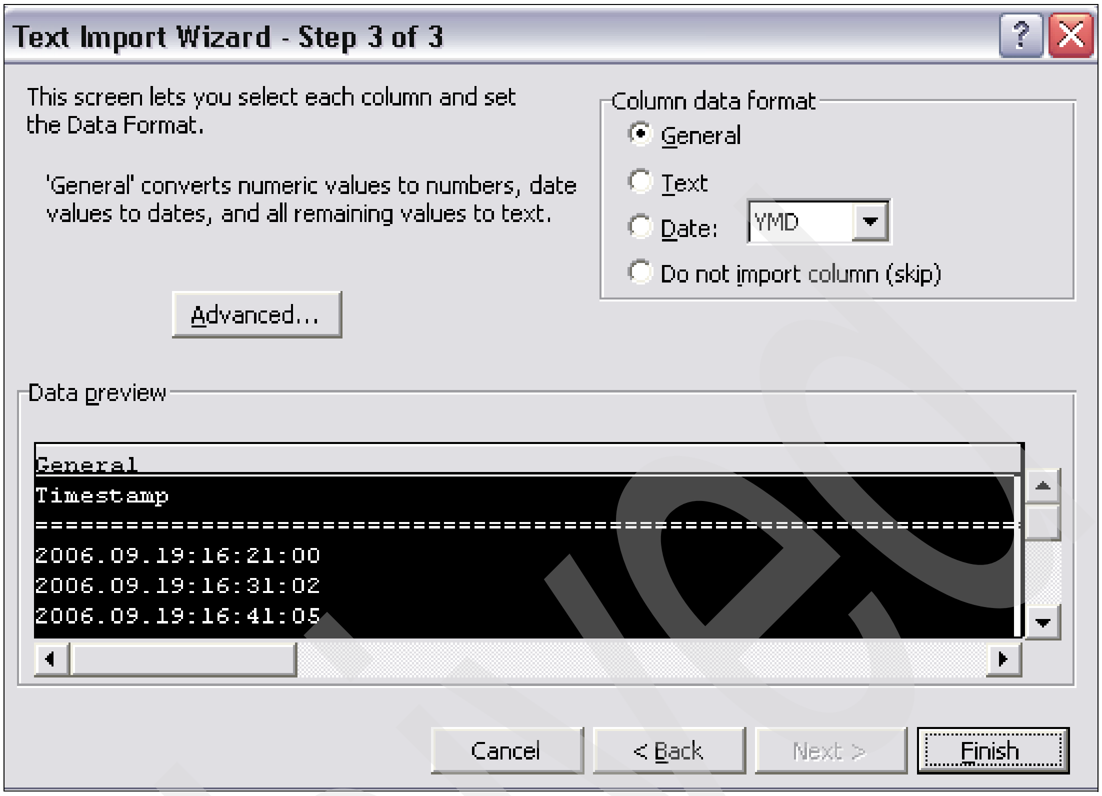

Text Import Wizard - Step 3 of 3

This screen lets you select each column and set the Data Format.

'General' converts numeric values to numbers, date values to dates, and all remaining values to text.

Column data format

● General

○ Text

○ Date: YMD

○ Do not import column (skip)

Data preview

```
General
Timestamp
================================================================
2006.09.19:16:21:00
2006.09.19:16:31:02
2006.09.19:16:41:05
```

*Figure 6-61 Select Finish to complete the wizard*

7. See Figure 6-62. Click **OK** to complete the task.

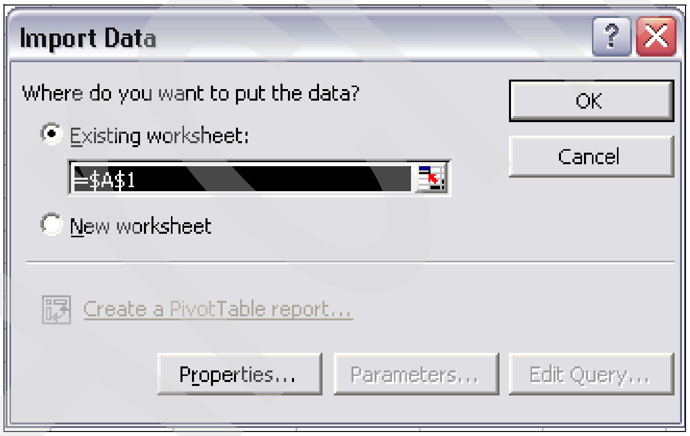

Import Data

Where do you want to put the data?

● Existing worksheet:

=$A$1

○ New worksheet

*Figure 6-62 Press OK to complete the task*

8. Figure 6-63 on page 289 shows the output.

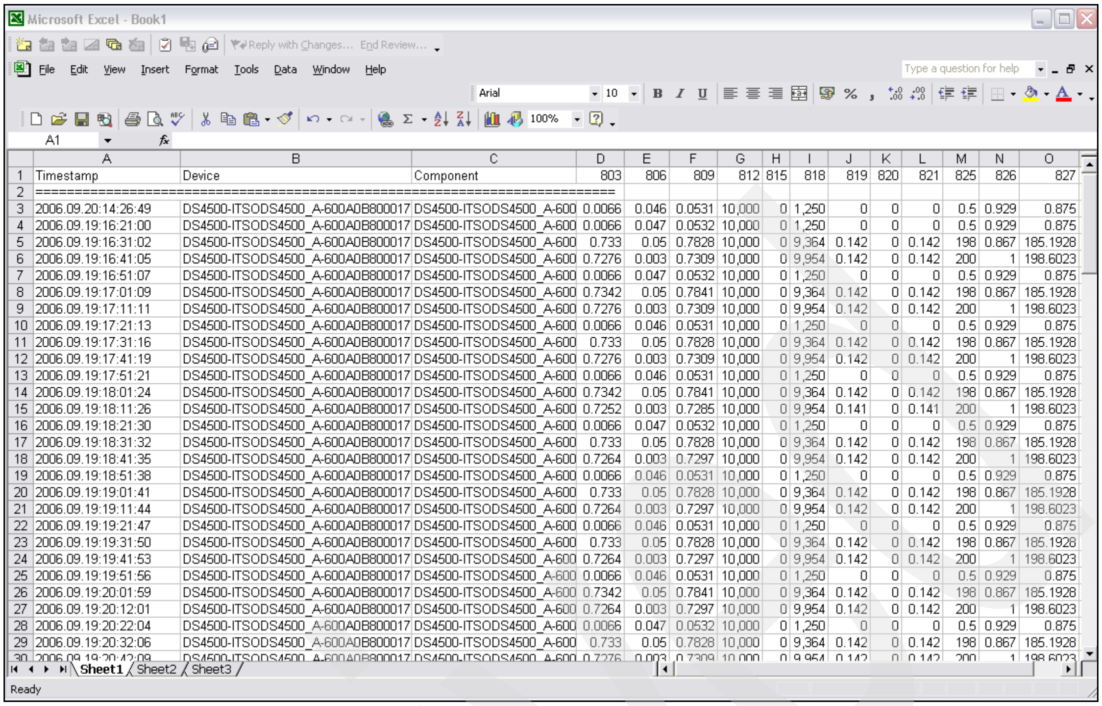
Microsoft Excel - Book1
Reply with Changes... End Review...
File Edit View Insert Format Tools Data Window Help
Type a question for help
Arial 10
100%
A1
Sheet1 Sheet2 Sheet3
Ready

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Component | 803 | 806 | 809 | 812 | 815 | 818 | 819 | 820 | 821 | 825 | 826 | 827 |
| 2 | ============================================================ | | | | | | | | | | | | | | |
| 3 | 2006.09.20:14:26:49 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.046 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 4 | 2006.09.19:16:21:00 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.047 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 5 | 2006.09.19:16:31:02 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.733 | 0.05 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 6 | 2006.09.19:16:41:05 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7276 | 0.003 | 0.7309 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |
| 7 | 2006.09.19:16:51:07 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.047 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 8 | 2006.09.19:17:01:09 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7342 | 0.05 | 0.7841 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 9 | 2006.09.19:17:11:11 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7276 | 0.003 | 0.7309 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |
| 10 | 2006.09.19:17:21:13 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.046 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 11 | 2006.09.19:17:31:16 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.733 | 0.05 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 12 | 2006.09.19:17:41:19 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7276 | 0.003 | 0.7309 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |
| 13 | 2006.09.19:17:51:21 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.046 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 14 | 2006.09.19:18:01:24 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7342 | 0.05 | 0.7841 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 15 | 2006.09.19:18:11:26 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7252 | 0.003 | 0.7285 | 10,000 | 0 | 9,954 | 0.141 | 0 | 0.141 | 200 | 1 | 198.6023 |
| 16 | 2006.09.19:18:21:30 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.047 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 17 | 2006.09.19:18:31:32 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.733 | 0.05 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 18 | 2006.09.19:18:41:35 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7264 | 0.003 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |
| 19 | 2006.09.19:18:51:38 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.046 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 20 | 2006.09.19:19:01:41 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.733 | 0.05 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 21 | 2006.09.19:19:11:44 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7264 | 0.003 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |
| 22 | 2006.09.19:19:21:47 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.046 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 23 | 2006.09.19:19:31:50 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.733 | 0.05 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 24 | 2006.09.19:19:41:53 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7264 | 0.003 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |
| 25 | 2006.09.19:19:51:56 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.046 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 26 | 2006.09.19:20:01:59 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7342 | 0.05 | 0.7841 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 27 | 2006.09.19:20:12:01 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7264 | 0.003 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |
| 28 | 2006.09.19:20:22:04 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.0066 | 0.047 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.929 | 0.875 |
| 29 | 2006.09.19:20:32:06 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.733 | 0.05 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.142 | 198 | 0.867 | 185.1928 |
| 30 | 2006.09.19:20:42:09 | DS4500-ITSODS4500_A-600A0B800017 | DS4500-ITSODS4500_A-600 | 0.7276 | 0.003 | 0.7309 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.142 | 200 | 1 | 198.6023 |

*Figure 6-63 Excel spreadsheet with unformatted columns*

9. The Excel document needs to be formatted to allow a template to be created. First, you need to delete row number 2 and then rename the metrics to more familiar headings. These metrics headings can be obtained by using the **`lsmetrics`** command. The Excel spreadsheet result is similar to Figure 6-64.

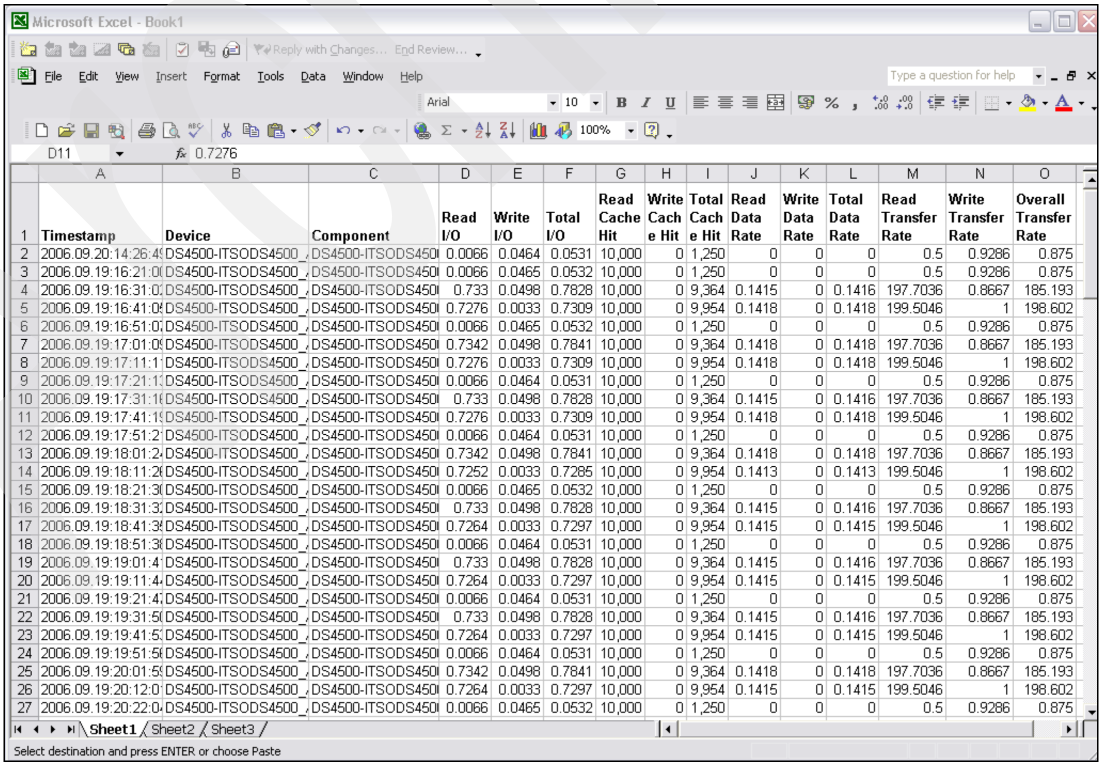

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | **Timestamp** | **Device** | **Component** | **Read I/O** | **Write I/O** | **Total I/O** | **Read Cache Hit** | **Write Cach e Hit** | **Total Cach e Hit** | **Read Data Rate** | **Write Data Rate** | **Total Data Rate** | **Read Transfer Rate** | **Write Transfer Rate** | **Overall Transfer Rate** |
| 2 | 2006.09.20:14:26:4 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 3 | 2006.09.19:16:21:0 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 4 | 2006.09.19:16:31:0 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.1415 | 0 | 0.1416 | 197.7036 | 0.8667 | 185.193 |
| 5 | 2006.09.19:16:41:0 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7276 | 0.0033 | 0.7309 | 10,000 | 0 | 9,954 | 0.1418 | 0 | 0.1418 | 199.5046 | 1 | 198.602 |
| 6 | 2006.09.19:16:51:0 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 7 | 2006.09.19:17:01:0 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7342 | 0.0498 | 0.7841 | 10,000 | 0 | 9,364 | 0.1418 | 0 | 0.1418 | 197.7036 | 0.8667 | 185.193 |
| 8 | 2006.09.19:17:11:1 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7276 | 0.0033 | 0.7309 | 10,000 | 0 | 9,954 | 0.1418 | 0 | 0.1418 | 199.5046 | 1 | 198.602 |
| 9 | 2006.09.19:17:21:1 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 10 | 2006.09.19:17:31:1 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.1415 | 0 | 0.1416 | 197.7036 | 0.8667 | 185.193 |
| 11 | 2006.09.19:17:41:1 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7276 | 0.0033 | 0.7309 | 10,000 | 0 | 9,954 | 0.1418 | 0 | 0.1418 | 199.5046 | 1 | 198.602 |
| 12 | 2006.09.19:17:51:2 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 13 | 2006.09.19:18:01:2 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7342 | 0.0498 | 0.7841 | 10,000 | 0 | 9,364 | 0.1418 | 0 | 0.1418 | 197.7036 | 0.8667 | 185.193 |
| 14 | 2006.09.19:18:11:2 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7252 | 0.0033 | 0.7285 | 10,000 | 0 | 9,954 | 0.1413 | 0 | 0.1413 | 199.5046 | 1 | 198.602 |
| 15 | 2006.09.19:18:21:3 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 16 | 2006.09.19:18:31:3 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.1415 | 0 | 0.1416 | 197.7036 | 0.8667 | 185.193 |
| 17 | 2006.09.19:18:41:3 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.1415 | 0 | 0.1415 | 199.5046 | 1 | 198.602 |
| 18 | 2006.09.19:18:51:3 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 19 | 2006.09.19:19:01:4 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.1415 | 0 | 0.1416 | 197.7036 | 0.8667 | 185.193 |
| 20 | 2006.09.19:19:11:4 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.1415 | 0 | 0.1415 | 199.5046 | 1 | 198.602 |
| 21 | 2006.09.19:19:21:4 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 22 | 2006.09.19:19:31:5 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.1415 | 0 | 0.1416 | 197.7036 | 0.8667 | 185.193 |
| 23 | 2006.09.19:19:41:5 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.1415 | 0 | 0.1415 | 199.5046 | 1 | 198.602 |
| 24 | 2006.09.19:19:51:5 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 25 | 2006.09.19:20:01:5 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7342 | 0.0498 | 0.7841 | 10,000 | 0 | 9,364 | 0.1418 | 0 | 0.1418 | 197.7036 | 0.8667 | 185.193 |
| 26 | 2006.09.19:20:12:0 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.1415 | 0 | 0.1415 | 199.5046 | 1 | 198.602 |
| 27 | 2006.09.19:20:22:0 | DS4500-ITSODS4500_ | DS4500-ITSODS450 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |

Select destination and press ENTER or choose Paste

*Figure 6-64 Excel spreadsheet with renamed columns*

10. From now on, you need to copy row 1 (the headings row) and paste it into a new book. The reason you do this is that if you attempt to delete the data below row 1 within Excel and save it as a template, when you reuse Excel, Excel will prompt you to refresh the data from the original source file. Thus, we need to copy row 1, the headings row, into a new book.
11. After you have copied the headings row into a new book, you can save this spreadsheet as a template.

### Timestamp

This extension inserts the current date and time into a message or an input field in the browser. A *timestamp* can refer to a time code or to a digitally signed timestamp. Timestamps are very useful for logging events.

Date, time, and their variants differ more than any other data types when you compare formats between *DATETIME* and *DATE* (see Example 6-26).

*Example 6-26 timestamp variants*

```
   2005-05-08 10:45
   Sat June 29 23:16:57 2005
   2005.08.03:10:45
```

The international standard date notation is *YYYY-MM-DD* where *YYYY* is the year, *MM* is the month of the year between `01 (January)` and `12 (December)`, and *DD* is the day of the month between `01` and `31`. For example, the third day of August in the year 1980 is written in the standard notation as `1980-08-03`.

The international standard notation for the time of day is *hh:mm:ss,* where *hh* is the number of complete hours that have passed since midnight (`00-24`), *mm* is the number of complete minutes that have passed since the start of the hour (`00-59`), and *ss* is the number of complete seconds since the start of the minute (`00-60`). If the hour value is 24, then the minute and second values must be zero.

For example, if time is `23:59:59`, this represents the time one second before midnight.

All of the time comparison procedures require the time objects to be of the same type. It is an error to use these procedures on time objects of different types. For the timestamp measurements, we need to convert the timestamp to meet international standard regardless of the format used by each country. We have developed a Visual Basic® Script (vbs) to collect the format that is used by your workstation and convert the timestamp to an international format (see Example 6-27).

*Example 6-27 Visual Basic script to convert timestamp to international format*

```
' Get Date and Time Separator String

  DS = Application.International(xlDateSeparator)
  TS = Application.International(xlTimeSeparator)
  If Application.International(xl24HourClock) Then
    AMPM = ""
  Else
    AMPM = " AM/PM"
  End If

  ' This loop runs until there is nothing in the next column
```

```
    Dim TimeStamp As String

    Do
       If Application.International(xlDateOrder) = 0 Then
             ActiveCell.NumberFormat = "mm" + DS + "dd" + DS + "yyyy hh" + TS +
"mm" + TS + "ss" + AMPM

       ElseIf Application.International(xlDateOrder) = 1 Then
             ActiveCell.NumberFormat = "dd" + DS + "mm" + DS + "yyyy hh" + TS +
"mm" + TS + "ss" + AMPM

       ElseIf Application.International(xlDateOrder) = 2 Then
            ActiveCell.NumberFormat = "yyyy" + DS + "mm" + DS + "dd hh" + TS +
"mm" + TS + "ss" + AMPM

      End If

       TimeStamp = Replace(ActiveCell.Value, ":", " ", 1, 1, vbTextCompare)
       TimeStamp = Replace(TimeStamp, ".", DS, 1, 2, vbTextCompare)
       TimeStamp = Replace(TimeStamp, ":", TS, 1, 2, vbTextCompare)
       ActiveCell.Value = TimeStamp
       ActiveCell.Offset(1, 0).Select

    Loop Until IsEmpty(ActiveCell)

ENDE:
    Application.ScreenUpdating = True
    Application.EnableEvents = True
```

You will need to copy this Visual Basic code into your Excel worksheet and save it as a macro. A *macro* automates a complex task. Excel macros can perform a complicated series of actions or simply record commonly used commands. Using the code Example 6-27 on page 290, we can fully automate the timestamp conversion.

These are the steps to create a macro:

1. Open the **Tools** menu, select **Macro**, and then select **Macros** (see Figure 6-65 on page 292).

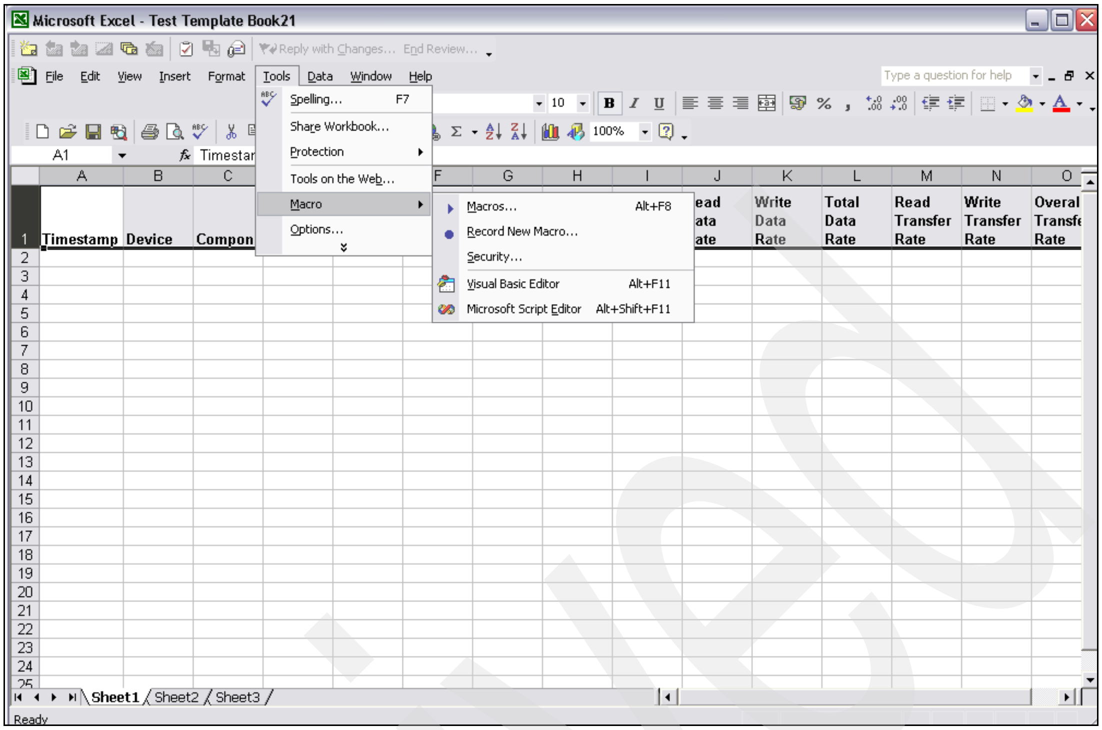

Figure 6-65 Select Macro

2. Complete the macro name in the box provided and click **Create**. Give the macro a descriptive name. In this case, we use *TimeStamp*. The macro can be available from only one worksheet or from any worksheet (see Figure 6-66).

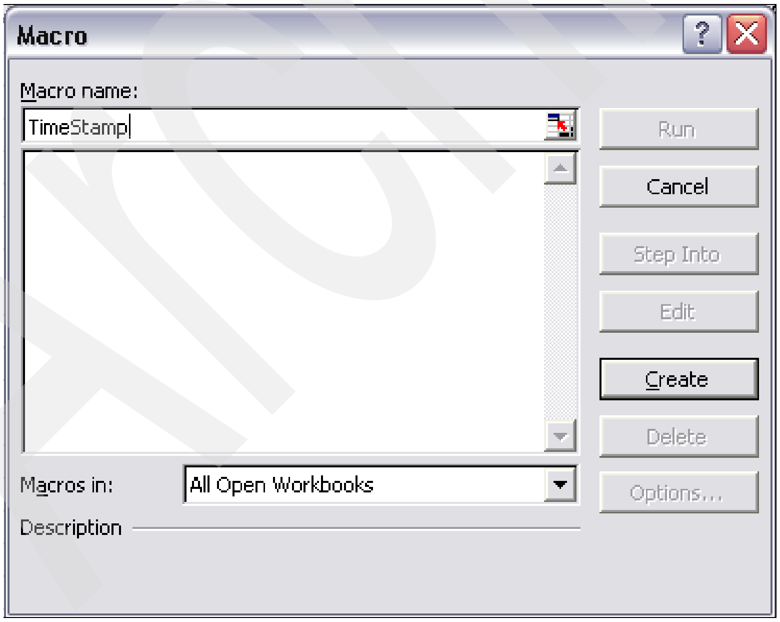

Macro name: TimeStamp

Macros in: All Open Workbooks

Figure 6-66 Macro name

3. Copy the code in Example 6-27 on page 290 into the Microsoft Visual Basic editor between the Sub line and the End Sub line (see Figure 6-67 on page 293).

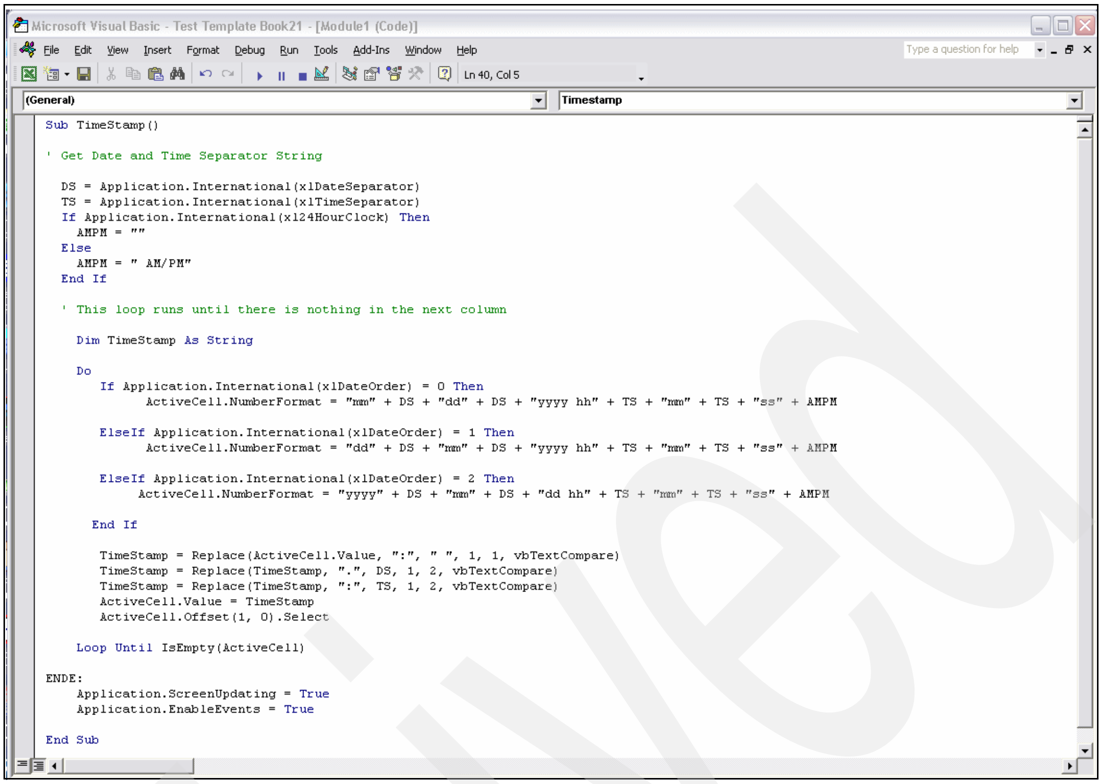
Microsoft Visual Basic - Test Template Book21 - [Module1 (Code)]
File Edit View Insert Format Debug Run Tools Add-Ins Window Help
Type a question for help
Ln 40, Col 5
(General)
Timestamp

```
Sub TimeStamp()

' Get Date and Time Separator String

  DS = Application.International(xlDateSeparator)
  TS = Application.International(xlTimeSeparator)
  If Application.International(xl24HourClock) Then
    AMPM = ""
  Else
    AMPM = " AM/PM"
  End If

  ' This loop runs until there is nothing in the next column

    Dim TimeStamp As String

    Do
       If Application.International(xlDateOrder) = 0 Then
             ActiveCell.NumberFormat = "mm" + DS + "dd" + DS + "yyyy hh" + TS + "mm" + TS + "ss" + AMPM

       ElseIf Application.International(xlDateOrder) = 1 Then
             ActiveCell.NumberFormat = "dd" + DS + "mm" + DS + "yyyy hh" + TS + "mm" + TS + "ss" + AMPM

       ElseIf Application.International(xlDateOrder) = 2 Then
            ActiveCell.NumberFormat = "yyyy" + DS + "mm" + DS + "dd hh" + TS + "mm" + TS + "ss" + AMPM

      End If

       TimeStamp = Replace(ActiveCell.Value, ":", " ", 1, 1, vbTextCompare)
       TimeStamp = Replace(TimeStamp, ".", DS, 1, 2, vbTextCompare)
       TimeStamp = Replace(TimeStamp, ":", TS, 1, 2, vbTextCompare)
       ActiveCell.Value = TimeStamp
       ActiveCell.Offset(1, 0).Select

    Loop Until IsEmpty(ActiveCell)

ENDE:
    Application.ScreenUpdating = True
    Application.EnableEvents = True

End Sub
```

*Figure 6-67 Copy Macro code in Visual Basic Editor*

4. After you copy the code into the editor, close the editor window, which brings you back into the Excel worksheet. The macro is saved automatically.
5. Now, you have successfully created a macro to convert the timestamp into an international standard.

## How to run the macro

After you have imported the data into your worksheet and you want to convert the timestamp into international standard and use it as a true value, you can run the macro. The macro is very simple to run.

1. Select **Tools** → **Macro** → **Macros** (Alt + F8) as shown in Figure 6-68 on page 294.

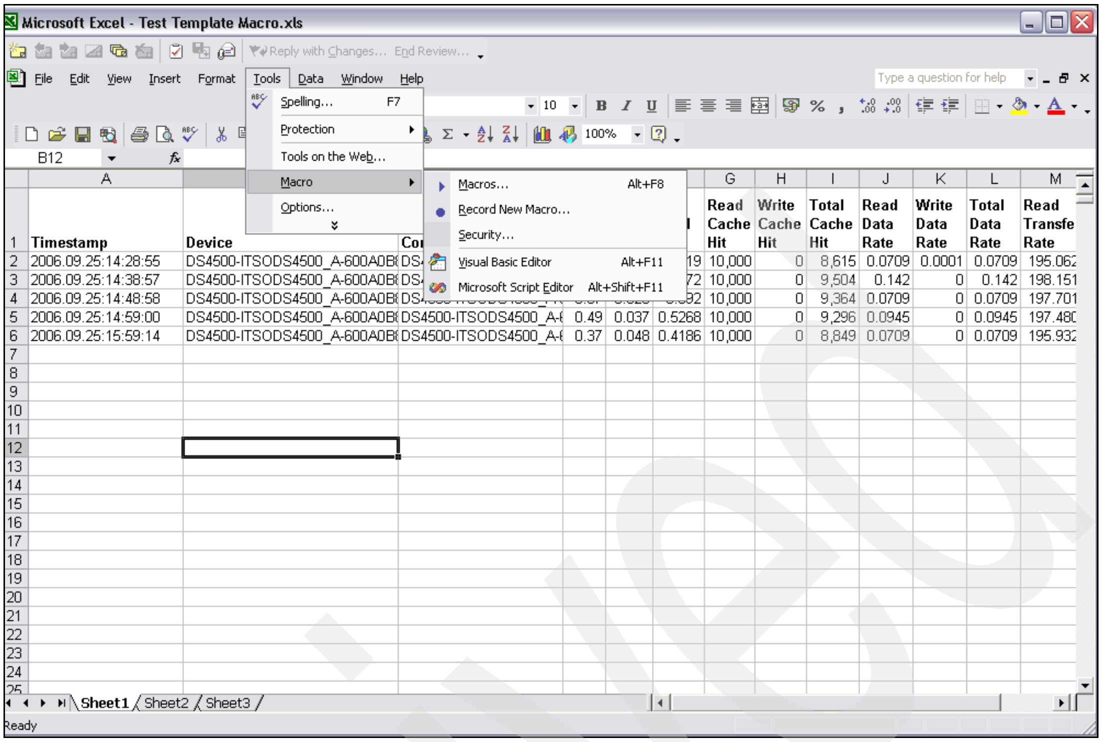

| | A | B | C | D | E | F | G | H | I | J | K | L | M |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Co | | | | Read Cache Hit | Write Cache Hit | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfe Rate |
| 2 | 2006.09.25:14:28:55 | DS4500-ITSODS4500_A-600A0B | DS | | | 19 | 10,000 | 0 | 8,615 | 0.0709 | 0.0001 | 0.0709 | 195.062 |
| 3 | 2006.09.25:14:38:57 | DS4500-ITSODS4500_A-600A0B | DS | | | 72 | 10,000 | 0 | 9,504 | 0.142 | 0 | 0.142 | 198.151 |
| 4 | 2006.09.25:14:48:58 | DS4500-ITSODS4500_A-600A0B | DS | | | 92 | 10,000 | 0 | 9,364 | 0.0709 | 0 | 0.0709 | 197.701 |
| 5 | 2006.09.25:14:59:00 | DS4500-ITSODS4500_A-600A0B | DS4500-ITSODS4500_A- | 0.49 | 0.037 | 0.5268 | 10,000 | 0 | 9,296 | 0.0945 | 0 | 0.0945 | 197.480 |
| 6 | 2006.09.25:15:59:14 | DS4500-ITSODS4500_A-600A0B | DS4500-ITSODS4500_A- | 0.37 | 0.048 | 0.4186 | 10,000 | 0 | 8,849 | 0.0709 | 0 | 0.0709 | 195.932 |

*Figure 6-68 Browse to run the macro*

2. Select the macro that you want to run. In this case, we select **Timestamp** and then click **Run** (see Figure 6-69).

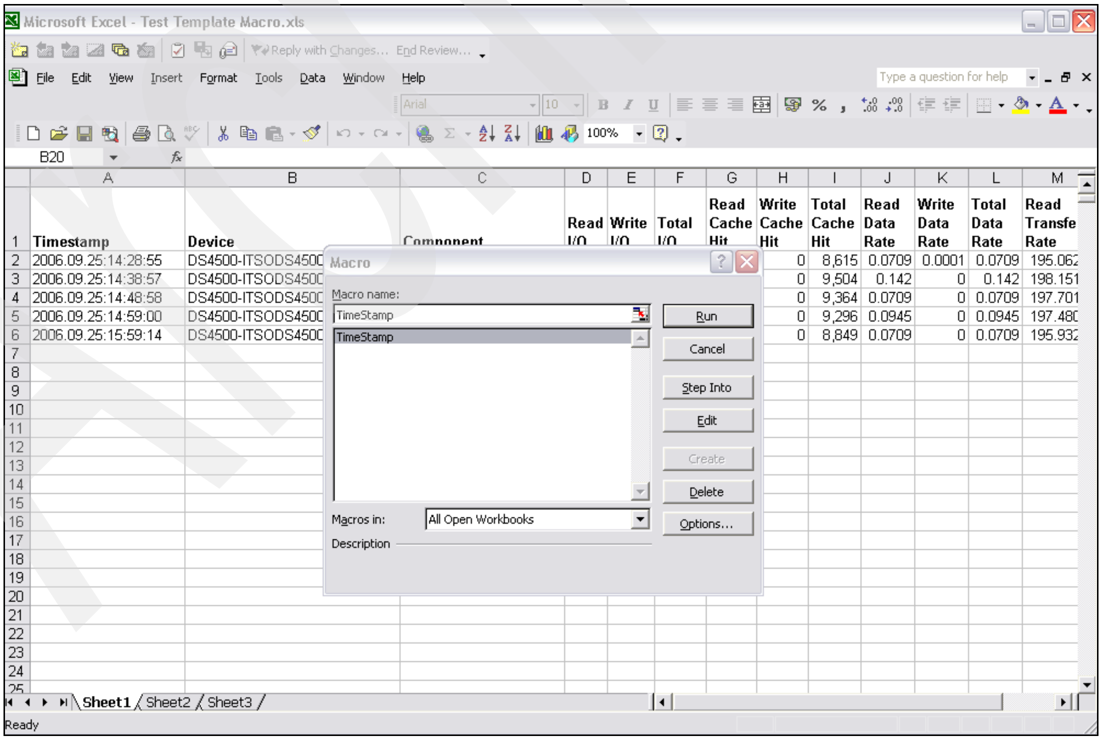

| | A | B | C | D | E | F | G | H | I | J | K | L | M |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Component | Read I/O | Write I/O | Total I/O | Read Cache Hit | Write Cache Hit | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfe Rate |
| 2 | 2006.09.25:14:28:55 | DS4500-ITSODS4500 | | | | | | 0 | 8,615 | 0.0709 | 0.0001 | 0.0709 | 195.062 |
| 3 | 2006.09.25:14:38:57 | DS4500-ITSODS4500 | | | | | | 0 | 9,504 | 0.142 | 0 | 0.142 | 198.151 |
| 4 | 2006.09.25:14:48:58 | DS4500-ITSODS4500 | | | | | | 0 | 9,364 | 0.0709 | 0 | 0.0709 | 197.701 |
| 5 | 2006.09.25:14:59:00 | DS4500-ITSODS4500 | | | | | | 0 | 9,296 | 0.0945 | 0 | 0.0945 | 197.480 |
| 6 | 2006.09.25:15:59:14 | DS4500-ITSODS4500 | | | | | | 0 | 8,849 | 0.0709 | 0 | 0.0709 | 195.932 |

Macro name: TimeStamp

TimeStamp

Macros in: All Open Workbooks

Description

*Figure 6-69 Select the macro to run*

3. This formats the timestamp column as shown in Figure 6-70 on page 295.

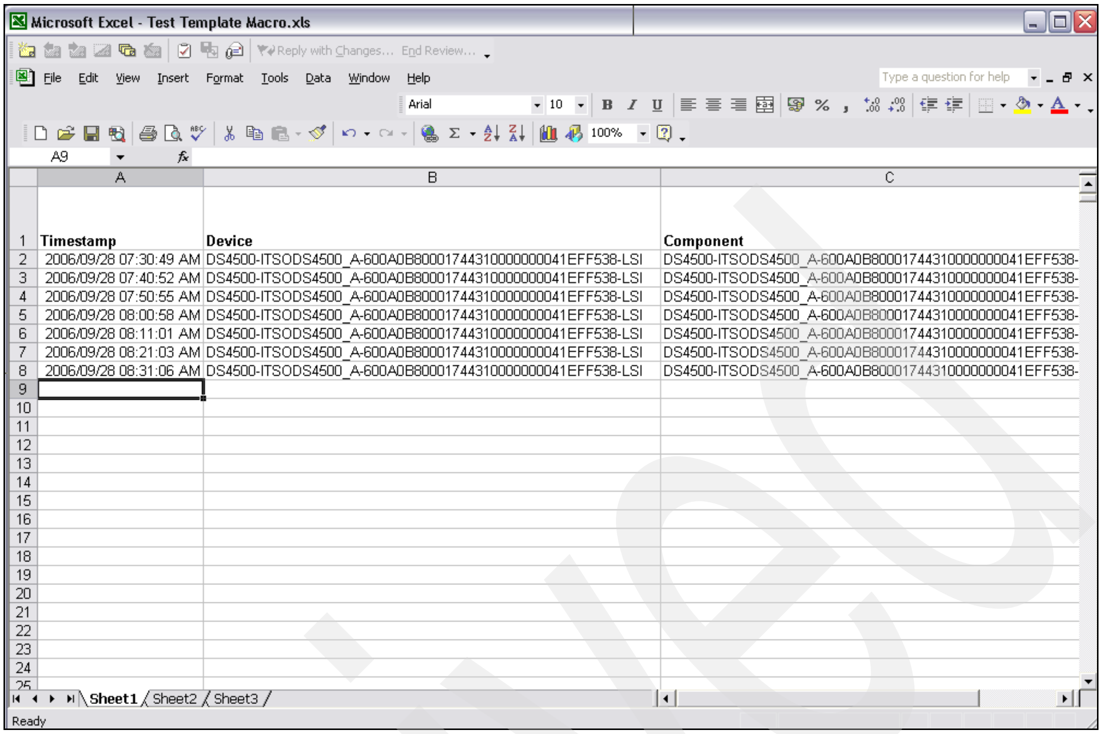
Microsoft Excel - Test Template Macro.xls
Reply with Changes... End Review...
File Edit View Insert Format Tools Data Window Help
Type a question for help
Arial
10
100%
A9
Sheet1
Sheet2
Sheet3
Ready

| | A | B | C |
|---|---|---|---|
| 1 | **Timestamp** | **Device** | **Component** |
| 2 | 2006/09/28 07:30:49 AM | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538- |
| 3 | 2006/09/28 07:40:52 AM | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538- |
| 4 | 2006/09/28 07:50:55 AM | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538- |
| 5 | 2006/09/28 08:00:58 AM | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538- |
| 6 | 2006/09/28 08:11:01 AM | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538- |
| 7 | 2006/09/28 08:21:03 AM | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538- |
| 8 | 2006/09/28 08:31:06 AM | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538-LSI | DS4500-ITSODS4500_A-600A0B800017443100000000041EFF538- |

*Figure 6-70 Timestamp column formatted*

4. You can save this as a complete template, which includes the performance report headings and the timestamp convertor.

## Creating a template

Using the Excel spreadsheet in Figure 6-70, we need create a template for later use either with or without the macro. This is dependent on the type of graphs that you want to create (see Figure 6-71 on page 296). Provide a descriptive name for the template.

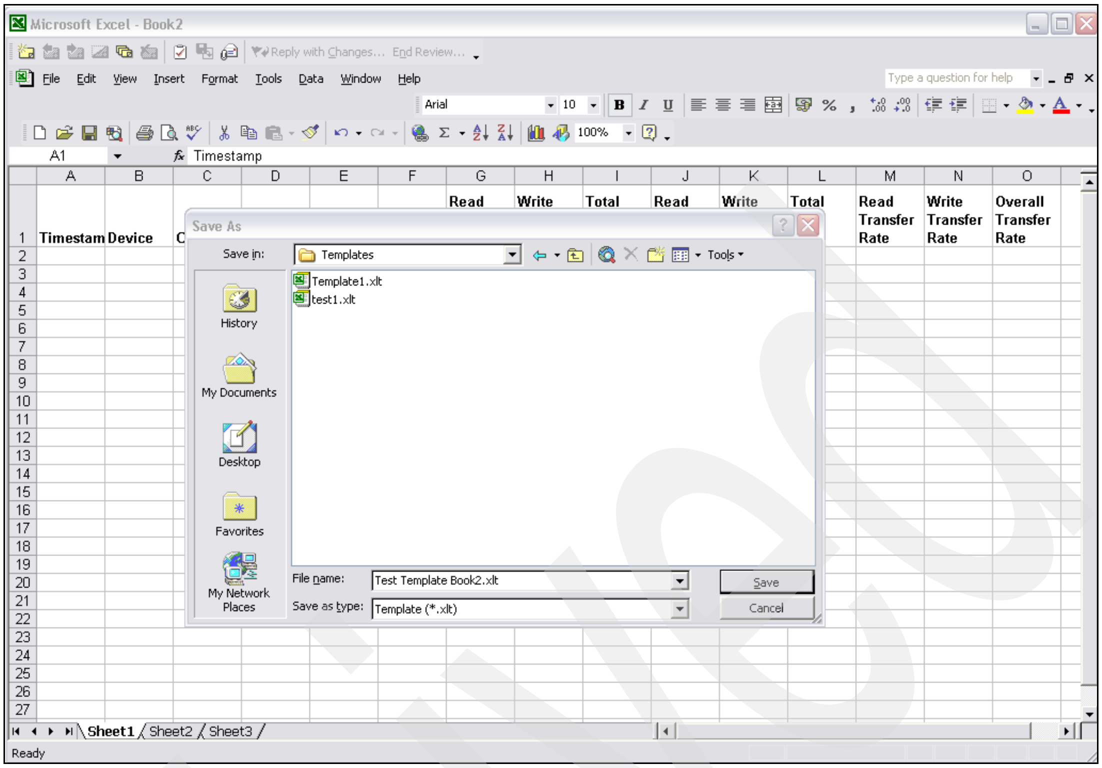

| A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Timestamp | Device | C |  |  |  | Read | Write | Total | Read | Write | Total | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |

Save in: Templates

Template1.xlt
test1.xlt

File name: Test Template Book2.xlt

Save as type: Template (*.xlt)

*Figure 6-71 Save the book as a template for reuse*

**Important:** Select **Template (*.xlt)** from the **Save Type As** drop-down box.

After you have saved this template, you can import extracted performance data into this spreadsheet.

**Restriction:** Only the same data criteria can be imported into the same template, for example, DS4000 subsystem metrics.

We have created a template into which to import your extracted data. You can modify these template headings to compensate for other metrics and reports; however, we have created the macro and the layout for you. The instructions to download the template are in Appendix A, “Additional material” on page 319.

## Importing into a saved template

Use the same methodology that we just described to import the extracted performance data into the saved template:

1. Open the saved template.
2. Select the source file.
3. Start the import wizard.
4. On the panel in Figure 6-72 on page 297, Text Import Wizard - Step 1 of 3, select **Start import at row** and choose row 3. This will remove the header from the source file and the additional separation line.

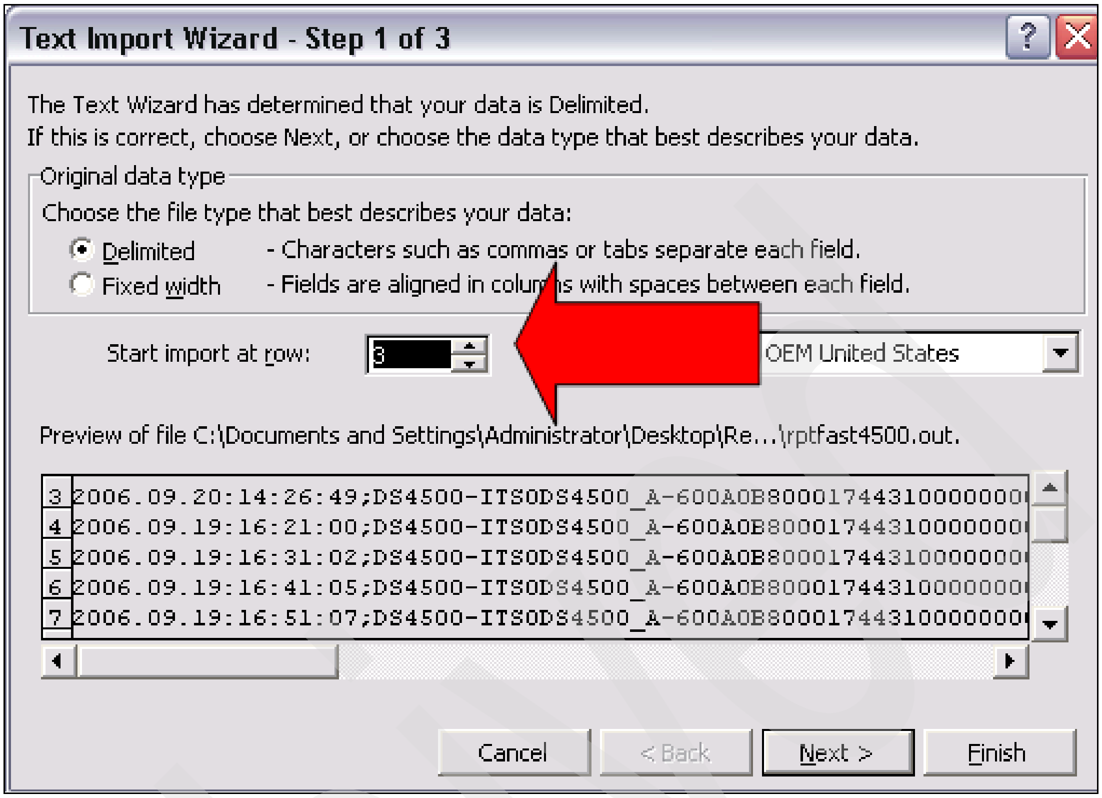

The Text Wizard has determined that your data is Delimited.
If this is correct, choose Next, or choose the data type that best describes your data.

Original data type
Choose the file type that best describes your data:

● Delimited - Characters such as commas or tabs separate each field.
○ Fixed width - Fields are aligned in columns with spaces between each field.

Start import at row: 3 File origin: OEM United States

Preview of file C:\Documents and Settings\Administrator\Desktop\Re...\rptfast4500.out.

```
3 2006.09.20:14:26:49;DS4500-ITSODS4500_A-600A0B8000174431000000000
4 2006.09.19:16:21:00;DS4500-ITSODS4500_A-600A0B8000174431000000000
5 2006.09.19:16:31:02;DS4500-ITSODS4500_A-600A0B8000174431000000000
6 2006.09.19:16:41:05;DS4500-ITSODS4500_A-600A0B8000174431000000000
7 2006.09.19:16:51:07;DS4500-ITSODS4500_A-600A0B8000174431000000000
```

*Figure 6-72 Change the start import row number in Excel*

5. Choose the delimiter to use when extracting the data in step 2 of 3.
6. Select **Finish**, which returns an output similar to Figure 6-73 on page 298.

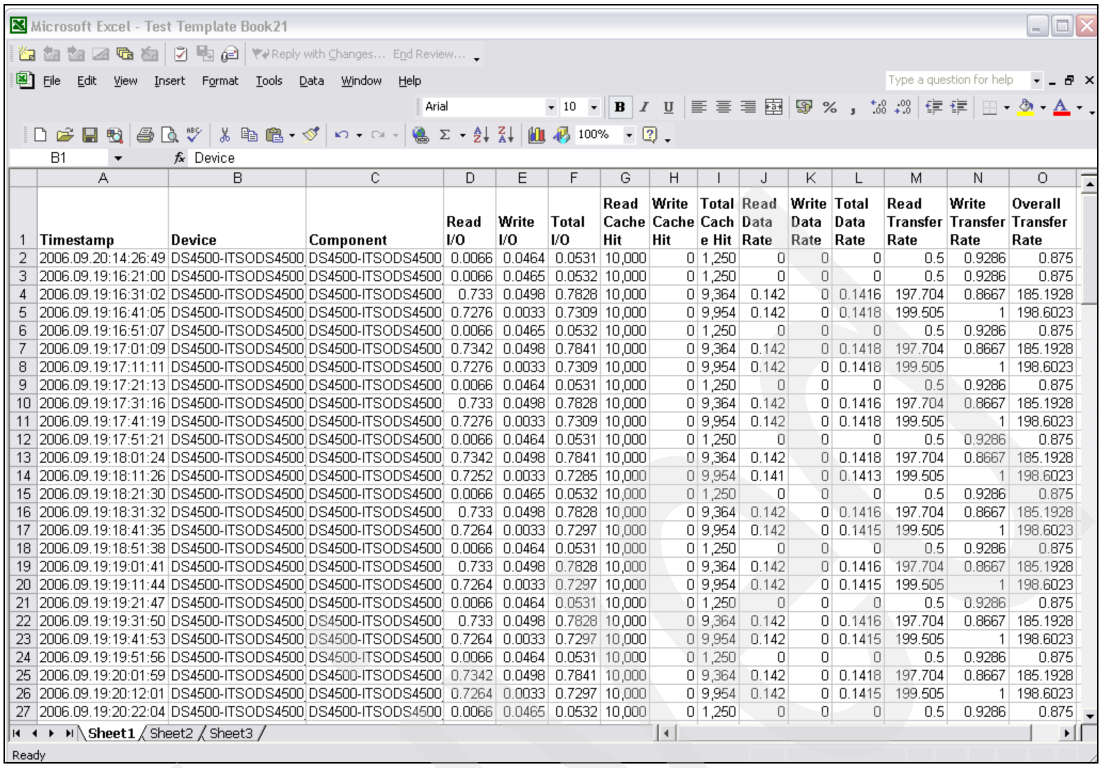
Microsoft Excel - Test Template Book21
Reply with Changes... End Review...
File Edit View Insert Format Tools Data Window Help
Type a question for help
Arial 10
100%
B1
Device
Sheet1 Sheet2 Sheet3
Ready

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Component | Read I/O | Write I/O | Total I/O | Read Cache Hit | Write Cache Hit | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |
| 2 | 2006.09.20:14:26:49 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 3 | 2006.09.19:16:21:00 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 4 | 2006.09.19:16:31:02 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1416 | 197.704 | 0.8667 | 185.1928 |
| 5 | 2006.09.19:16:41:05 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7276 | 0.0033 | 0.7309 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.1418 | 199.505 | 1 | 198.6023 |
| 6 | 2006.09.19:16:51:07 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 7 | 2006.09.19:17:01:09 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7342 | 0.0498 | 0.7841 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1418 | 197.704 | 0.8667 | 185.1928 |
| 8 | 2006.09.19:17:11:11 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7276 | 0.0033 | 0.7309 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.1418 | 199.505 | 1 | 198.6023 |
| 9 | 2006.09.19:17:21:13 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 10 | 2006.09.19:17:31:16 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1416 | 197.704 | 0.8667 | 185.1928 |
| 11 | 2006.09.19:17:41:19 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7276 | 0.0033 | 0.7309 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.1418 | 199.505 | 1 | 198.6023 |
| 12 | 2006.09.19:17:51:21 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 13 | 2006.09.19:18:01:24 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7342 | 0.0498 | 0.7841 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1418 | 197.704 | 0.8667 | 185.1928 |
| 14 | 2006.09.19:18:11:26 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7252 | 0.0033 | 0.7285 | 10,000 | 0 | 9,954 | 0.141 | 0 | 0.1413 | 199.505 | 1 | 198.6023 |
| 15 | 2006.09.19:18:21:30 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 16 | 2006.09.19:18:31:32 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1416 | 197.704 | 0.8667 | 185.1928 |
| 17 | 2006.09.19:18:41:35 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.1415 | 199.505 | 1 | 198.6023 |
| 18 | 2006.09.19:18:51:38 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 19 | 2006.09.19:19:01:41 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1416 | 197.704 | 0.8667 | 185.1928 |
| 20 | 2006.09.19:19:11:44 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.1415 | 199.505 | 1 | 198.6023 |
| 21 | 2006.09.19:19:21:47 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 22 | 2006.09.19:19:31:50 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.733 | 0.0498 | 0.7828 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1416 | 197.704 | 0.8667 | 185.1928 |
| 23 | 2006.09.19:19:41:53 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.1415 | 199.505 | 1 | 198.6023 |
| 24 | 2006.09.19:19:51:56 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0464 | 0.0531 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |
| 25 | 2006.09.19:20:01:59 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7342 | 0.0498 | 0.7841 | 10,000 | 0 | 9,364 | 0.142 | 0 | 0.1418 | 197.704 | 0.8667 | 185.1928 |
| 26 | 2006.09.19:20:12:01 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.7264 | 0.0033 | 0.7297 | 10,000 | 0 | 9,954 | 0.142 | 0 | 0.1415 | 199.505 | 1 | 198.6023 |
| 27 | 2006.09.19:20:22:04 | DS4500-ITSODS4500 | DS4500-ITSODS4500 | 0.0066 | 0.0465 | 0.0532 | 10,000 | 0 | 1,250 | 0 | 0 | 0 | 0.5 | 0.9286 | 0.875 |

*Figure 6-73 Final output of importing data to a template*

The data extracted from TPC is now available for analysis. Futhermore, you can use this data to create graphs.

Depending on the request and the analysis that is required, you can convert the timestamp into the international standard or you can use it as a string.

### Creating graphs

We can now create readable graphs by using the data that has been extracted from TPC and imported into the Excel template. Use graphs to determine relationships by plotting large numbers of data points and observing the grouping or clustering of the data points. The template helps you create a macro to copy data from an Excel spreadsheet to any application. You can record Excel macros so that you can reuse these templates for creating graphs in the future for reproducing graphs of the same type, such as in our example, subsystem reports for DS4500.

After creating your template and importing the data extracted from TPC by using the TPCTOOL, we need to understand what type of report you want to generate. In this example, we show you how to create a performance report using the `timestamp`, `Read I/O`, `Write I/O`, and `Total I/O`.

> **Note:** For this example, we have used the macro to convert the timestamp into international standard, and we use the timestamp as a data value and not as a string.

To create our graph:

1. Select the columns that you want to plot onto the graph, for example, `Timestamp`, `Read I/O`, `Write I/O`, and `Total I/O`. In Figure 6-74 on page 299, we only select a few line entries of data.

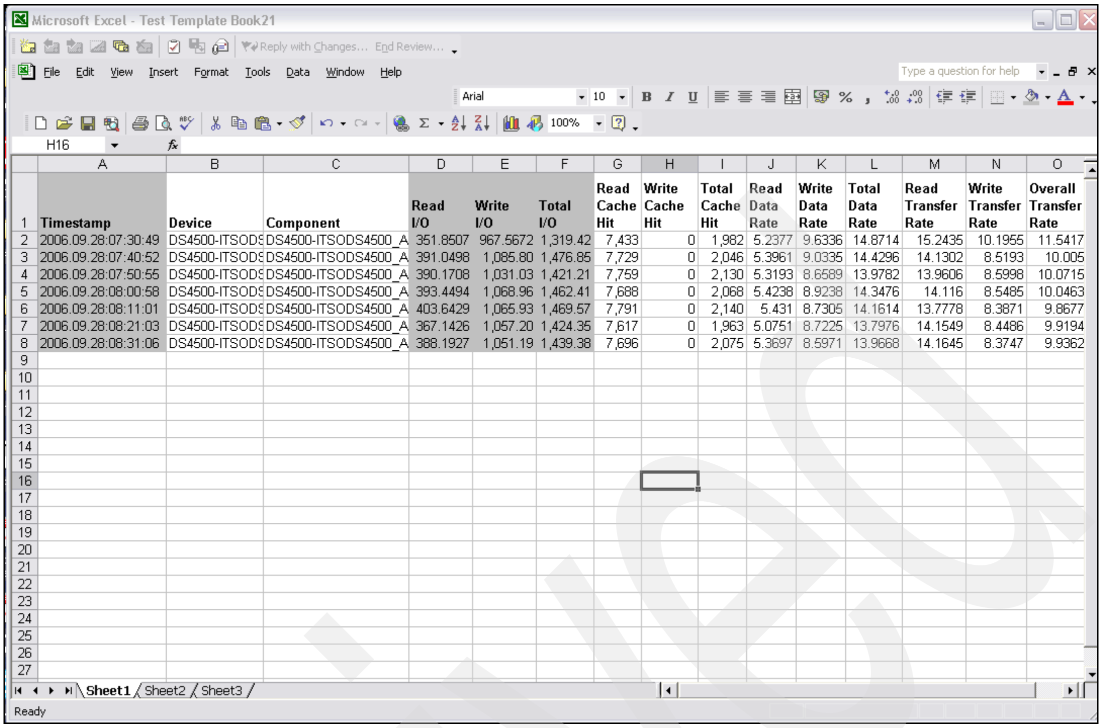
Microsoft Excel - Test Template Book21
Reply with Changes... End Review...
File Edit View Insert Format Tools Data Window Help
Type a question for help
Arial 10
100%
H16
Sheet1 Sheet2 Sheet3
Ready

|  | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Component | Read I/O | Write I/O | Total I/O | Read Cache Hit | Write Cache Hit | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |
| 2 | 2006.09.28:07:30:49 | DS4500-ITSODS | DS4500-ITSODS4500_A | 351.8507 | 967.5672 | 1,319.42 | 7,433 | 0 | 1,982 | 5.2377 | 9.6336 | 14.8714 | 15.2435 | 10.1955 | 11.5417 |
| 3 | 2006.09.28:07:40:52 | DS4500-ITSODS | DS4500-ITSODS4500_A | 391.0498 | 1,085.80 | 1,476.85 | 7,729 | 0 | 2,046 | 5.3961 | 9.0335 | 14.4296 | 14.1302 | 8.5193 | 10.005 |
| 4 | 2006.09.28:07:50:55 | DS4500-ITSODS | DS4500-ITSODS4500_A | 390.1708 | 1,031.03 | 1,421.21 | 7,759 | 0 | 2,130 | 5.3193 | 8.6589 | 13.9782 | 13.9606 | 8.5998 | 10.0715 |
| 5 | 2006.09.28:08:00:58 | DS4500-ITSODS | DS4500-ITSODS4500_A | 393.4494 | 1,068.96 | 1,462.41 | 7,688 | 0 | 2,068 | 5.4238 | 8.9238 | 14.3476 | 14.116 | 8.5485 | 10.0463 |
| 6 | 2006.09.28:08:11:01 | DS4500-ITSODS | DS4500-ITSODS4500_A | 403.6429 | 1,065.93 | 1,469.57 | 7,791 | 0 | 2,140 | 5.431 | 8.7305 | 14.1614 | 13.7778 | 8.3871 | 9.8677 |
| 7 | 2006.09.28:08:21:03 | DS4500-ITSODS | DS4500-ITSODS4500_A | 367.1426 | 1,057.20 | 1,424.35 | 7,617 | 0 | 1,963 | 5.0751 | 8.7225 | 13.7976 | 14.1549 | 8.4486 | 9.9194 |
| 8 | 2006.09.28:08:31:06 | DS4500-ITSODS | DS4500-ITSODS4500_A | 388.1927 | 1,051.19 | 1,439.38 | 7,696 | 0 | 2,075 | 5.3697 | 8.5971 | 13.9668 | 14.1645 | 8.3747 | 9.9362 |

*Figure 6-74 Select columns to create a graph*

2. You need to utilize the chart wizard to plot the values onto the graph. Select **Insert** → **Chart Wizard** as shown in Figure 6-75.

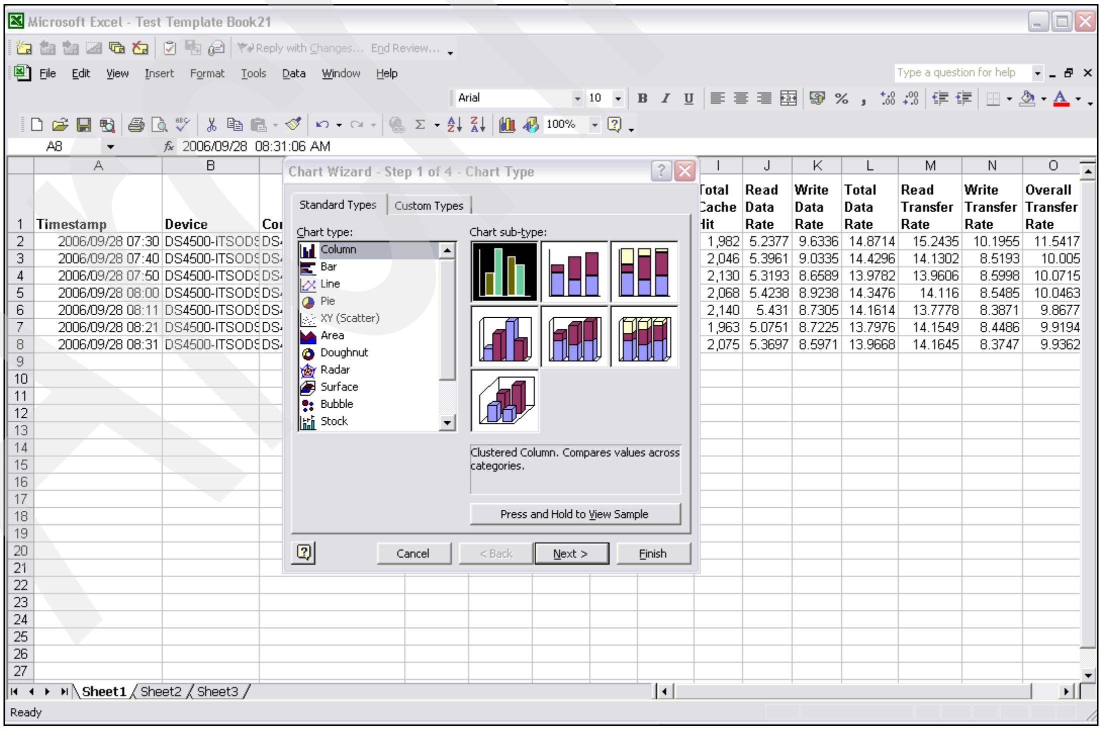

A8 2006/09/28 08:31:06 AM

|  | A | B | C | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Co | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |
| 2 | 2006/09/28 07:30 | DS4500-ITSODS | DS | 1,982 | 5.2377 | 9.6336 | 14.8714 | 15.2435 | 10.1955 | 11.5417 |
| 3 | 2006/09/28 07:40 | DS4500-ITSODS | DS | 2,046 | 5.3961 | 9.0335 | 14.4296 | 14.1302 | 8.5193 | 10.005 |
| 4 | 2006/09/28 07:50 | DS4500-ITSODS | DS | 2,130 | 5.3193 | 8.6589 | 13.9782 | 13.9606 | 8.5998 | 10.0715 |
| 5 | 2006/09/28 08:00 | DS4500-ITSODS | DS | 2,068 | 5.4238 | 8.9238 | 14.3476 | 14.116 | 8.5485 | 10.0463 |
| 6 | 2006/09/28 08:11 | DS4500-ITSODS | DS | 2,140 | 5.431 | 8.7305 | 14.1614 | 13.7778 | 8.3871 | 9.8677 |
| 7 | 2006/09/28 08:21 | DS4500-ITSODS | DS | 1,963 | 5.0751 | 8.7225 | 13.7976 | 14.1549 | 8.4486 | 9.9194 |
| 8 | 2006/09/28 08:31 | DS4500-ITSODS | DS | 2,075 | 5.3697 | 8.5971 | 13.9668 | 14.1645 | 8.3747 | 9.9362 |

Clustered Column. Compares values across categories.

*Figure 6-75 Begin Chart Wizard*

3. Select the Chart type that you want to use to create the graph.

## Graphs (line graphs and scatter plots)

*Line graphs* provide an excellent way to map independent variables and dependent variables that are both quantitative. When both variables are quantitative, the line segment that connects two points on the graph expresses a slope, which can be interpreted visually relative to the slope of other lines or expressed as a precise mathematical formula. *Scatter plots* are similar to line graphs in that they start with mapping quantitative data points. The difference is that with a scatter plot, the individual points are not connected directly with a line but instead, the points express a trend. You can see this trend directly through the distribution of points or with the addition of a regression line. You use a statistical tool to mathematically express a trend in the data.

> **Tip:** Using the XY scatter graph provides you with the most realistic graph, because both axes, X and Y, are true value plot points on the graph, instead of points. However, you can produce the graphs in any type for analysis. The graph types that you produce depend on the analysis that you will perform.

We use the XY scatter graph to analyze and produce the graph (Step 1 of 4) as shown in Figure 6-76.

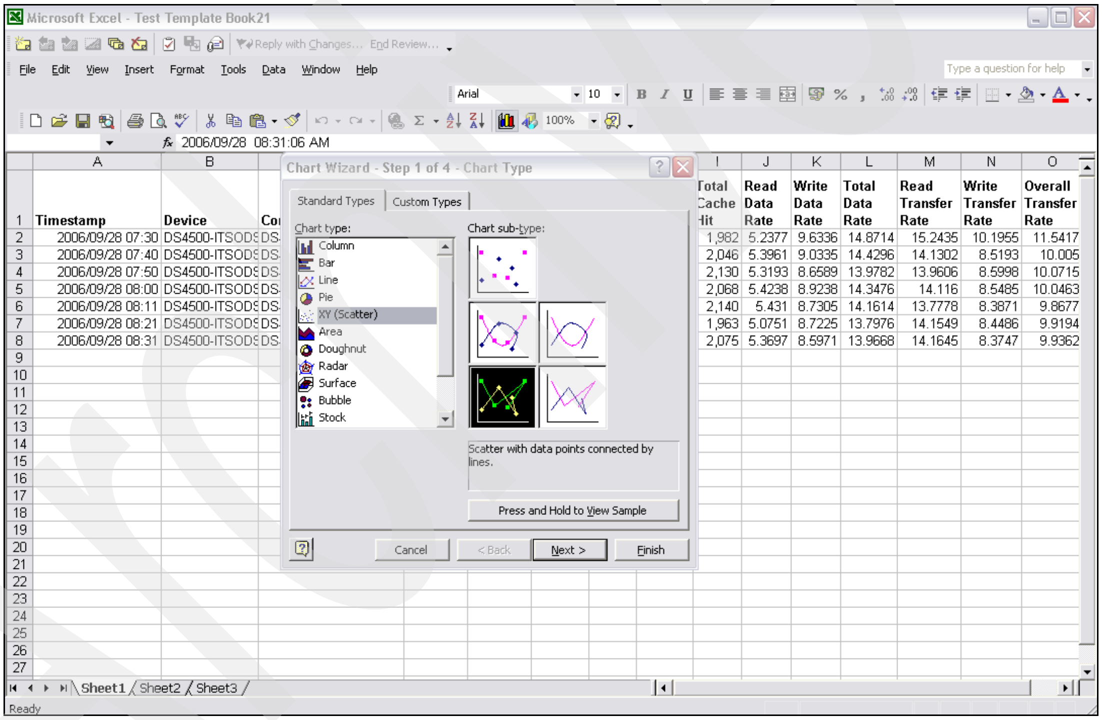

2006/09/28 08:31:06 AM

Scatter with data points connected by lines.

|  | A | B | C | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Co | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |
| 2 | 2006/09/28 07:30 | DS4500-ITSODS | DS | 1,982 | 5.2377 | 9.6336 | 14.8714 | 15.2435 | 10.1955 | 11.5417 |
| 3 | 2006/09/28 07:40 | DS4500-ITSODS | DS | 2,046 | 5.3961 | 9.0335 | 14.4296 | 14.1302 | 8.5193 | 10.005 |
| 4 | 2006/09/28 07:50 | DS4500-ITSODS | DS | 2,130 | 5.3193 | 8.6589 | 13.9782 | 13.9606 | 8.5998 | 10.0715 |
| 5 | 2006/09/28 08:00 | DS4500-ITSODS | DS | 2,068 | 5.4238 | 8.9238 | 14.3476 | 14.116 | 8.5485 | 10.0463 |
| 6 | 2006/09/28 08:11 | DS4500-ITSODS | DS | 2,140 | 5.431 | 8.7305 | 14.1614 | 13.7778 | 8.3871 | 9.8677 |
| 7 | 2006/09/28 08:21 | DS4500-ITSODS | DS | 1,963 | 5.0751 | 8.7225 | 13.7976 | 14.1549 | 8.4486 | 9.9194 |
| 8 | 2006/09/28 08:31 | DS4500-ITSODS | DS | 2,075 | 5.3697 | 8.5971 | 13.9668 | 14.1645 | 8.3747 | 9.9362 |

*Figure 6-76 Selecting XY Scatter graph*

4. Select **Next** after you choose the graph type. The next window confirms the data range (step 2 of 4). Click **Next** as shown in Figure 6-77 on page 301.

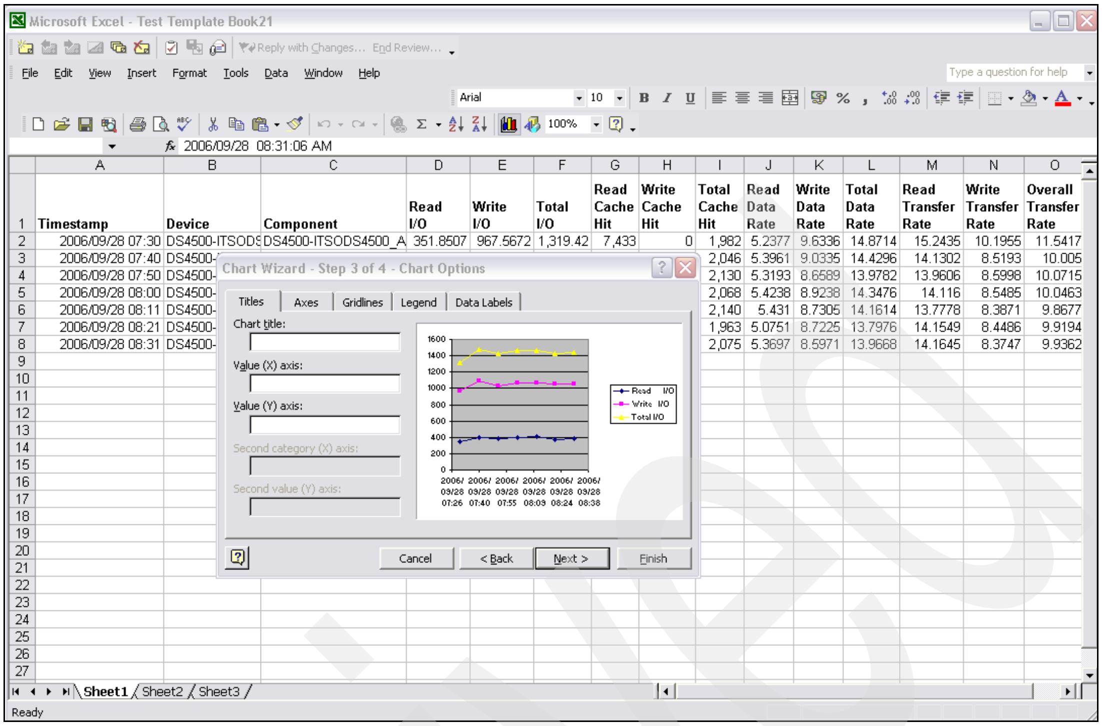

fx 2006/09/28 08:31:06 AM

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Component | Read I/O | Write I/O | Total I/O | Read Cache Hit | Write Cache Hit | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |
| 2 | 2006/09/28 07:30 | DS4500-ITSODS | DS4500-ITSODS4500_A | 351.8507 | 967.5672 | 1,319.42 | 7,433 | 0 | 1,982 | 5.2377 | 9.6336 | 14.8714 | 15.2435 | 10.1955 | 11.5417 |
| 3 | 2006/09/28 07:40 | DS4500- | | | | | | | 2,046 | 5.3961 | 9.0335 | 14.4296 | 14.1302 | 8.5193 | 10.005 |
| 4 | 2006/09/28 07:50 | DS4500- | | | | | | | 2,130 | 5.3193 | 8.6589 | 13.9782 | 13.9606 | 8.5998 | 10.0715 |
| 5 | 2006/09/28 08:00 | DS4500- | | | | | | | 2,068 | 5.4238 | 8.9238 | 14.3476 | 14.116 | 8.5485 | 10.0463 |
| 6 | 2006/09/28 08:11 | DS4500- | | | | | | | 2,140 | 5.431 | 8.7305 | 14.1614 | 13.7778 | 8.3871 | 9.8677 |
| 7 | 2006/09/28 08:21 | DS4500- | | | | | | | 1,963 | 5.0751 | 8.7225 | 13.7976 | 14.1549 | 8.4486 | 9.9194 |
| 8 | 2006/09/28 08:31 | DS4500- | | | | | | | 2,075 | 5.3697 | 8.5971 | 13.9668 | 14.1645 | 8.3747 | 9.9362 |

Chart title:

Value (X) axis:

Value (Y) axis:

*Figure 6-77 Step 3 of 4 - Chart Options*

5. You need to enter the chart title, the X value axis title, and the Y value axis title. Use descriptive titles for each graph as shown in Figure 6-78.

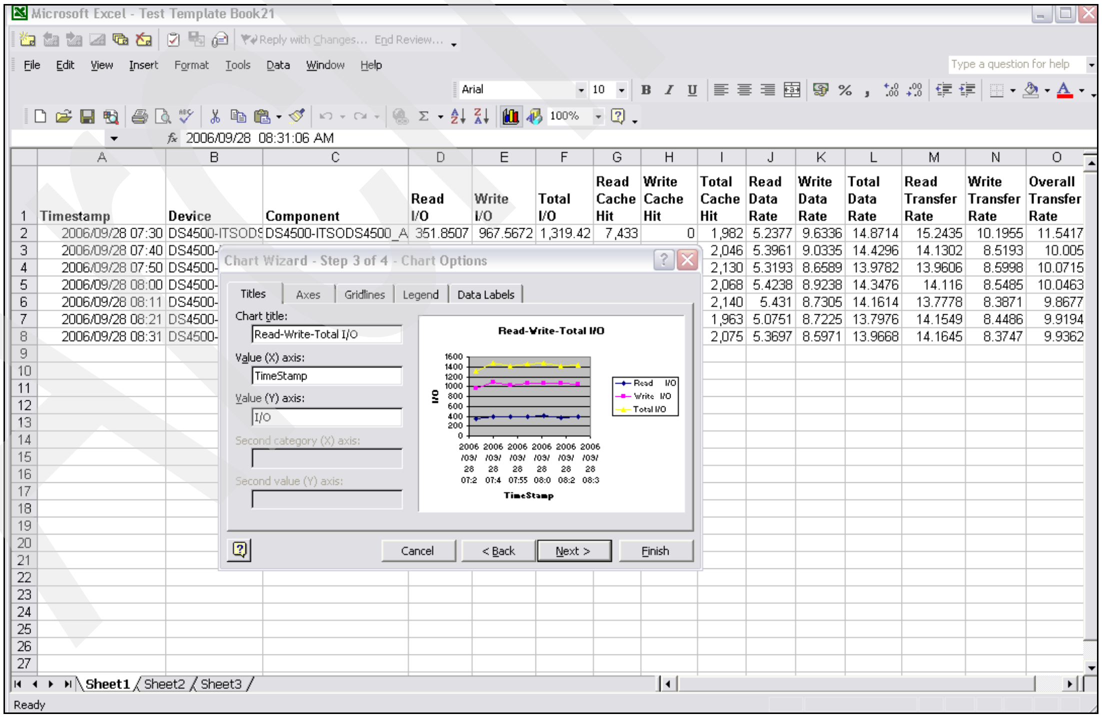
Microsoft Excel - Test Template Book21
File Edit View Insert Format Tools Data Window Help
Type a question for help
Arial 10
100%
Chart Wizard - Step 3 of 4 - Chart Options
Titles Axes Gridlines Legend Data Labels
Second category (X) axis:
Second value (Y) axis:
Read-Write-Total I/O
I/O
1600
1400
1200
1000
800
600
400
200
0
TimeStamp
Read I/O
Write I/O
Total I/O
Cancel
< Back
Next >
Finish
Sheet1 Sheet2 Sheet3
Ready

Chart title: Read-Write-Total I/O

Value (X) axis: TimeStamp

Value (Y) axis: I/O

fx 2006/09/28 08:31:06 AM

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Component | Read I/O | Write I/O | Total I/O | Read Cache Hit | Write Cache Hit | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |
| 2 | 2006/09/28 07:30 | DS4500-ITSODS | DS4500-ITSODS4500_A | 351.8507 | 967.5672 | 1,319.42 | 7,433 | 0 | 1,982 | 5.2377 | 9.6336 | 14.8714 | 15.2435 | 10.1955 | 11.5417 |
| 3 | 2006/09/28 07:40 | DS4500- | | | | | | | 2,046 | 5.3961 | 9.0335 | 14.4296 | 14.1302 | 8.5193 | 10.005 |
| 4 | 2006/09/28 07:50 | DS4500- | | | | | | | 2,130 | 5.3193 | 8.6589 | 13.9782 | 13.9606 | 8.5998 | 10.0715 |
| 5 | 2006/09/28 08:00 | DS4500- | | | | | | | 2,068 | 5.4238 | 8.9238 | 14.3476 | 14.116 | 8.5485 | 10.0463 |
| 6 | 2006/09/28 08:11 | DS4500- | | | | | | | 2,140 | 5.431 | 8.7305 | 14.1614 | 13.7778 | 8.3871 | 9.8677 |
| 7 | 2006/09/28 08:21 | DS4500- | | | | | | | 1,963 | 5.0751 | 8.7225 | 13.7976 | 14.1549 | 8.4486 | 9.9194 |
| 8 | 2006/09/28 08:31 | DS4500- | | | | | | | 2,075 | 5.3697 | 8.5971 | 13.9668 | 14.1645 | 8.3747 | 9.9362 |

*Figure 6-78 Input chart titles*

6. The next panel (step 4 of 4) allows you to create the graph in the same sheet or in a new sheet. This is entirely up to you. If you produce multiple graphs from the same template,

then we recommend that you place each graph in a new sheet. In this example, we use the same sheet. The graph is now generated as shown in Figure 6-79.

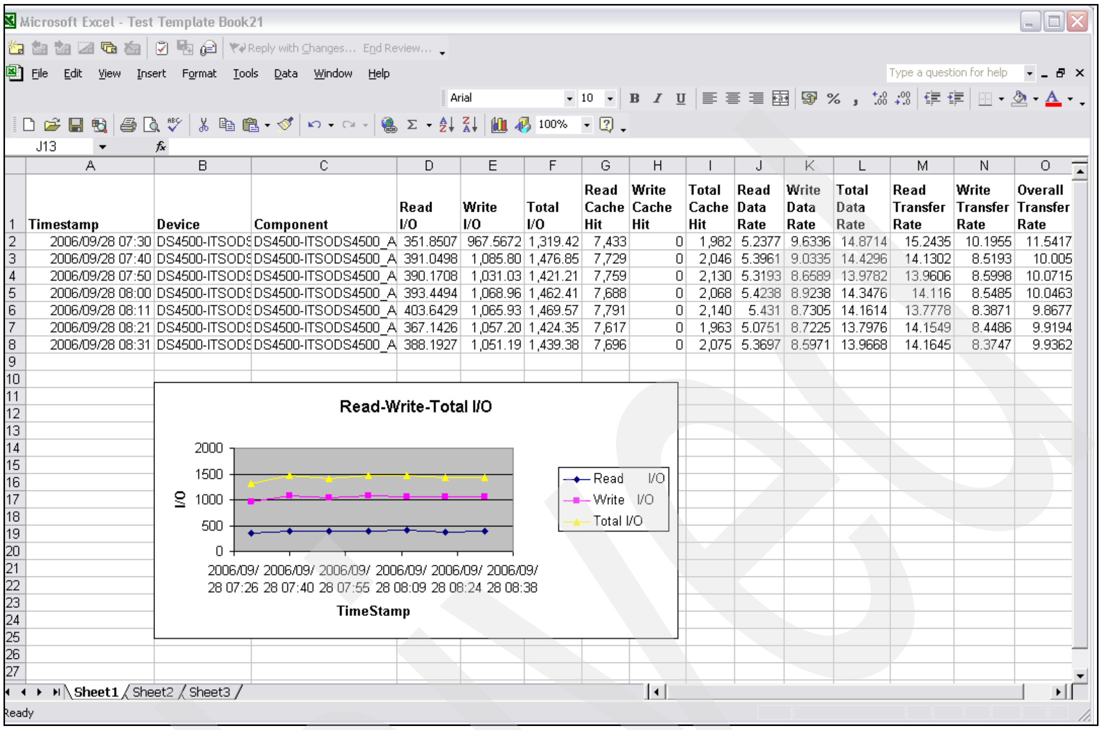

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Timestamp | Device | Component | Read I/O | Write I/O | Total I/O | Read Cache Hit | Write Cache Hit | Total Cache Hit | Read Data Rate | Write Data Rate | Total Data Rate | Read Transfer Rate | Write Transfer Rate | Overall Transfer Rate |
| 2 | 2006/09/28 07:30 | DS4500-ITSODS | DS4500-ITSODS4500_A | 351.8507 | 967.5672 | 1,319.42 | 7,433 | 0 | 1,982 | 5.2377 | 9.6336 | 14.8714 | 15.2435 | 10.1955 | 11.5417 |
| 3 | 2006/09/28 07:40 | DS4500-ITSODS | DS4500-ITSODS4500_A | 391.0498 | 1,085.80 | 1,476.85 | 7,729 | 0 | 2,046 | 5.3961 | 9.0335 | 14.4296 | 14.1302 | 8.5193 | 10.005 |
| 4 | 2006/09/28 07:50 | DS4500-ITSODS | DS4500-ITSODS4500_A | 390.1708 | 1,031.03 | 1,421.21 | 7,759 | 0 | 2,130 | 5.3193 | 8.6589 | 13.9782 | 13.9606 | 8.5998 | 10.0715 |
| 5 | 2006/09/28 08:00 | DS4500-ITSODS | DS4500-ITSODS4500_A | 393.4494 | 1,068.96 | 1,462.41 | 7,688 | 0 | 2,068 | 5.4238 | 8.9238 | 14.3476 | 14.116 | 8.5485 | 10.0463 |
| 6 | 2006/09/28 08:11 | DS4500-ITSODS | DS4500-ITSODS4500_A | 403.6429 | 1,065.93 | 1,469.57 | 7,791 | 0 | 2,140 | 5.431 | 8.7305 | 14.1614 | 13.7778 | 8.3871 | 9.8677 |
| 7 | 2006/09/28 08:21 | DS4500-ITSODS | DS4500-ITSODS4500_A | 367.1426 | 1,057.20 | 1,424.35 | 7,617 | 0 | 1,963 | 5.0751 | 8.7225 | 13.7976 | 14.1549 | 8.4486 | 9.9194 |
| 8 | 2006/09/28 08:31 | DS4500-ITSODS | DS4500-ITSODS4500_A | 388.1927 | 1,051.19 | 1,439.38 | 7,696 | 0 | 2,075 | 5.3697 | 8.5971 | 13.9668 | 14.1645 | 8.3747 | 9.9362 |

*Figure 6-79 Performance graph for a subsystem*

Looking at the performance graph in Figure 6-79, we notice a consistent I/O rate throughout the time frame. For an I/O graph such as this one, we look for spikes and then drill deeper into the subsystem to determine the causes of these spikes.

It is quite easy to use Excel to produce performance graphs. You can produce performance graphs for other subsystems, switches, and SVCs.

## 6.5 Repocopy

*Repocopy* enables you to make a copy of your TPC database repository.

> **Note:** Do not use repocopy as a backup mechanism or backup tool. It is merely for troubleshooting and analysis.

**Repocopy** allows you to export all of the tables of the database repository. It also allows you to import the tables for another database. **Repocopy** is also a very useful tool when development needs to review your database tables for programming bugs and error conditions.

Never import **repocopy** into an existing production TPC database, because this will cause database corruption.

> **Note:** All tables are exported into a text file format.

## 6.5.1 How to run repocopy

The steps are:

1. You need to run the tool from the following directory:
   - For Windows:

     C:\Program Files\IBM\TPC\data\server\tools\
   - For UNIX or Linux:

     /opt/IBM/TPC/data/server/tools
2. Run **`repocopy.bat`** (Windows) or **`repocopy.sh`** (UNIX or Linux) as shown in Figure 6-80.

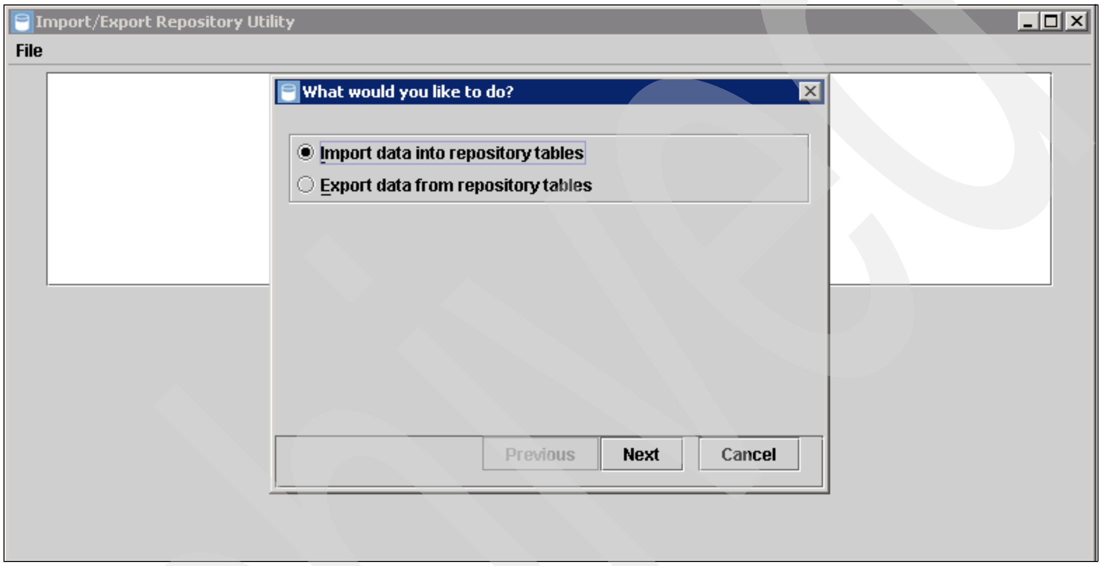

● Import data into repository tables

○ Export data from repository tables

*Figure 6-80 Run Repocopy and select the function you want*

3. Select which function you want, either Import Data or Export Data. Click **Next**.

### Exporting the data repository

To export the data repository:

1. To export the data from the repository table, run the **`repocopy`** command.
2. Select **Export data from repository tables** (See Figure 6-81 on page 304).

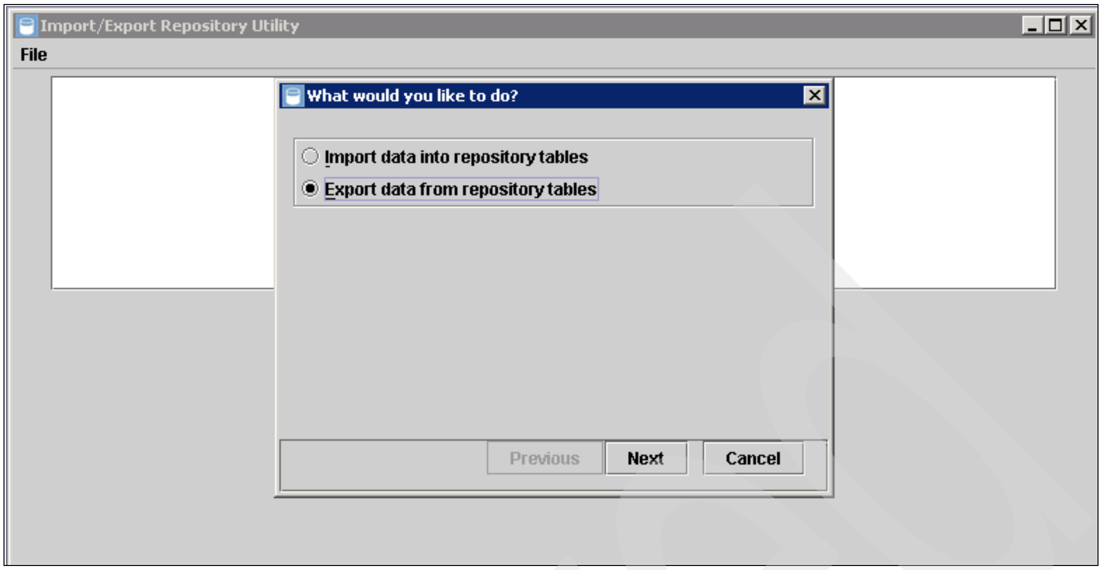

○ Import data into repository tables
● Export data from repository tables

*Figure 6-81   Select the Export data from repository tables function*

3. Click **Next** (see Figure 6-82).

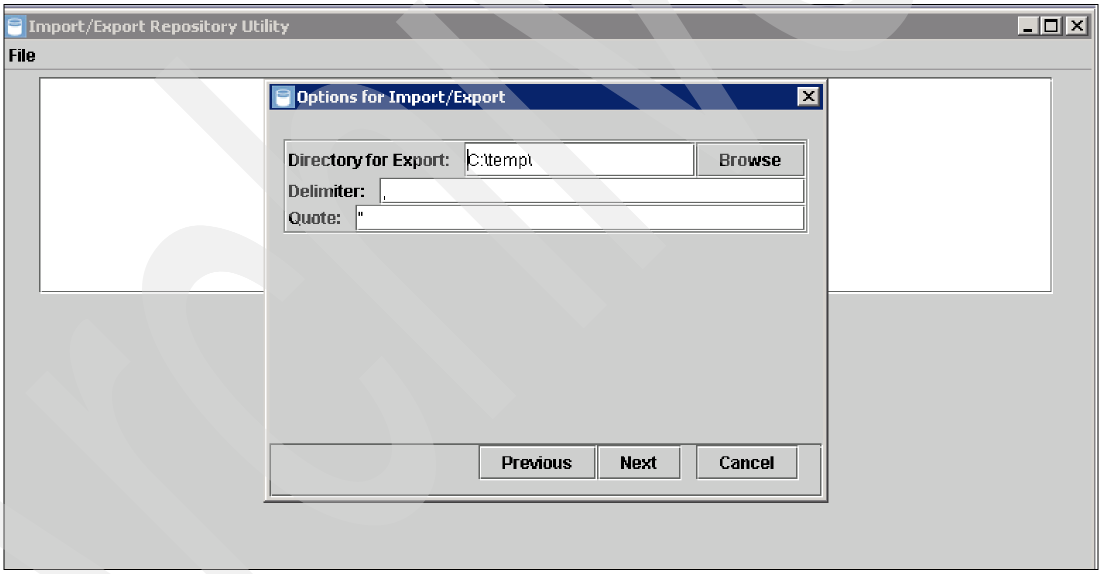

Directory for Export: C:\temp\
Delimiter: ,
Quote: "

*Figure 6-82   Options for Import/Export*

4. Enter the information requested for the following fields:
   - **Directory for Export:** Enter the directory where you want to export the repository.

     > **Note:** Always enter the location of the file to which you want to export or browse to a level of the folder that is one level higher than the export location.

   - **Delimiter:** Enter a delimiter character (the default is a comma).
   - **Quote:** Enter the symbol that will contain the string data (double quotation marks are the default).
5. Click **Next**. The Connection Properties panel displays. See Figure 6-83 on page 305.

6. In this step, the utility reads the information that is detected in the server.config file. You can locate this server configuration file under the \tpc\data\config directory (see Figure 6-83).

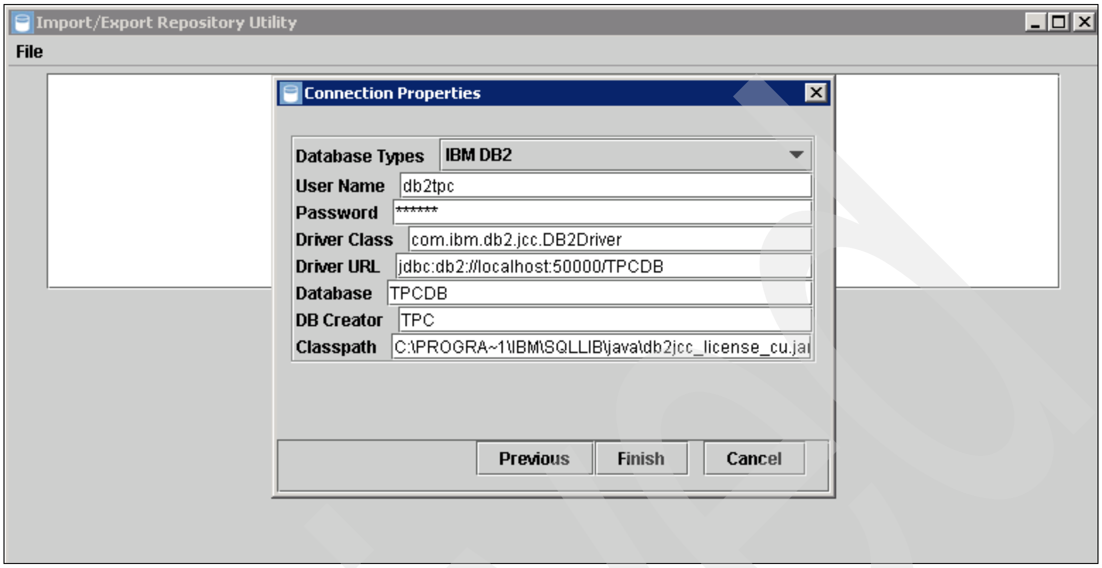

Database Types: IBM DB2
User Name: db2tpc
Password: ******
Driver Class: com.ibm.db2.jcc.DB2Driver
Driver URL: jdbc:db2://localhost:50000/TPCDB
Database: TPCDB
DB Creator: TPC
Classpath: C:\PROGRA~1\IBM\SQLLIB\java\db2jcc_license_cu.jar

*Figure 6-83 Enter the database connection properties information*

**Note:** You can export data from another database. You merely edit the fields in the Connection Properties window.

7. Click **Finish**. The database information displays (see Figure 6-84).

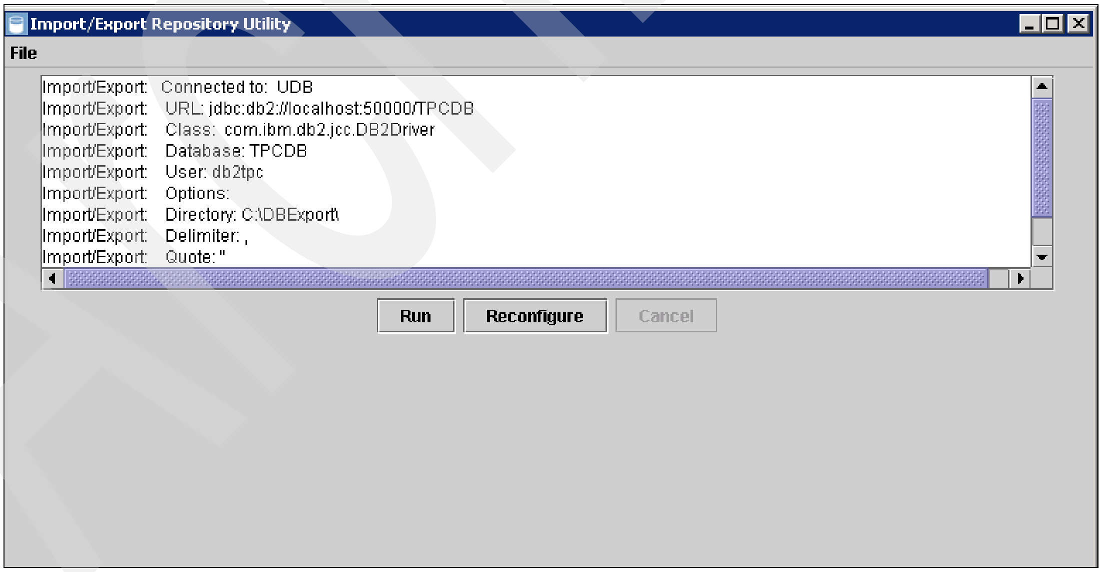

```
Import/Export:  Connected to:  UDB
Import/Export:   URL: jdbc:db2://localhost:50000/TPCDB
Import/Export:   Class:  com.ibm.db2.jcc.DB2Driver
Import/Export:   Database: TPCDB
Import/Export:   User: db2tpc
Import/Export:   Options:
Import/Export:   Directory: C:\DBExport\
Import/Export:   Delimiter: ,
Import/Export:   Quote: "
```

*Figure 6-84 Database information displays*

8. Click **Run.** TPC provides you with the option to reconfigure your selection. A progress bar displays, and then, a list of tables that are being exported displays (see Figure 6-85 on page 306).

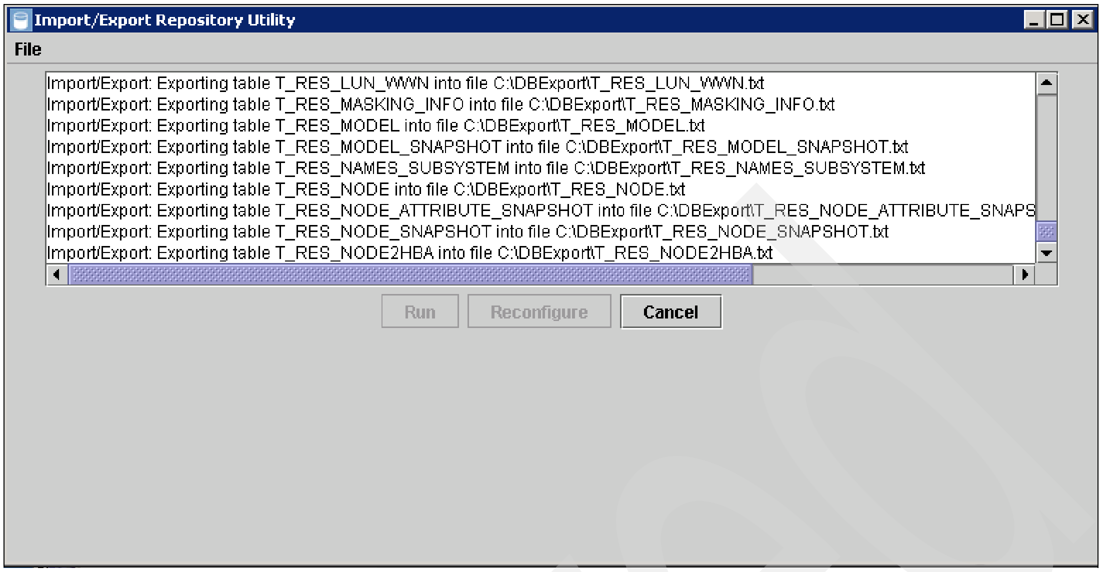

```
Import/Export: Exporting table T_RES_LUN_WWN into file C:\DBExport\T_RES_LUN_WWN.txt
Import/Export: Exporting table T_RES_MASKING_INFO into file C:\DBExport\T_RES_MASKING_INFO.txt
Import/Export: Exporting table T_RES_MODEL into file C:\DBExport\T_RES_MODEL.txt
Import/Export: Exporting table T_RES_MODEL_SNAPSHOT into file C:\DBExport\T_RES_MODEL_SNAPSHOT.txt
Import/Export: Exporting table T_RES_NAMES_SUBSYSTEM into file C:\DBExport\T_RES_NAMES_SUBSYSTEM.txt
Import/Export: Exporting table T_RES_NODE into file C:\DBExport\T_RES_NODE.txt
Import/Export: Exporting table T_RES_NODE_ATTRIBUTE_SNAPSHOT into file C:\DBExport\T_RES_NODE_ATTRIBUTE_SNAPS
Import/Export: Exporting table T_RES_NODE_SNAPSHOT into file C:\DBExport\T_RES_NODE_SNAPSHOT.txt
Import/Export: Exporting table T_RES_NODE2HBA into file C:\DBExport\T_RES_NODE2HBA.txt
```

*Figure 6-85 The export utility progress bar*

9. The display window displays a message when the export completes (see Figure 6-86).

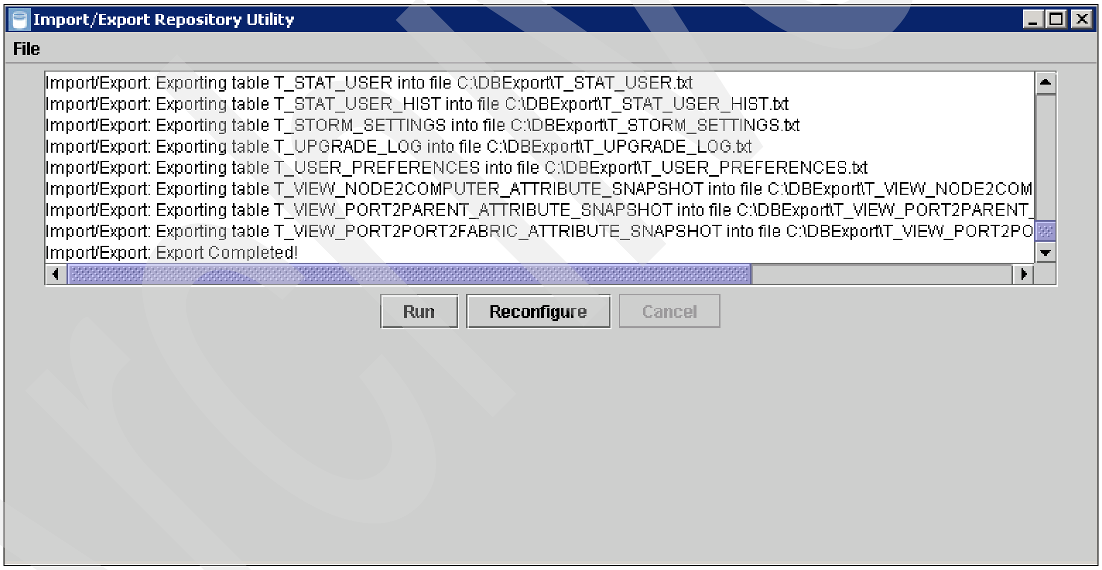

```
Import/Export: Exporting table T_STAT_USER into file C:\DBExport\T_STAT_USER.txt
Import/Export: Exporting table T_STAT_USER_HIST into file C:\DBExport\T_STAT_USER_HIST.txt
Import/Export: Exporting table T_STORM_SETTINGS into file C:\DBExport\T_STORM_SETTINGS.txt
Import/Export: Exporting table T_UPGRADE_LOG into file C:\DBExport\T_UPGRADE_LOG.txt
Import/Export: Exporting table T_USER_PREFERENCES into file C:\DBExport\T_USER_PREFERENCES.txt
Import/Export: Exporting table T_VIEW_NODE2COMPUTER_ATTRIBUTE_SNAPSHOT into file C:\DBExport\T_VIEW_NODE2COM
Import/Export: Exporting table T_VIEW_PORT2PARENT_ATTRIBUTE_SNAPSHOT into file C:\DBExport\T_VIEW_PORT2PARENT_
Import/Export: Exporting table T_VIEW_PORT2PORT2FABRIC_ATTRIBUTE_SNAPSHOT into file C:\DBExport\T_VIEW_PORT2PO
Import/Export: Export Completed!
```

*Figure 6-86 Display window message indicating that the export has completed*

### Importing the data repository

To import the data repository:

1. To import the data from the repository table, run the **`repocopy`** command.
2. Select **Import data into repository tables** (see Figure 6-87 on page 307).

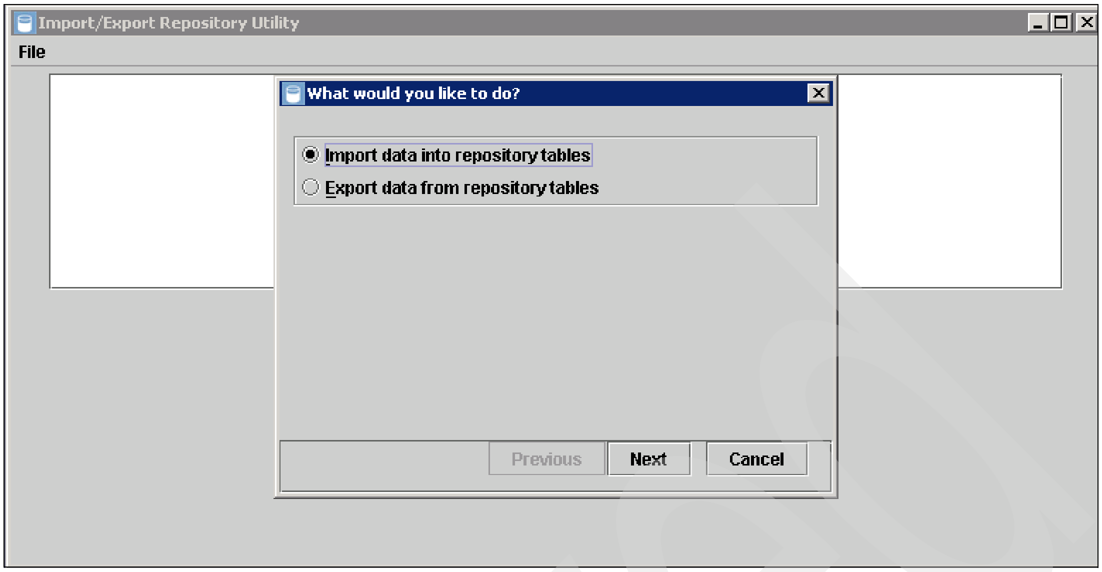

● Import data into repository tables
○ Export data from repository tables

*Figure 6-87 Select Import data from repository tables*

3. Click **Next** to continue.
4. Enter the Options for Import fields (see Figure 6-88 on page 308):
   - **Directory for Import:** Enter the directory where the saved repository tables are stored.

     > **Note:** Always type the location of the file that you want to import or browse to one level of the folder higher than the import location.

   - **Delimiter:** Enter a delimiter character that has been used in the export utility (the default is a comma).
   - **Quote:** Enter the symbol that will contain the string data that used in the export utility (double quotation marks are the default).
   - **Delete before Inserting:** Check this option if you want to delete any existing data.

> **Tip:** Always check the Delete before Inserting option, because this provides you with a clean database into which to import the exported data.

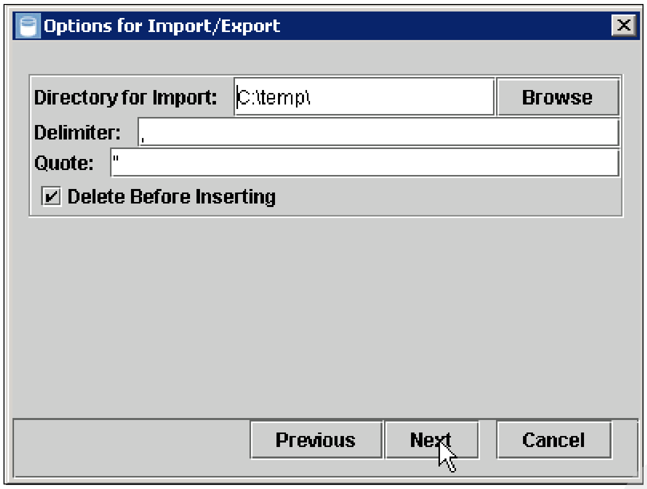

Directory for Import: C:\temp\
Delimiter: ,
Quote: "
☑ Delete Before Inserting

*Figure 6-88 Enter the Import fields*

5. Click **Next**. The Connection Properties panel displays (see Figure 6-89).

   Enter the following fields:

   – Database Types
   – User Name
   – Password
   – Driver Class
   – Driver URL
   – Database
   – DB Creator
   – Classpath

> **Note:** All of these fields are populated by default from DB2 configuration files. Only modify these fields if required.

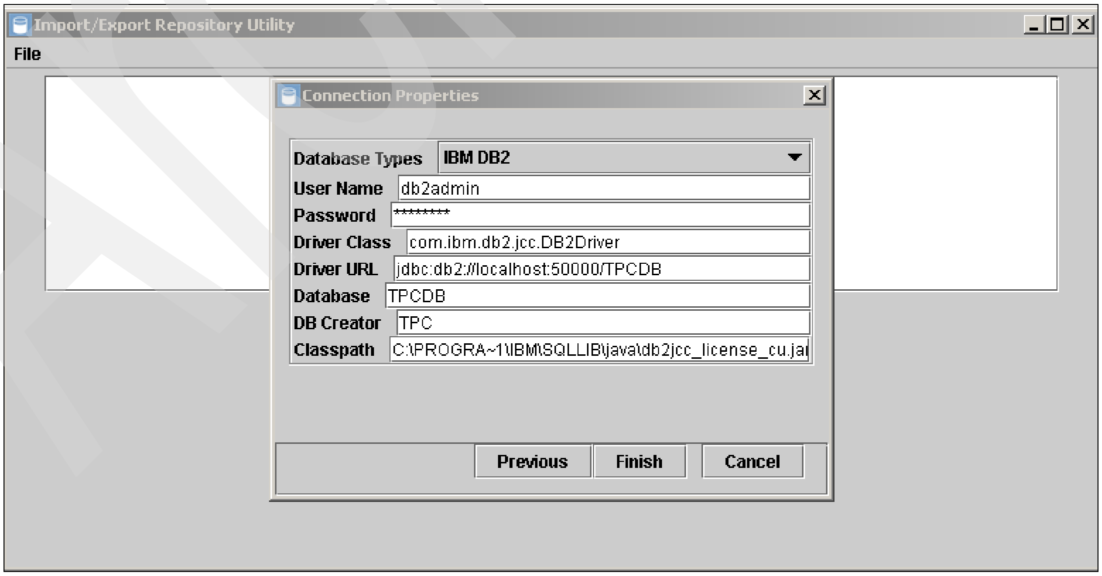

Database Types: IBM DB2
User Name: db2admin
Password: *********
Driver Class: com.ibm.db2.jcc.DB2Driver
Driver URL: jdbc:db2://localhost:50000/TPCDB
Database: TPCDB
DB Creator: TPC
Classpath: C:\PROGRA~1\IBM\SQLLIB\java\db2jcc_license_cu.jar

*Figure 6-89 Enter the Connection Properties for importing the repository*

**Important:** Always enter the database information for the new database that you created and not the information of the exported data.

6. Click **Finish**. The Import/Export Repository Utility displays the information entered in the previous panel (see Figure 6-90).

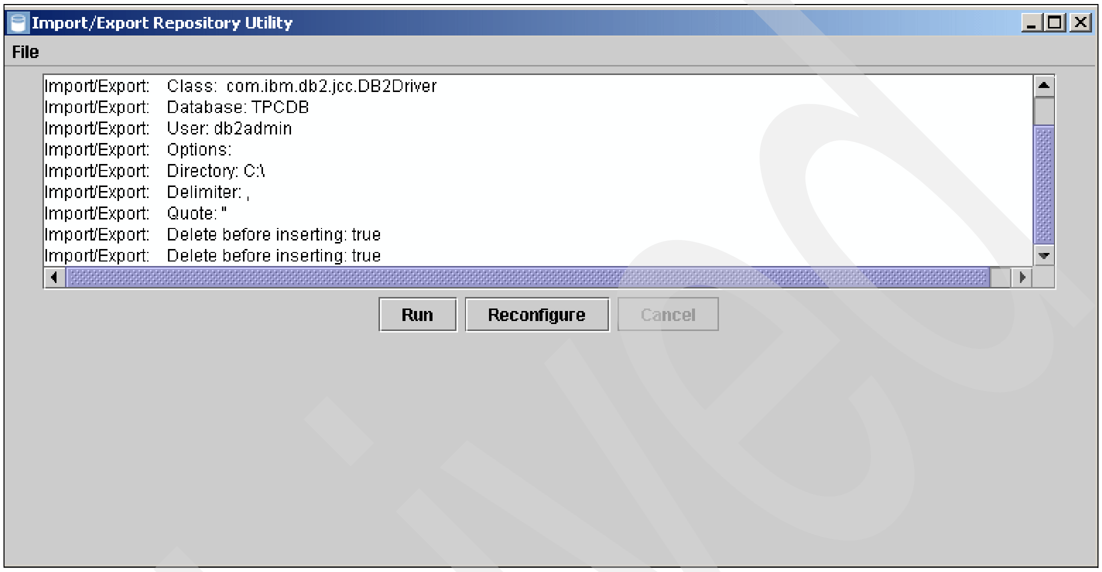

```
Import/Export:   Class:  com.ibm.db2.jcc.DB2Driver
Import/Export:   Database: TPCDB
Import/Export:   User: db2admin
Import/Export:   Options:
Import/Export:   Directory: C:\
Import/Export:   Delimiter: ,
Import/Export:   Quote: "
Import/Export:   Delete before inserting: true
Import/Export:   Delete before inserting: true
```

*Figure 6-90 Display of the information that was entered in the previous panel*

7. Click **Run** to start the import. A window displays with a progress bar of the import. After the import is complete, the window displays that the import has completed (see Figure 6-91).

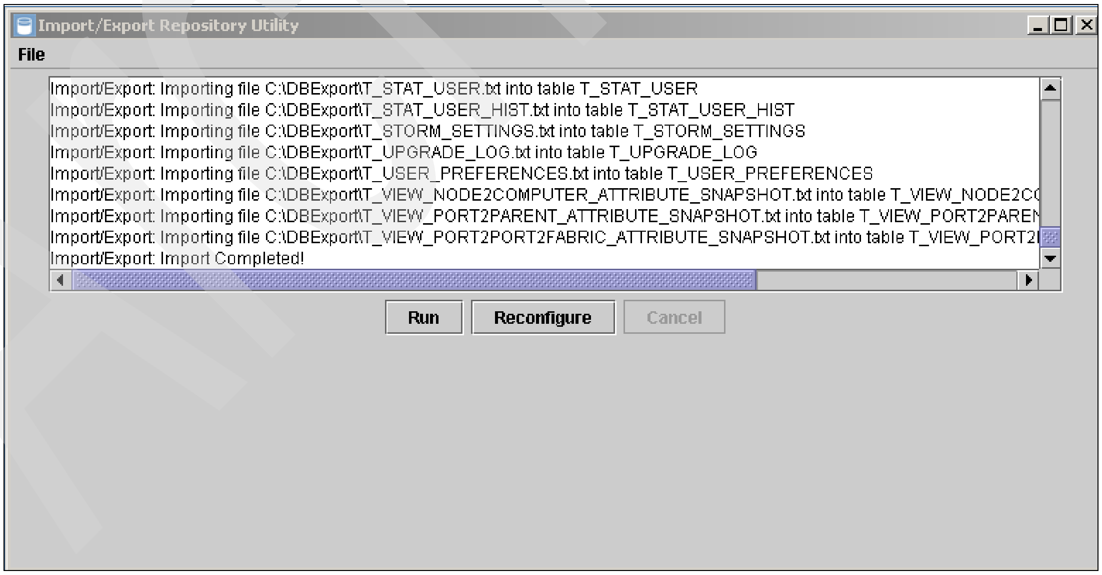

```
Import/Export: Importing file C:\DBExport\T_STAT_USER.txt into table T_STAT_USER
Import/Export: Importing file C:\DBExport\T_STAT_USER_HIST.txt into table T_STAT_USER_HIST
Import/Export: Importing file C:\DBExport\T_STORM_SETTINGS.txt into table T_STORM_SETTINGS
Import/Export: Importing file C:\DBExport\T_UPGRADE_LOG.txt into table T_UPGRADE_LOG
Import/Export: Importing file C:\DBExport\T_USER_PREFERENCES.txt into table T_USER_PREFERENCES
Import/Export: Importing file C:\DBExport\T_VIEW_NODE2COMPUTER_ATTRIBUTE_SNAPSHOT.txt into table T_VIEW_NODE2C
Import/Export: Importing file C:\DBExport\T_VIEW_PORT2PARENT_ATTRIBUTE_SNAPSHOT.txt into table T_VIEW_PORT2PARE
Import/Export: Importing file C:\DBExport\T_VIEW_PORT2PORT2FABRIC_ATTRIBUTE_SNAPSHOT.txt into table T_VIEW_PORT2
Import/Export: Import Completed!
```

*Figure 6-91 The utility verifies the progress and the success of the task*

8. The **`repocopy`** utility also has a DOS window running in the background to verify the steps of the import (see Figure 6-92 on page 310).

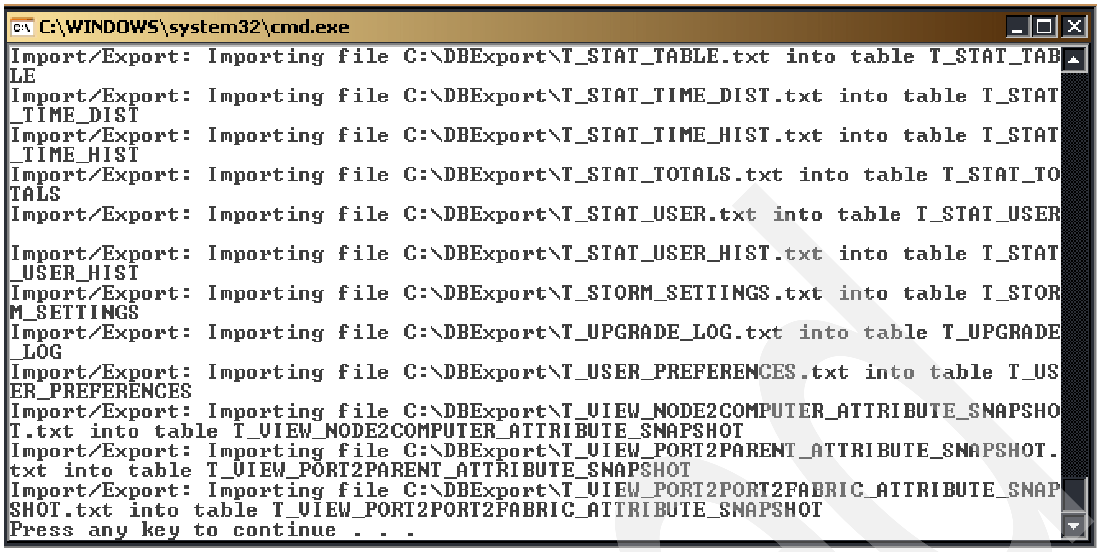

```
Import/Export: Importing file C:\DBExport\T_STAT_TABLE.txt into table T_STAT_TAB
LE
Import/Export: Importing file C:\DBExport\T_STAT_TIME_DIST.txt into table T_STAT
_TIME_DIST
Import/Export: Importing file C:\DBExport\T_STAT_TIME_HIST.txt into table T_STAT
_TIME_HIST
Import/Export: Importing file C:\DBExport\T_STAT_TOTALS.txt into table T_STAT_TO
TALS
Import/Export: Importing file C:\DBExport\T_STAT_USER.txt into table T_STAT_USER

Import/Export: Importing file C:\DBExport\T_STAT_USER_HIST.txt into table T_STAT
_USER_HIST
Import/Export: Importing file C:\DBExport\T_STORM_SETTINGS.txt into table T_STOR
M_SETTINGS
Import/Export: Importing file C:\DBExport\T_UPGRADE_LOG.txt into table T_UPGRADE
_LOG
Import/Export: Importing file C:\DBExport\T_USER_PREFERENCES.txt into table T_US
ER_PREFERENCES
Import/Export: Importing file C:\DBExport\T_VIEW_NODE2COMPUTER_ATTRIBUTE_SNAPSHO
T.txt into table T_VIEW_NODE2COMPUTER_ATTRIBUTE_SNAPSHOT
Import/Export: Importing file C:\DBExport\T_VIEW_PORT2PARENT_ATTRIBUTE_SNAPSHOT.
txt into table T_VIEW_PORT2PARENT_ATTRIBUTE_SNAPSHOT
Import/Export: Importing file C:\DBExport\T_VIEW_PORT2PORT2FABRIC_ATTRIBUTE_SNAP
SHOT.txt into table T_VIEW_PORT2PORT2FABRIC_ATTRIBUTE_SNAPSHOT
Press any key to continue . . .
```

*Figure 6-92 DOS window running in the background*

## 6.6 Brocade SMI-S agent management

The Brocade Storage Management Initiative Specification (SMI-S) agent provides access to the management facilities in Brocade fabrics.

In the ever changing SAN environment, clients constantly increase the number of SAN switches and fabrics in their environment, therefore, these switches and fabrics need to be monitored and managed from a central reporting and management tool. TPC provides you with this functionality.

The most commonly asked questions are:

- How do you add additional switches for the TPC server to monitor and manage after you have already installed your Brocade SMI-S agent?
- Which Brocade SMI-S agent logs do you need to collect if you are encountering problems?
- How do you manage the security authentication between the CIMOM and the TPC server?

The Brocade SMI-S agent has a built-in utility to assist in these tasks. The utility is called the *Configuration Tool*.

> **Note:** These features are only available on Brocade SMI-S agent Version 110.4.0. For more information, go to:
>
> `http://www.brocade.com/support/SMIAGENT.jsp`

To locate the Configuration tool, go to:

- For Windows:

  <*SMIAgent*>\agent\server\jserver\bin\
- For Linux and UNIX:

  <*SMIAgent*>/agent/server/jserver/bin

### How to run ProviderUtil

On Windows:

1. Open a command window.
2. Set your system to the path C:\SMIAgent\agent\server\jserver\bin
3. Run the Configurationtool.bat file.

On UNIX:

1. Open a command window.
2. Set your system to the path opt/SMIAgent/agent/server/jserver/bin
3. Run the Configurationtool.sh file.

This opens the Brocade SMI Agent Configuration Tool as shown in Figure 6-93.

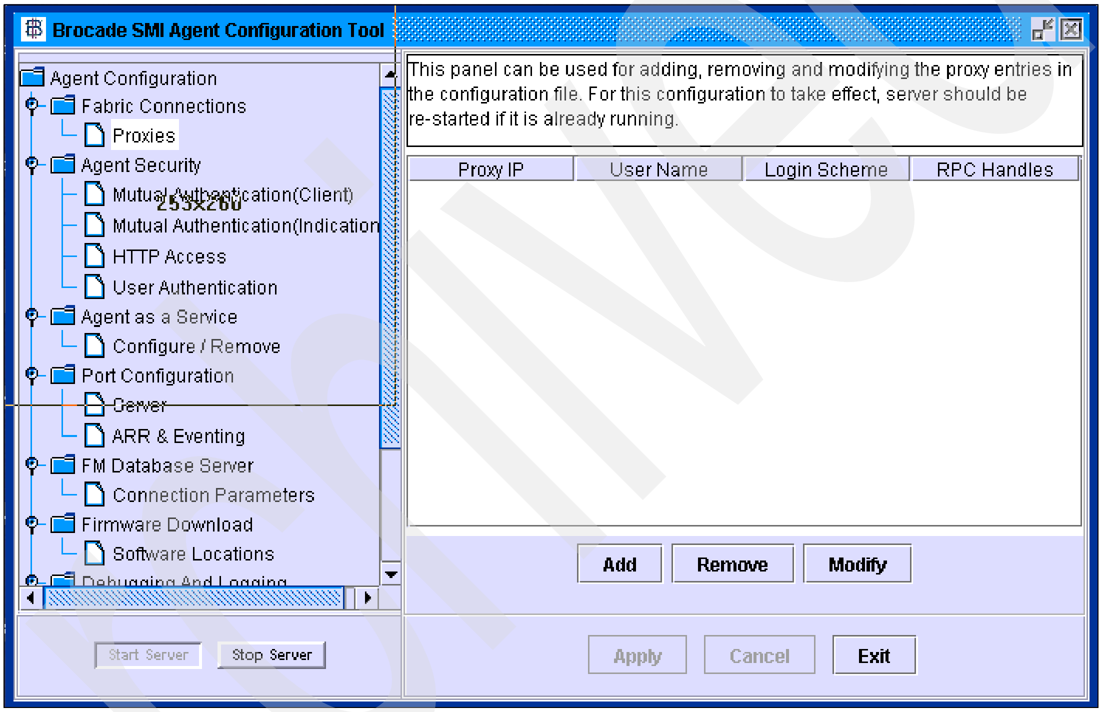

This panel can be used for adding, removing and modifying the proxy entries in the configuration file. For this configuration to take effect, server should be re-started if it is already running.

| Proxy IP | User Name | Login Scheme | RPC Handles |
|---|---|---|---|

*Figure 6-93 Brocade SMI Agent Configuration Tool*

## 6.6.1 Add additional switches to SMI Agent

In previous versions to Version 110.4.0, you needed to update the provider.xml file to add an additional switch. Brocade has built this function into the configuration tool along with other functions.

### How to run

To run the Configuration Tool:

1. Launch the Configuration Tool.
2. The Brocade SMI Agent Configuration Tool window opens. Select **Proxies** if you have already added other switches, they display in the right panel (see Figure 6-94 on page 312).

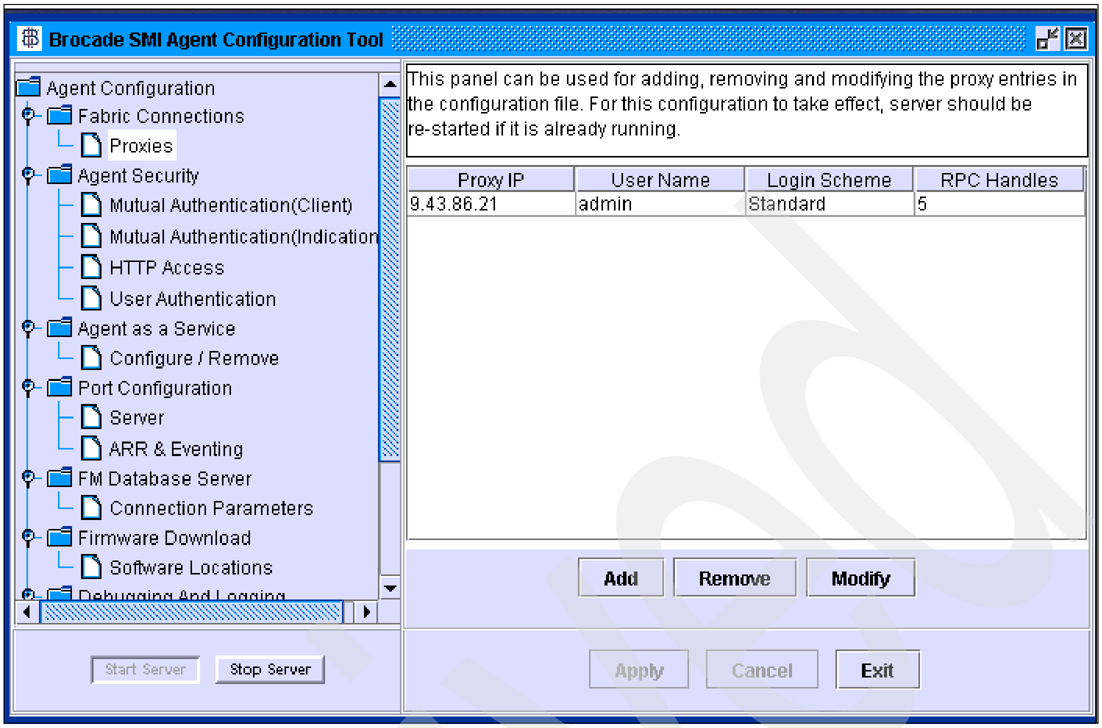

This panel can be used for adding, removing and modifying the proxy entries in the configuration file. For this configuration to take effect, server should be re-started if it is already running.

| Proxy IP | User Name | Login Scheme | RPC Handles |
|---|---|---|---|
| 9.43.86.21 | admin | Standard | 5 |

*Figure 6-94 Brocade SMI Agent with added proxies*

3. To add additional switches (proxies), click **Add**. A new window with the proxy configuration connections appears. Here, you enter the following information (see Figure 6-95):

   – Proxy (switch) IP

   – User name (of the switch)

   – Password (of the switch)

   – Login-Scheme: **Standard** (default)

   – No.of RPC Handles: **5** (default)

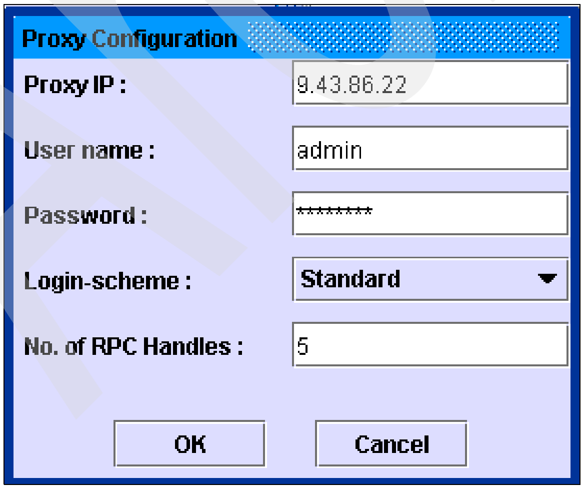

| | |
|---|---|
| Proxy IP : | 9.43.86.22 |
| User name : | admin |
| Password : | ******** |
| Login-scheme : | Standard |
| No. of RPC Handles : | 5 |

*Figure 6-95 Add switch details*

- Click **OK.**

- Click **Apply**. This adds the switch to the configuration. For the configuration changes to take effect, we stop and start the Brocade SMI Agent service. You can use the Stop and Start Server options available on the lower left of the tool panel, but we recommend that you exit the tool, which reminds you if you need to restart the Agent. Select **No** and restart the service manually (see Figure 6-96).

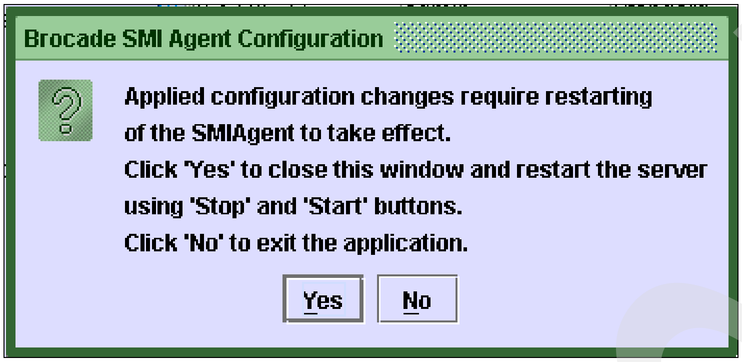
Brocade SMI Agent Configuration

Applied configuration changes require restarting
of the SMIAgent to take effect.
Click 'Yes' to close this window and restart the server
using 'Stop' and 'Start' buttons.
Click 'No' to exit the application.

*Figure 6-96 Select No to exit*

**Note:** By using the Configuration Tool, the provider.xml file gets updated automatically. All switch passwords are encrypted.

## 6.6.2 Collect logs for analysis

The Configuration Tool provides you with the functionality of gathering logs and other information that you need to debug the Brocade SMI Agent.

The tool collects the following information about the Agent:

- Provider.xml
- Jserver.properties
- CIMOM and provider log files
- SMIAgentconfig.xml
- cimom.properties
- Operating system name, version, and agent information (systemInfo.txt)

All of the logs that are created are stored in an SMISupportFiles.zip file. The location of the .zip file is defined by the users prior to running this process. You must run this command from the machine on which the Brocade SMI Agent is installed.

### How to run

The steps to gathering the logs are:

1. Launch the Configuration Tool.
2. The Brocade SMI Agent Configuration Tool window opens. Select **Agent Support Show** → **Collect Information** (see Figure 6-97 on page 314).

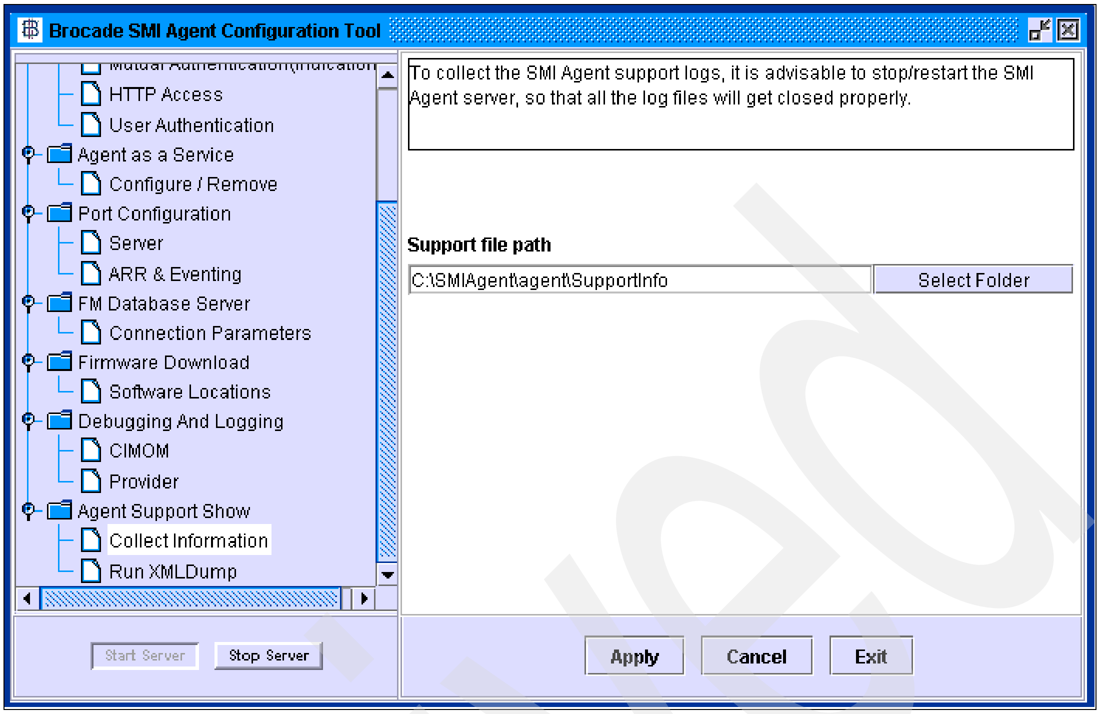

To collect the SMI Agent support logs, it is advisable to stop/restart the SMI Agent server, so that all the log files will get closed properly.

Support file path

C:\SMIAgent\agent\SupportInfo

*Figure 6-97 Brocade SMI Agent Collect Information*

3. In the right pane, you can select the folder where you want to put the output to be saved. Select the location.
4. Click **Apply** to generate the .zip file.

**Note:** The tool also allows you to generate an XML dump of the leaf classes.

## 6.6.3 How to enable Brocade SMI-S authentication

One benefit of enhanced security is that exclusive individuals control to access to the SMI-S agent. Security settings for the security areas are supported and set by the Security Configuration tool.

The security level that we set only allows local users that have an account (domain or local) access to the SMI-S agent. A *local user* is an account that can be granted permissions and rights from your computer. We need to ensure that this user account is created on the computer. Figure 6-98 on page 315 shows that we have created the local user BrocadeCIM to use in this example.

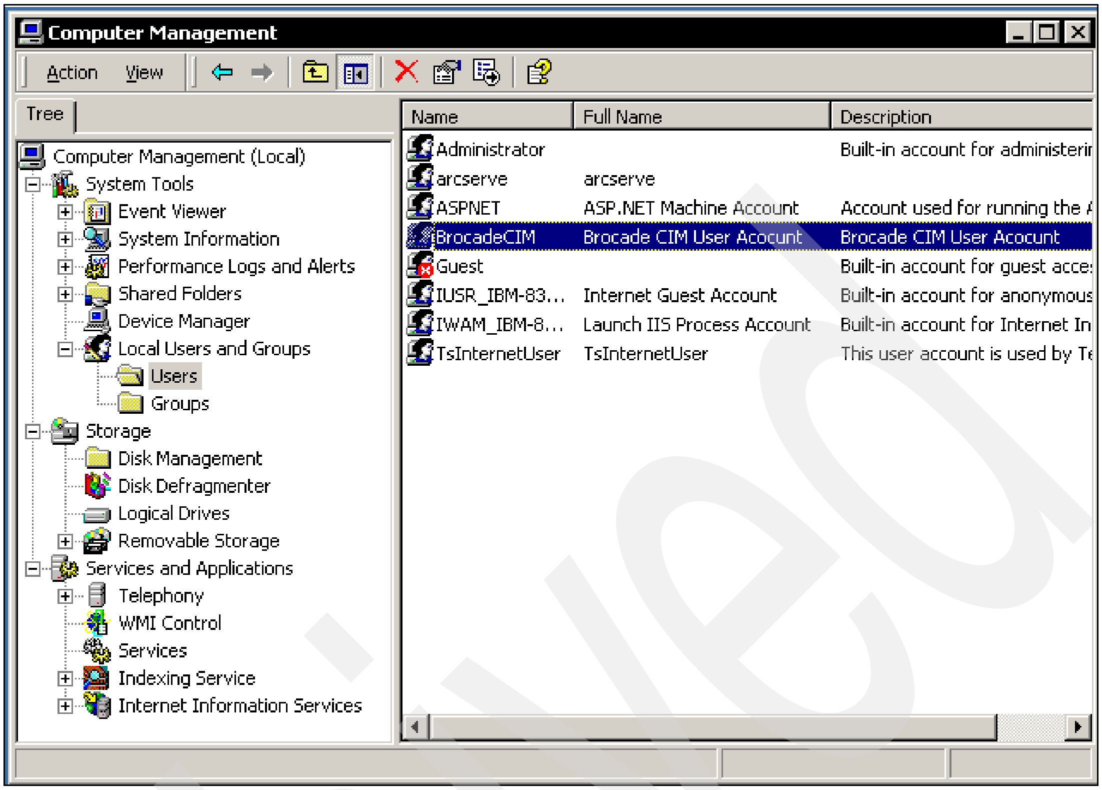

| Name | Full Name | Description |
|---|---|---|
| Administrator | | Built-in account for administerir |
| arcserve | arcserve | |
| ASPNET | ASP.NET Machine Account | Account used for running the A |
| BrocadeCIM | Brocade CIM User Account | Brocade CIM User Account |
| Guest | | Built-in account for guest acce: |
| IUSR_IBM-83... | Internet Guest Account | Built-in account for anonymous |
| IWAM_IBM-8... | Launch IIS Process Account | Built-in account for Internet In |
| TsInternetUser | TsInternetUser | This user account is used by Te |

*Figure 6-98 List of local users*

## How to run

To run user authentication:

1. Launch the Configuration Tool.
2. The Brocade SMI Agent Configuration Tool window opens. Select **Agent Security** → **User Authentication** (see Figure 6-99 on page 316).

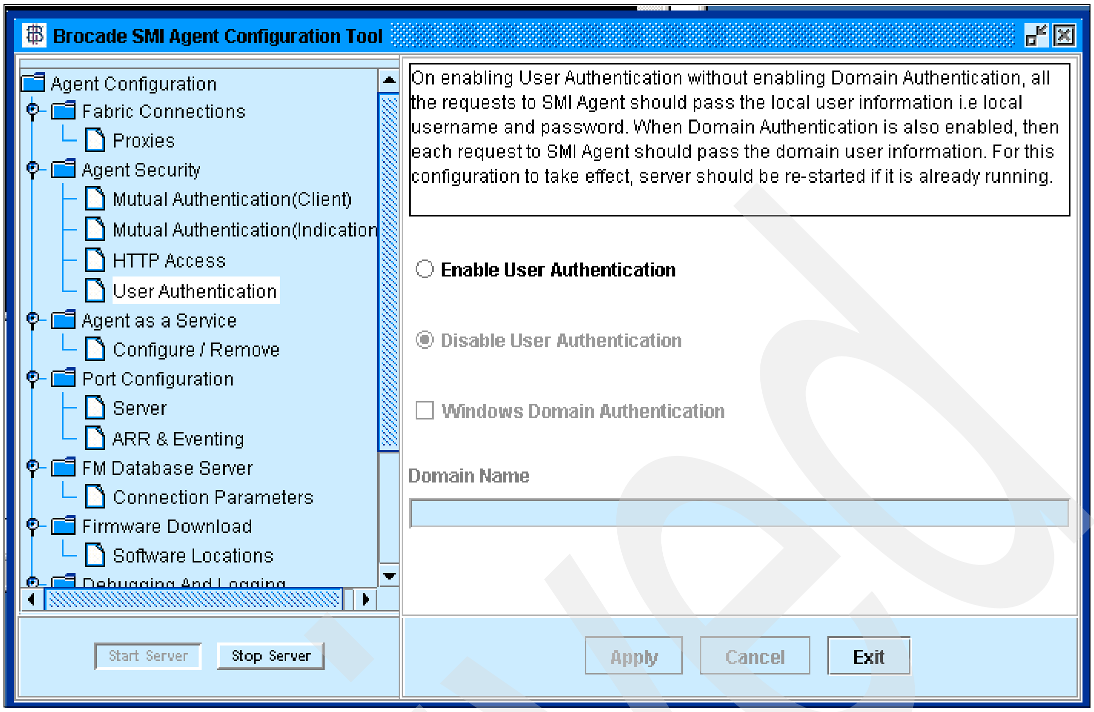

On enabling User Authentication without enabling Domain Authentication, all the requests to SMI Agent should pass the local user information i.e local username and password. When Domain Authentication is also enabled, then each request to SMI Agent should pass the domain user information. For this configuration to take effect, server should be re-started if it is already running.

○ Enable User Authentication

● Disable User Authentication

□ Windows Domain Authentication

Domain Name

*Figure 6-99 Brocade SMI Agent User Authentication*

3. By default, **Disable User Authentication** is set to allow anyone access to the SMI-S Agent. You need to change in order to enable user authentication. Select **Enable User Authentication,** and click **Apply**. A message appears indicating that you have enabled user authentication, and you need to supply the user name and password as shown in Figure 6-100. Click **OK** to continue.

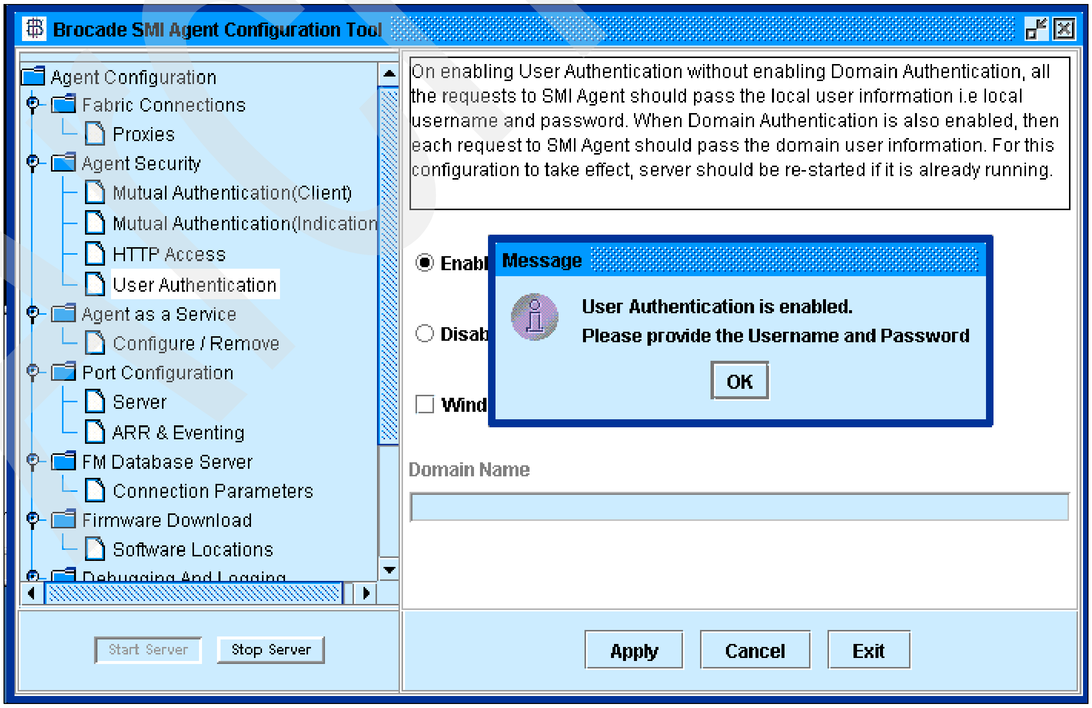

On enabling User Authentication without enabling Domain Authentication, all the requests to SMI Agent should pass the local user information i.e local username and password. When Domain Authentication is also enabled, then each request to SMI Agent should pass the domain user information. For this configuration to take effect, server should be re-started if it is already running.

● Enabl

○ Disab

□ Wind

User Authentication is enabled.
Please provide the Username and Password

Domain Name

*Figure 6-100 Enable User Authentication*

4. A User Authentication window prompts you for a user name and password. Enter the user name and password. Click **OK**.

**Note:** Although you cannot see to verify the user name and password that you enter, be very careful to enter them correctly.

5. You need to restart the Brocade SMI Agent service, exit the tool, and restart the service manually for the changes to take effect:

   a. This now allows you to add the Brocade CIMOM to your TPC server with user authentication now enabled (see Figure 6-101).

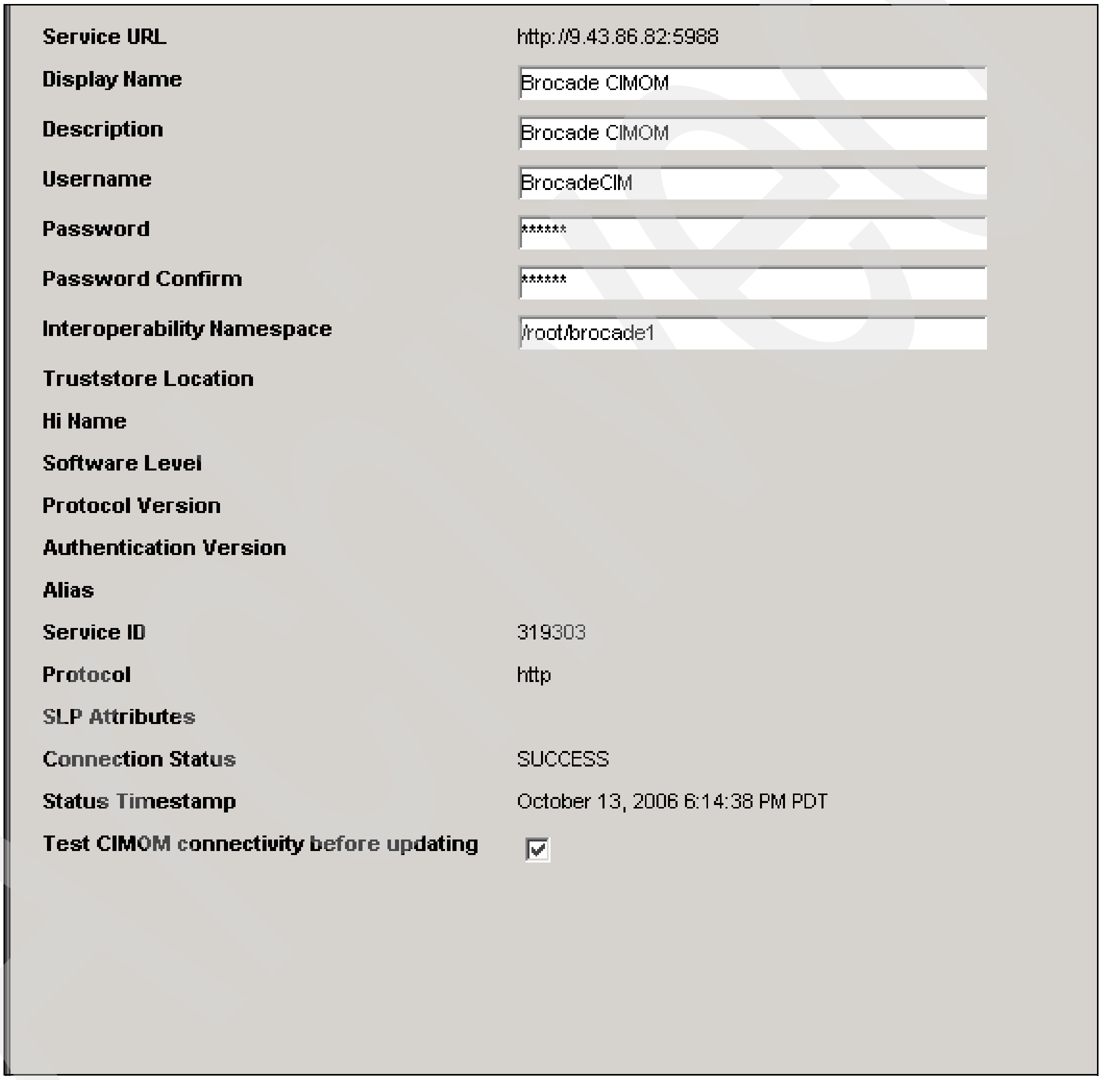

| | |
|---|---|
| **Service URL** | http://9.43.86.82:5988 |
| **Display Name** | Brocade CIMOM |
| **Description** | Brocade CIMOM |
| **Username** | BrocadeCIM |
| **Password** | ****** |
| **Password Confirm** | ****** |
| **Interoperability Namespace** | /root/brocade1 |
| **Truststore Location** | |
| **Hi Name** | |
| **Software Level** | |
| **Protocol Version** | |
| **Authentication Version** | |
| **Alias** | |
| **Service ID** | 319303 |
| **Protocol** | http |
| **SLP Attributes** | |
| **Connection Status** | SUCCESS |
| **Status Timestamp** | October 13, 2006 6:14:38 PM PDT |
| **Test CIMOM connectivity before updating** | ☑ |

*Figure 6-101 Brocade CIMOM with user authentication*

**Tip:** If you have a Brocade CIMOM already configured on your TPC server, you need to delete the CIMOM entry and recreate it with the correct user name and password to enable user authentication.

# A

# Additional material

This IBM Redbooks publication refers to additional material that you can download from the Internet as we describe here.

## Locating the Web material

The Web material associated with this IBM Redbooks publication is available in softcopy on the Internet from the IBM Redbooks publications Web server. Point your Web browser to:

`ftp://www.redbooks.ibm.com/redbooks/SG247348`

Alternatively, you can go to the IBM Redbooks publications Web site at:

`ibm.com/redbooks`

Select the **Additional materials** and open the directory that corresponds with this IBM Redbooks form number, SG247348.

## Using the Web material

The additional Web material that accompanies this IBM Redbooks publication includes the following files:

| File name | Description |
|---|---|
| **SG247348.zip** | TPCTOOL macros |

### System requirements for downloading the Web material

The following system configuration is recommended:

**Hard disk space**: 30 MB minimum

## How to use the Web material

Create a subdirectory (folder) on your workstation, and unzip the contents of the Web material .zip file into this folder.

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this IBM Redbooks publication.

## IBM Redbooks publications

For information about ordering these publications, see “How to get IBM Redbooks publications” on page 322. Note that some of the documents referenced here might be available in softcopy only.

- *IBM TotalStorage Productivity Center V3.1: The Next Generation*, SG24-7194

## Other publications

These publications are also relevant as further information sources:

- *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference,* SC09-4831
- *IBM TotalStorage Productivity Center Command-Line Interface Reference,* GC32-1777
- *IBM TotalStorage UltraScalable Tape Library 3584 Planning and Operator Guide,* GA32-0408
- *IBM TotalStorage Productivity Center Messages,* GC32-7194
- IC49190 *Need procedures for changes to IBM TotalStorage Productivity Center: DB2 password changes; uninstallation of Data agents; uninstallation of Tivoli Common Agent*

  You can locate this document at:

  http://www-1.ibm.com/support/docview.wss?rs=1133&context=SS8JB5&context=SSWQP2&dc=DB500&q1=password&uid=swg21236490&loc=en_US&cs=utf-8&lang=en
- *IBM TotalStorage Productivity Center for Data V3 - Best Practices*:
  Chapter: Changing the DB2 password for the Data Server
  Chapter: Changing the DB2 password for the Device Server
  Chapter: Changing the host authentication password for the Device Server

  You can locate this document at:

  http://www-1.ibm.com/support/docview.wss?rs=597&context=STCRLM4&context=SSMMUP&dc=DA4A10&uid=ssg1S7001491&loc=en_US&cs=utf-8&lang=en

## Online resources

These Web sites are also relevant as further information sources:

- Messages link

  http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp
- TPC Web site

  http://www-03.ibm.com/servers/storage/support/software/tpc

- Brocade SMI-S Agent Web site

  http://www.brocade.com/support/SMIAGENT.jsp

## How to get IBM Redbooks publications

You can search for, view, or download IBM Redbooks publications, Redpapers, Hints, and Tips, draft publications, and Additional materials, as well as order hardcopy IBM Redbooks publications or CD-ROMs, at this Web site:

**ibm.com**/redbooks

## Help from IBM

IBM Support and downloads

**ibm.com**/support

IBM Global Services

**ibm.com**/services

# Index

Redbooks

IBM TotalStorage Productivity Center Advanced Topics

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

# IBM TotalStorage Productivity Center Advanced Topics

Redbooks

**Learn to effectively use TotalStorage Productivity Center**

**Create customized reports**

**Maintain the database repository**

You have installed and performed the basic customization of IBM TotalStorage Productivity Center. You have successfully completed performance data collection and have generated reports. But how do you best use the TotalStorage Productivity Center to manage your storage infrastructure?

This IBM Redbooks publication shows how to best set up TotalStorage Productivity Center based on the storage environment infrastructure, and then manage that storage infrastructure with TotalStorage Productivity Center. It includes experiences from client accounts and our own internal experiences. This book includes the following topics:

- TotalStorage Productivity Center installation considerations (number of servers, database placement, firewall considerations, and agent deployment)
- CIMOM management (how many are required and how to customize them)
- Performance monitoring (setting up thresholds and alerts, gathering data, and which reports to use)
- Ceating custom reports with TPCTOOL
- Maintaining the TotalStorage Productivity Center implementation (data retention, database backup, and debugging tools)

This Redbooks publication is intended for storage administrators, who are responsible for the performance and growth of the IT storage infrastructure.

## INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

## BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**
**ibm.com**/redbooks

SG24-7348-00 ISBN 0738486396