

IBM System Storage SAN Volume Controller 7.7.0 release – IP quorum GUI configuration guidelines

This white paper provides a brief introduction to the IBM® System Storage® SAN Volume Controller (SVC) cluster with the Internet Protocol (IP) based quorum support feature in the 7.6.0 release along with the new GUI enhancements in release 7.7.0. This white paper covers the GUI steps to configure IP quorum. The CLI version of IP quorum deployment is still valid but version 7.7.0 provides an attractive and a simpler GUI to configure the IP quorum application. This white paper includes the details of deployment with screen captures of IP-based quorum deployment using the GUI.

Introduction

A quorum disk is a managed disk (MDisk) that contains an exclusive reserved area for use by the system. The system uses quorum disks to break a tie when exactly half the nodes in the system remain after a failure and half of the nodes die at the same time. This is referred to as a split-brain situation. There are three candidate quorum disks that exist in the system. However, only one quorum disk is active at any time.

In other words, a quorum disk is a disk that contains a reserved area used exclusively for cluster membership management. The quorum disk is accessed when it is necessary to determine which half of the clustered system continues to read and write data. Quorum disks can either be MDisk or drives in the storage controller cluster. There are two main use cases of quorum disk.

- As a tiebreaker if there is a storage area network (SAN) fault (when half of the nodes in the cluster die at the same time)
- To hold a copy of important system configuration data

Table of contents

<i>Introduction</i>	1
<i>About this white paper</i>	2
<i>Assumptions</i>	2
<i>What is IP-based quorum support?</i> ..	3
<i>GUI-based deployment steps</i>	4
<i>Summary</i>	8
<i>Resources</i>	8
<i>About the authors</i>	8

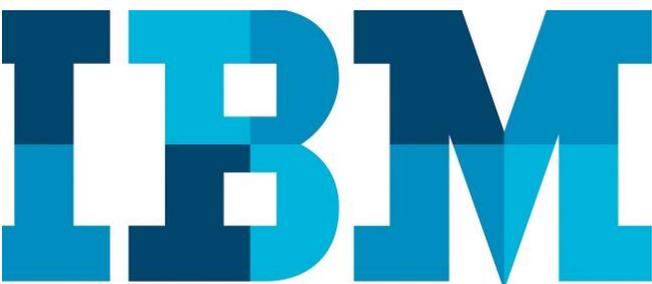
Overview

Challenge

Versions before IBM SVC 7.7.0 lacked the graphical user interface to configure IP quorum feature, and because of this, IP quorum configuration and monitoring was arduous.

Solution

IBM SVC release 7.7.0 brings the much-needed attractive user interface to configure and monitor IP quorum, thus providing an easier management model for IP quorums.



The quorum disk saves data slightly more than 256 MB, which is reserved to store important system configuration data on each quorum disk candidate. The storage controller cluster system automatically selects the active quorum disk from the pool of assigned quorum disk candidates. If a tiebreaker condition occurs, one-half portion of the system nodes, which can reserve the quorum disk after the split-brain condition occurred, lock the quorum disk and continue to operate. At this stage, the other-half portion stops its operation. This design prevents both the sites from claiming ownership of the cluster and start their own configuration nodes.

When MDisks are added to the SVC cluster, the system checks the MDisks to see whether they can be used as quorum disks. If the MDisks fulfill the requirements, then SVC assigns the MDisks that are added to the system as quorum candidates. One of these MDisks is selected as the active quorum disk. An internal mechanism of SVC places the quorum candidates on separate disk subsystems, if possible. After the quorum disk is selected, no attempt is made to ensure that the other quorum candidates are presented to the system through separate disk subsystems or not. This mechanism of quorum disk placement verification and adjustment to separate storage systems (if possible) reduce the dependency from a single storage system and can increase the quorum disk availability significantly.

About this white paper

The intent of this white paper is to be a support for self-guided configuration of IP quorum with SVC or Storwize controllers for the storage administrators and IT professionals, and to walk them through different configurations, considerations, and workflows required to set up the IP-based quorum using GUI.

This paper intends to provide an overview of the steps required to successfully evaluate and deploy IP quorum using the GUI. This white paper is not a replacement for product documentation and recommends users to refer to the IBM Redbooks® product documentation, IBM Knowledge Center, and CLI Guides at:

ibm.com/support/knowledgecenter/STPVGU_7.7.0/com.ibm.storage.svc.console.770.doc/svc_ichome_770.html

Assumptions

The following assumptions are considered while writing this white paper.

- The storage area network (SAN) is configured according to the product documentation and the infrastructure is ready to support the IP-quorum installation and view the results.
- The user has the basic understanding and awareness of quorum and IP-based quorum.

The following table summarizes the abbreviations or terminologies mentioned in this paper.

SVC and Storwize terminology	Brief description
SAN storage virtualization	Process of taking an external (possibly third-party) storage controller under the control of SVC or IBM Storwize® by exporting it to existing SVC or Storwize systems to be used for creating pools of storage for the purpose of presenting them to the hosts or applications as logical units (LUs)
iSCSI	A protocol standard by Internet Engineering Task Force (IETF) designed to interconnect SCSI devices using the TCP/IP protocol. It is described in detail in the RFC 3720 at: https://www.ietf.org/rfc/rfc3720.txt
SAN	Storage area network
NAS	Network-attached storage
FC	Fibre Channel
FCoE	Fibre Channel over Ethernet
IOPS	Input/Output operations per second
SVC/IBM SAN Volume Controller	Unless explicitly specified, a general term that is used to describe all applicable SVC platforms – IBM SVC (CG8, CF8), IBM Storwize V7000, Storwize V5000, Storwize V3700, Storwize V9000, and IBM PureFlex™ System – storage nodes

Table 1: Terminology and description

What is IP-based quorum support?

With respect to quorum disk, IP quorum is a main feature released in IBM SVC 7.6.0 release. IP-based quorum support can enable the use of a low cost IP network-attached host as a quorum disk for simplified implementation and operation. IBM SVC currently uses storage from MDisks or drives for quorum resolution. As stated earlier, normally three MDisks are either automatically or manually selected to be the quorum disk candidates. Storage system exclusively reserves some area in each of these disks to store quorum data. In the case of a quorum loss, which is a split-brain scenario (where half of the cluster cannot see the other half) the quorum disks are used to break the tie. The first half in a split-brain scenario, to reach the quorum disks, assumes ownership of the cluster and locks the disks. All the nodes in a cluster must have access to the quorum disks.

In this implementation design, no extra hardware or networking is required beyond what is normally provisioned within a cluster [that is, Fibre Chanel (FC) or serial-attached SCSI (SAS) attached storage]. But, in case of a stretched cluster environment, the need for accessibility to the quorum device in the case of a site failure necessitates the presence of a third independent domain for quorum resolution. In versions before 7.6.0, the requirement of a

third site was mandate over FC connectivity to the third site with some associated storage controller at the third site. Maintaining this third site and storage controller over FC makes the system costly in case of site recovery implementation of IBM SVC.

To overcome this limitation of maintaining a third site with FC connectivity along with a site 3 controller, the alternative is to implement Ethernet-attached quorum servers at the third site that can be run on hosts. Ethernet connectivity is generally easier and more economical to provide than FC connectivity, and hosts are typically less expensive than fully fledged network-attached storage controllers. This implementation of using host application over Ethernet connection can reduce the implementation and maintenance cost.

GUI-based deployment steps

The IP quorum application is a Java™ application that runs on a host at the third site. The IP network is used for communication between the IP quorum application and node canisters in the system. An IP quorum application can also act as the quorum device for systems that are configured with a standard topology that do not have any external storage configured. This configuration gives a system tie-break capability, automatically resuming I/O processing if half of the system's node canisters are inaccessible. The maximum number of IP quorum applications that can be deployed is five. Applications can be deployed on multiple hosts to provide redundancy. IBM SVC version 7.7.0 provides GUI panels to generate and view the IP-based quorums.

Perform the following steps to configure IP quorum:

1. Open the SVC cluster GUI using the supported web browser. After logging in, you will be directed to the home page of IBM SVC software.
2. On the home page, click **Settings** → **System** and click the **IP Quorum** tab.

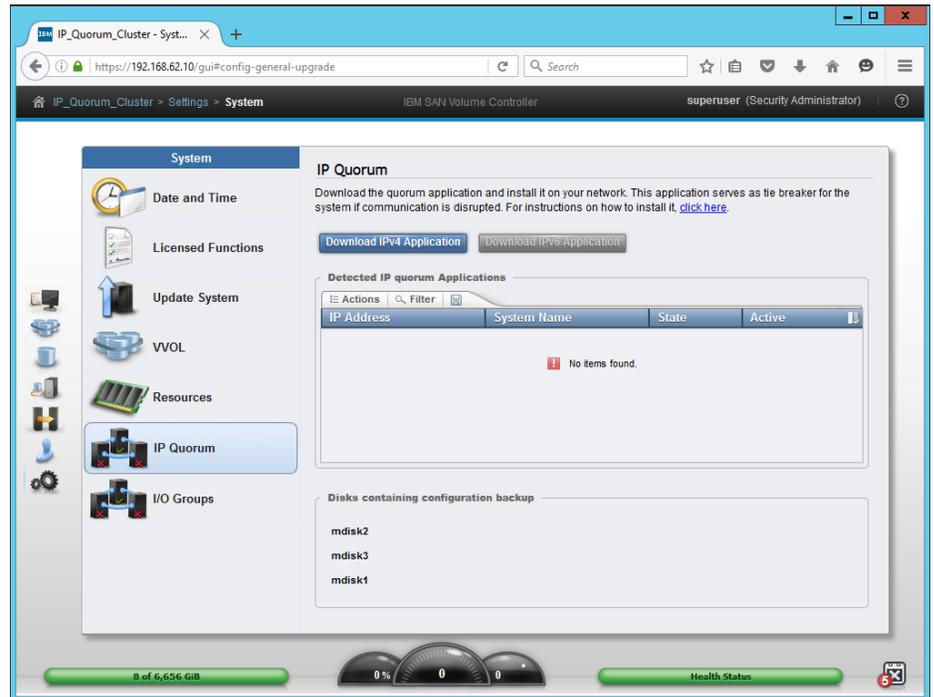


Figure 1: GUI panel for quorum disk management

3. Click **Download IPv4 Application** (or **Download IPv6 Application** if you are using IPv6 setup). This step will complete with errors if the service IP on node is not set. This step will provide a signed JAR file as output to the dumps directory on the config node and prompt user to download it locally.

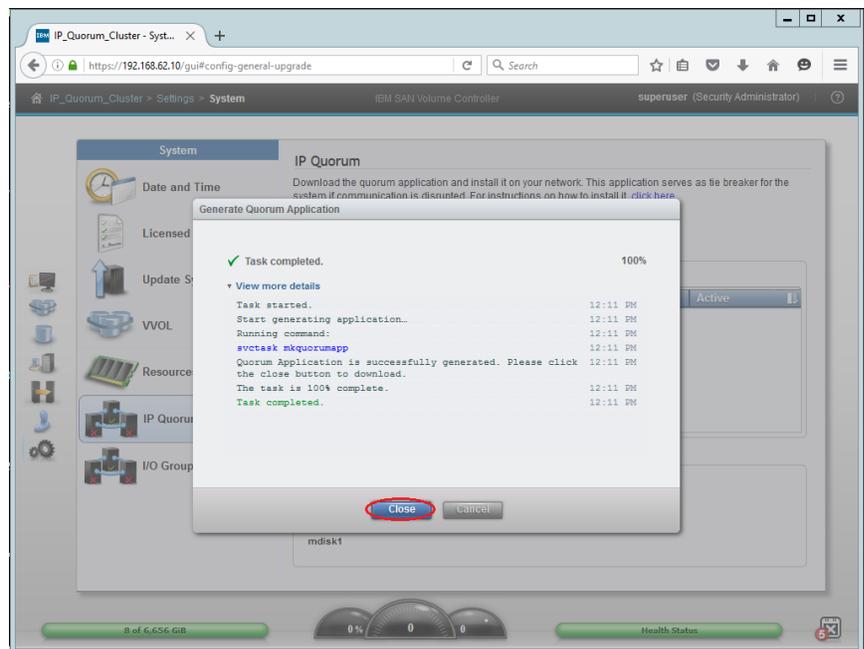


Figure 1: Generating ip_quorum.jar using GUI

4. Click **Save File** to save the ip_quorum.jar file to your computer.

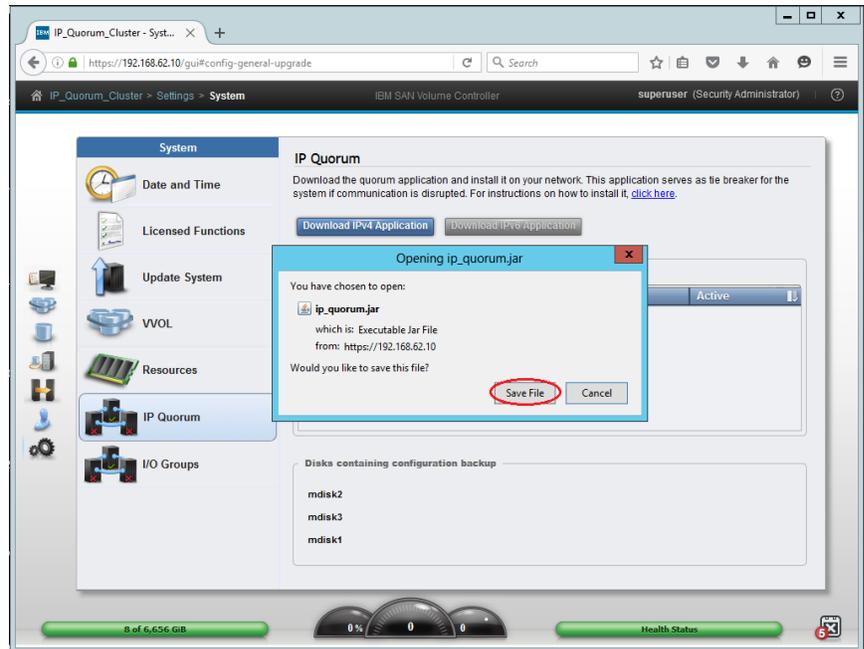


Figure 3: Saving the Java application file

5. Copy the JAR file generated in the previous step to the supported Linux® server where IBM Java 7.1 or IBM Java 8 is already being installed. The JAR file contains the application and all additional files needed for the quorum application to work with the SVC cluster. The JAR file must be regenerated when the user changes a service IP, when the cluster is assigned a new Secure Sockets Layer (SSL) certificate, or when a node is permanently added or removed from the cluster. Only one JAR file needs to be generated, and can then be deployed on multiple hosts.
6. Initialize the quorum application on the host by running the quorum ip_quorum.jar file. The app creates a log file in the current directory, which can be used to assist in debugging issues. During usage, you can check the output from the app for any potential errors.

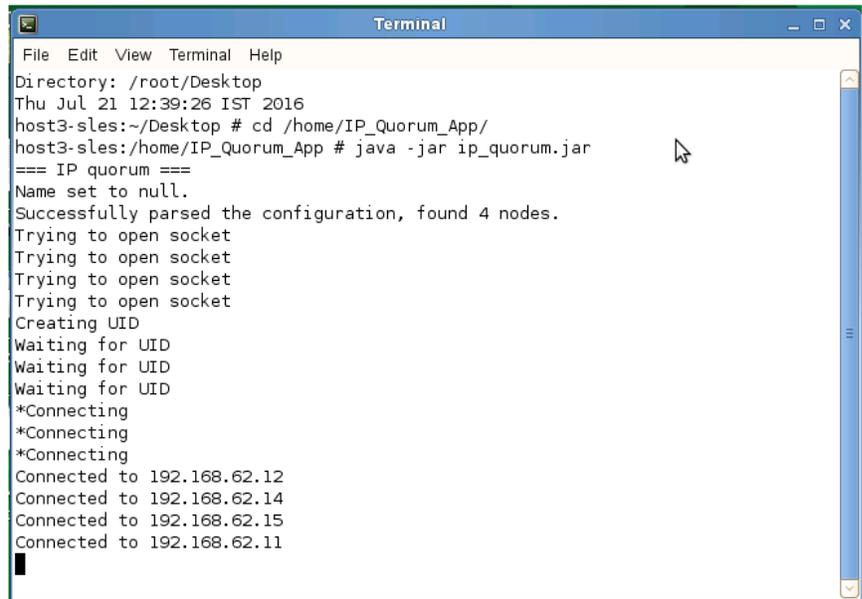


Figure 4: Initializing the IP-quorum application

7. After starting the Java application, notice that the IP quorum disk is listed in the **Detected IP quorum Applications** section of the GUI window.

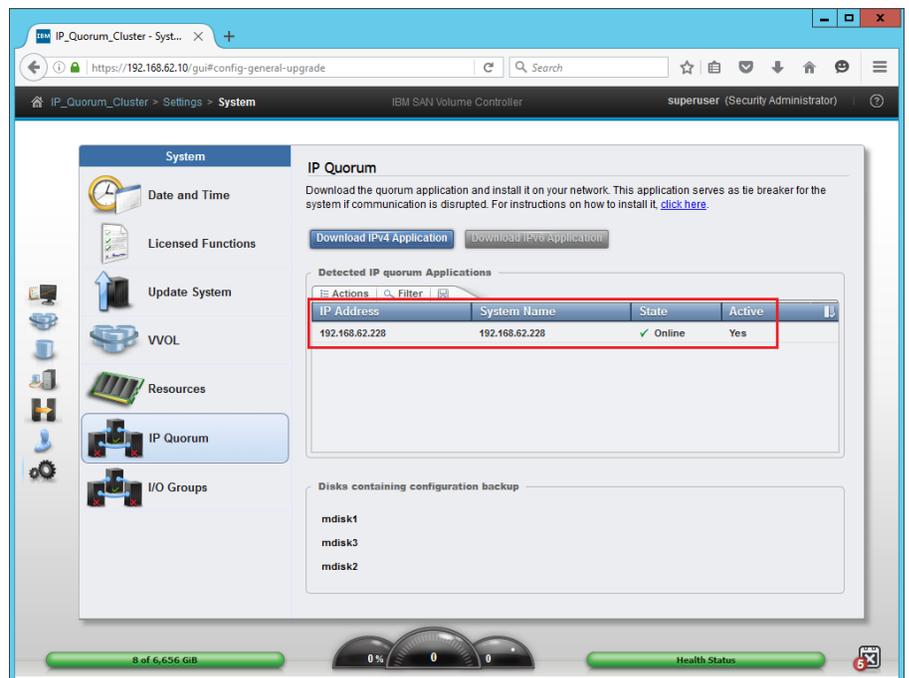


Figure 5: Additional quorum disks listed in the IP quorum section of GUI panel

The IP quorum setup is now complete and a new quorum disk is visible to the cluster and can be used by the SVC cluster for tiebreak and other purposes.

Summary

This white paper provided a brief introduction to the IBM System Storage SVC cluster with the IP-based quorum support in the 7.6.0 release. Improvements in the quorum disk area enable to use the IP-based quorum hosted over a Linux server. This feature can be effectively used in a stretched cluster or IBM HyperSwap® topology configurations, or both, where a third site is mandatory to break the tie in case of a split-brain scenario. This IP-based quorum saves the Fibre Channel infrastructure for the third site and can be used over an IP network. This feature is easy to deploy and users can migrate their current FC-based third-site quorum to an IP-based quorum by adding the IP quorum and then removing the older FC quorum at the third site. This feature supports only one JAR file to deploy from multiple Linux servers to have more than one IP quorum available. This paper provides a complete configuration guideline using the GUI, which is added in release 7.7.0 of IBM SVC. This addition graphical user interface provides a simpler and a less time consuming configuration of the IP quorum feature.

Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld®
ibm.com/partnerworld/systems
- IBM Power Development Platform
ibm.com/partnerworld/pdp
- IBM Power Systems Knowledge Center
<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp>
- IBM Redbooks
ibm.com/redbooks

About the authors

Kushal S Patel is a software engineer in the IBM Systems and Storwize team. You can reach Kushal at kushpate@in.ibm.com.

Sarvesh S Patel is a senior staff software engineer in IBM Systems SVC and Storwize team. You can reach Sarvesh at sarvepat@in.ibm.com.

Shiness George is a senior software engineer in the IBM Systems and Storwize team. You can reach Shiness at shingeor@in.ibm.com.



©Copyright IBM Corporation 2016
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of the Internal Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

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

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.

Please recycle

