



IBM TotalStorage Productivity Center Topology Viewer

*A Scalable Master Console for Managing Storage
Infrastructure*

Table of Contents

Table of Contents	2
Introduction	3
Process	4
Scalability concerns	5
Semantic Zooming/Progressive Disclosure	5
Multi-level, multi-perspective layouts	6
Grouping/Aggregation	8
Overlays	9
Pinning	9
Synchronized Views	10
Summary	11
IBM TotalStorage Productivity Center Overview	11

Introduction

Enterprise-scale storage environments are typically very complex with tens of switches and thousands of ports interconnecting servers, with storage subsystems containing hundreds to thousands of LUNs, tape libraries and other devices. All of these are potentially distributed across several data centers and use technologies such as virtualization. In addition, the complexity of the work of storage administrators, system designers and architects increases as new infrastructure is added to support corporate information demands.

A general challenge for storage administrators is “to have the environment mapped out in your head.” Common practice is to produce large spreadsheets or maps, where updates are done manually by a few administrators as they map and track the system environment. In today’s IT environment, rapid change is causing significant storage growth, which is leading to more frequent storage provisioning. A topology viewer with live updates is required to keep up with these changes.

The challenge is to develop the right views for the various administration tasks (e.g., monitoring, diagnostics, and configuration). In some cases, administrators need to see the big picture, to quickly assess the business impact to any outages or performance bottlenecks. In other cases, details are needed on the storage components to troubleshoot underlying problems. Many administration tasks require administrators to keep an overview of the entire environment but quickly access very low-level parts of the system. For instance, while an administrator keeps an eye on the overall performance of a group of server hosts, he or she may need to quickly examine the throughput of an individual switch port. End-to-end diagnostics is crucial as administrators trace the source of a problem from a host to a switch to a storage subsystem. In diagnostics, as well as in configuration tasks, the ability to organize, analyze and display information is extremely important. The challenge organizations face is being able to find management tools that provide a clear picture of the current infrastructure, display health and performance information in an easy to understand manner, and allow storage administrators to make changes in real-time. The topology viewer of

IBM TotalStorage® Productivity Center is designed to help provide this capability, delivering a common management console to give storage administrators the ability to handle a number of tasks from a single tool, including:

- *Monitoring: Virtualizing the infrastructure and accessing health at-a-glance and logs*
- *Planning: Anticipating problem areas and trends*
- *Configuration: Providing allocation, zoning and masking*
- *Reporting: Displaying capacity, utilization and performance*
- *Problem determination: Providing aggregated status, drilling down and identifying impacted resources*

Process

Traditional storage management tools have had challenges in offering scalable views of storage infrastructures that allow organizations to see both “the forest and the trees.” Many of today’s current tools only span a limited scope of the overall environment—just the SAN fabric topology or just a particular storage component. These tools fail to give storage administrators a global view of their storage infrastructures, including the hosts, the SAN fabric, as well as a single view of storage systems. To fully manage the environment, administrators want to drill down into the details, such as looking at a poorly performing SAN switch port or a volume on a storage array. Most storage management tools do not provide this capability today without losing the context. Today’s tools either do not provide a scalable enterprise view of storage resources or they do not provide enough detail about a specific storage device.

History has demonstrated that graphical topology viewers do not scale to the complexity of typical storage administration systems. A typical problem a storage administrator might face in managing the storage infrastructure is whether the management tool can scale to handle growing amounts of data. For example, a question might be: How does a management console viewer scale for a single Cisco MDS 9509 Multilayer Director with up to 224 Fibre Channel ports in a single rack? The design principles of the topology viewer of IBM TotalStorage Productivity Center are based on addressing these scalability concerns.

Scalability concerns

The basic premise of the IBM topology viewer was to provide the right level of detail a storage administrator would need to manage the configuration of one device as well as an overall storage environment made up of hundreds, if not thousands, of devices. One particular challenge the topology viewer is designed to address is to represent a very large Fibre Channel switch at a high level as well as detailed level. This includes the ability to render each port on this switch with added status and connections to other devices without creating a cluttered view.

Semantic Zooming/Progressive Disclosure

IBM TotalStorage Productivity Center is designed to take advantage of a number of advanced techniques to display information, including semantic zooming and progressive disclosure that disclose information progressively. Semantic Zooming is a visualization technique for rendering very high-density data by adaptively changing the level of data abstraction. While graphical zooming changes the scale of the graphical representation of the objects in a view, semantic zooming changes the level of information abstraction with a corresponding representation of information at that level. It is often employed in conjunction with progressive disclosure, which provides task-specific presentation and interaction in a sequence of displays. Much of this capability was achieved in the IBM TotalStorage Productivity Center topology viewer by anticipating the steps administrators would take in completing tasks and creating displays to support completing those tasks quickly. The IBM TotalStorage Productivity Center User Interface is designed to preserve the information an administrator has viewed through layouts, similar to the way Internet Web sites provide tabs to access common tasks a user may use.

The Productivity Center topology viewer includes a number of design features, such as multi-level and multi-perspective layouts, grouping, pinning, and overlays. For instance, connections between devices are only shown when one of the end points is selected (**Figure 1**). Grouped devices, at first, reveal aggregate information, and additional details are shown when expanded (**Figure 2**). The following sections look at key features of the topology viewer that we believe are mostly unique for storage management tools today.

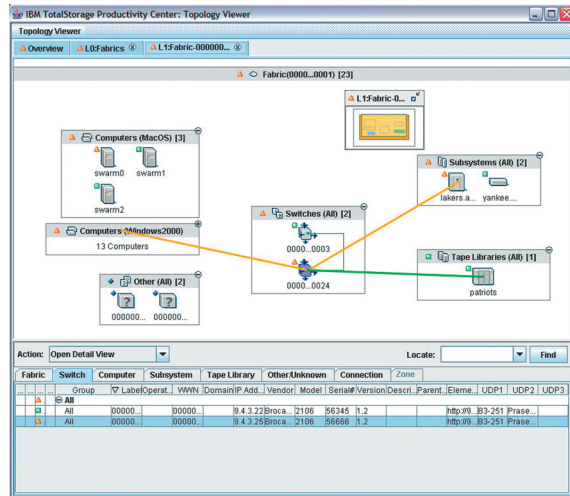


Figure 1 – IBM TotalStorage Productivity Center Topology Viewer includes the ability to display information on varying levels. Connections are drawn when an end point is selected. Connections to collapsed groups are grouped (indicated with thicker line) and show aggregated status.

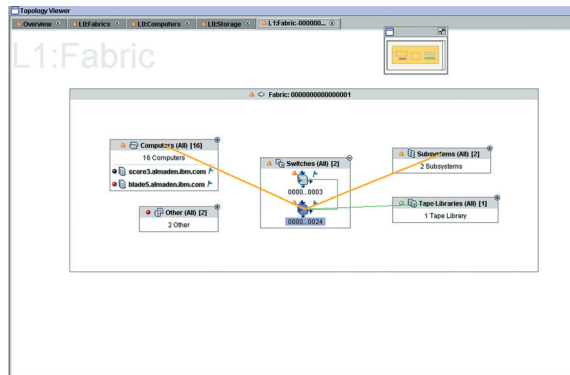


Figure 2 – A number of computers grouped by location (e.g. B3-251) shown in collapsed and expanded states. The health of individual computers is aggregated to determine the health of the group.

Multi-level, multi-perspective layouts

The viewer provides multiple views (server, fabric, and storage-centric) of the system topology with varying levels of abstraction (overview, group, single devices). Initially, users are shown an overview of their entire systems environment, where devices are grouped by type (**Figure 3**).

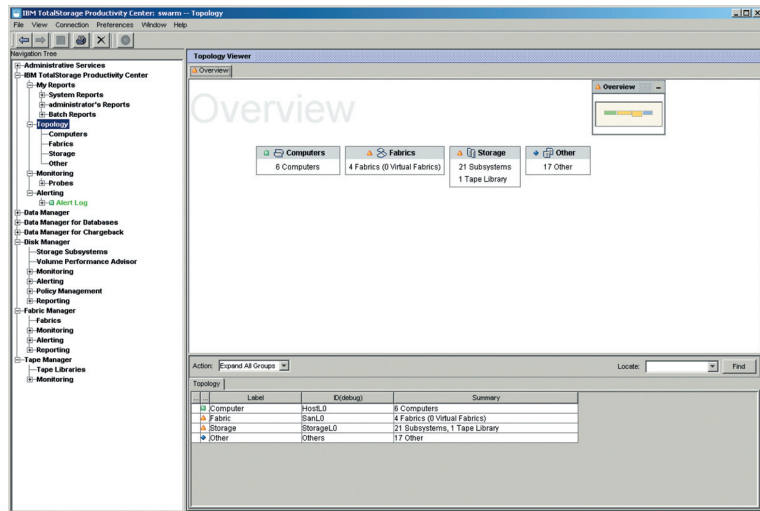


Figure 3 – Overview level with several pinned items.

In the event of a problem, users can use aggregated status to trace information to troubled devices by drilling down further into the environment—such as initially to fabric groups (Figure 4) and then eventually to a switch in a fabric. The administrator can quickly recall a specific view without having to back in/out of panel hierarchies or lose context.

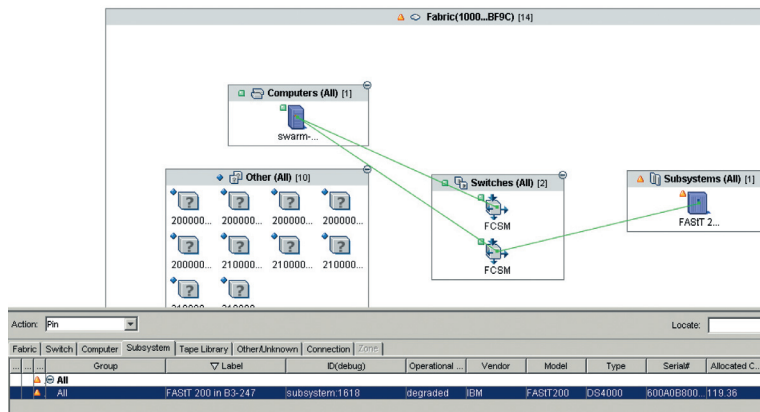


Figure 4 – Drilling into “Fabrics” showing groups of fabrics with constituent entities summarized.

A key challenge was to design the display of information at each level and perspective in such a way as to support users in their tasks, as best as possible, while at the same time reducing complexity across various displays. In this effort IBM Almaden researchers consulted with systems experts to determine the layout of environments for various tasks. To support a scalable and intuitive approach, a number of navigation tools were designed to take the users from one level and perspective to another, either horizontally, vertically, or diagonally of what might be considered a grid of topology layouts.

Grouping/Aggregation

In all of the layouts, devices in the environment are organized into a number of task-dependent groups. At each level and perspective a default grouping is initially provided, but the viewer is designed to allow users to regroup devices dynamically to suit their tasks. Users can focus not only on a smaller number of devices but also on devices that are relevant to the task at hand. For example, an administrator can regroup hosts first by status, identify critical entities, and then regroup again by operating system, or by a user-defined location property to gain different perspectives on the problem. This provides a foundation for problem determination capabilities in IBM TotalStorage Productivity Center that will be expanded in future releases of this product.

The topology viewer is also designed so that storage administrators can collapse or expand groups in place (**Figure 2**). Collapsed groups show a summary of their contents, which allows users to survey the contents as well as spot important information on devices, such as degraded status, even at higher levels (see Pinning). Aggregation of information of devices supports administrators monitoring a large number of entities even at higher levels and helps guide administrators to the source of a problem at lower levels.

Connections to devices such as storage systems, tape libraries, storage switches and hosts in a collapsed group are shown as aggregated thick lines (representing more than one connection) that also provide aggregated status information (**Figure 1**). That typically occurs when a group is collapsed. It can also occur when there are multiple connections between a computer and a switch. These thick lines are groups of connections. And like groups, they aggregate health status (for the connections).

In the design of the topology viewer, straight lines are used between devices (ports) to indicate connections, thus making connections easier to trace. As connections are shown only on demand, the clutter problem common in graphical topology diagrams is avoided in most cases. This makes the topology viewer more scalable than many other storage management tools because information, such as relationships and connectivity of devices, is only available to administrators in an on demand approach—allowing the display to scale easily for more complex storage environments.

Overlays

In order to reduce visual complexity, IBM Almaden researchers used an overlay concept, which adds views with task-specific information such as health and performance status. The types of overlays available include Performance, Health and Zone Membership. Overlays are rendered as green, orange or red icons next to entities and can be independently turned on and off at any point. In essence, each overlay has a number of icons for its status. Most commonly used is the health overlay, and there you typically see red (critical), green (normal) and orange (for aggregation of normal and critical, for instance). Overlay status is then aggregated for groups up the hierarchy of devices. In effect, this turns higher-level views into a dashboard view for the entire environment. In addition, hovering on overlay icons reveals more detailed information at each level (**Figure 5**).

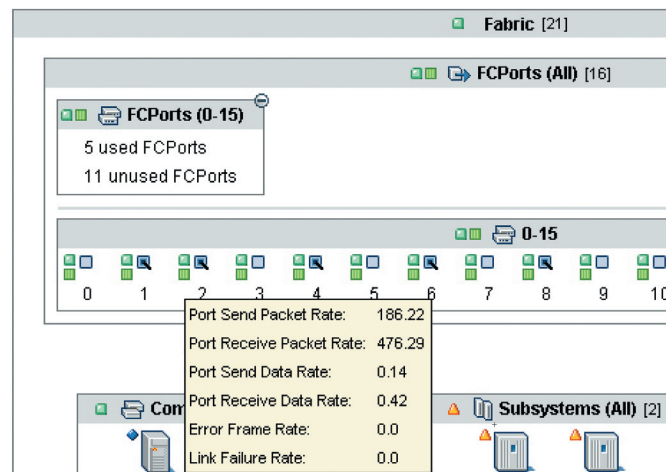


Figure 5 – Switch ports showing health and performance information. The roll-over pop-up window shows port performance information.

Pinning

Pinning is another display technique that provides a rapid view and access of user-selected devices at any level view of the topology (**Figure 6**). Typically, when administrators want to keep an eye on problematic devices over extended periods of time, they can pin such devices, thus marking these devices with a flag. A pinned device is considered to be of special interest to the user and is made visible even in collapsed groups. These devices of interest are rendered in a list along with their status and icons below the content summary of the group. Clicking on a pinned device shifts the view to the lowest level view suitable for the pinned device, as a shortcut. Or, if there are crucial storage systems or SAN switches that need to have mission-critical performance and availability, administrators can highlight these devices so they can be

monitored at a glance. Additionally, with the role-based administration capability in IBM TotalStorage Productivity Center, specific storage arrays or switches can be pinned for different kinds of roles managing storage infrastructure (Fabric or Disk operators or administrators).

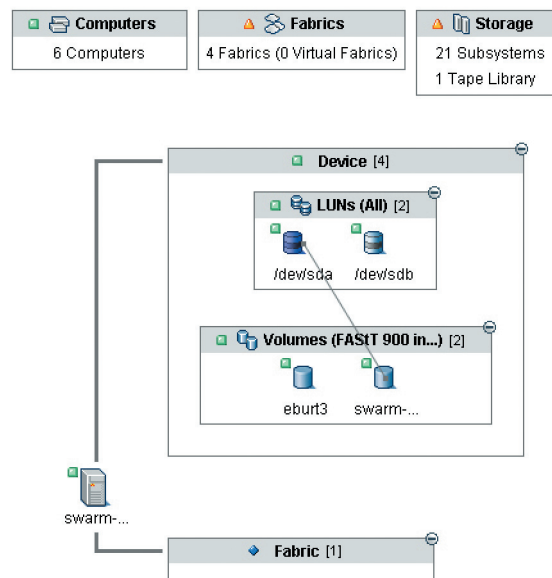


Figure 6 – Collapsed and expanded group of computers containing pinned elements.

Through pinning, users can monitor low-level devices while they are maintaining an overview of the larger environment. Unlike a “favorites” list, pinned devices are shown within their group, in context, and not in a separated list on the side.

Pinning is designed to track only a handful of devices. The purpose of pinning is to make a few devices always visible, always reachable.

Synchronized Views

One of the design goals of the IBM TotalStorage Productivity Center topology viewer was to make the topology both more accessible and usable. Another key feature of this viewer is the synchronization of data presented in both a visual view of the topology as well as in tables to give administrators choices in collecting information about their storage environments.

The topology viewer combines a graphical view with a synchronized tabular view, such that each action performed in one view is mirrored in the other view. For example, selecting a device in the graphical view also selects the same entity in the tabular view (**Figure 4**). Aside from making the graphical topology more accessible, it turns out that some actions are simply easier to do in a tabular view, such as scanning a long list of device names or generalized sorting. Used together, graphical and tabular views provide synergy that helps administrators effectively view and navigate their environment.

Summary

The IBM TotalStorage Productivity Center topology viewer is designed to deliver an advanced management console that gives storage administrators a scalable, single view of their storage infrastructures. This viewer is a departure from many of the other storage management tools today because it provides a task-focused approach to displaying information. Information is on demand for storage administrators, helping them to do the monitoring, configuration, planning, reporting and problem determination tasks of managing storage infrastructure from a centralized, real-time view of their environments. A number of key functions make this management console distinct in the storage industry by supporting enterprise-scale storage administration while still providing detailed monitoring of devices. The viewer employs high-level design principles such as semantic zooming and progressive information disclosure. It also supports new display concepts such as dynamic grouping, pinning, overlays, and synchronized views to help provide a scalable, usable, and accessible storage management platform. This viewer complements the broader capabilities of IBM TotalStorage Productivity Center that is designed to provide a simplified, centralized management tool to manage the performance, availability and capacity management of storage infrastructures.

IBM TotalStorage Productivity Center Overview

IBM TotalStorage Productivity Center includes:

- *IBM TotalStorage Productivity Center for Fabric*
- *IBM TotalStorage Productivity Center for Data*
- *IBM TotalStorage Productivity Center for Disk*
- *IBM TotalStorage Productivity Center for Replication*
- *IBM TotalStorage Productivity Center Limited Edition*
- *IBM Standard Edition (Disk, Data and Fabric), which is a package of TotalStorage Productivity Center for Disk, Data and Fabric offered at a reduced price.*

IBM TotalStorage Productivity Center is an integrated suite that includes a single user interface to manage capacity, storage networks, storage systems and replication services. Productivity Center is designed to provide both role-based administration and single-sign-on capabilities. Additionally, the suite is designed to provide the foundation for storage service-level management by offering storage system and SAN performance and availability management. This includes connectivity reporting between file systems and physical disk as well as SAN and disk subsystem failure and audit logging.

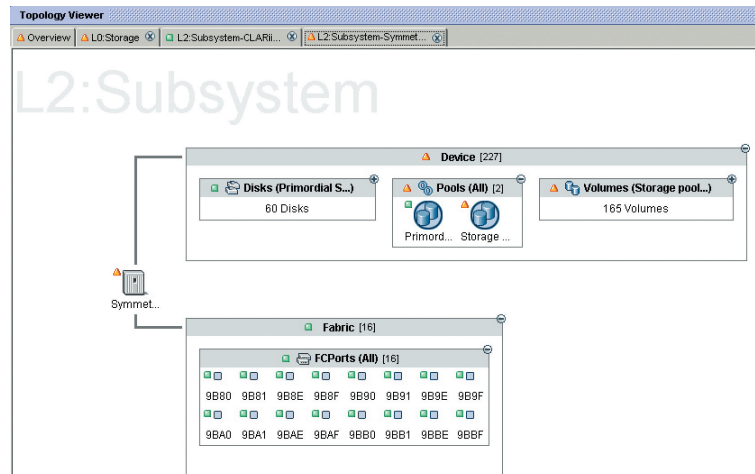


Figure 7— A deeper look at how the topology viewer is leveraged in Productivity Center for Disk—in this case looking at a storage array at a detailed level across the disks, storage pools, and volumes as well as controller ports and connectivity.

For more information:

Contact Jamie Gruener at
gruener@us.ibm.com or go to
ibm.com/storage



© Copyright IBM Corporation 2006

IBM Systems and Technology Group
5600 Cottle Road
San Jose, CA 95193
U.S.A.

September 2006
All Rights Reserved

Tivoli and TotalStorage are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

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

This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only and is subject to change without notice. IBM shall have no responsibility to update such information.

IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein. Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

MB, GB and TB equal 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, where referring to storage capacity. Actual storage capacity will vary based upon many factors and may be less than stated. Some numbers given for storage capacities give capacity in native mode followed by capacity using data compression technology.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESSED OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided.

References in this document to IBM products, programs or services do not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM program or product in this document is not intended to state or imply that only that program may be used. Any functionally equivalent program or product that does not infringe IBM's intellectual property rights may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

TSW02067-USEN-00