

Use of rose-storage package?

distributed storage as one of its use cases



Hello!

I am Martin Pína

I'm a certified MikroTik trainer and my trademark is that I'm on a time slip most of the time, so let's get to it. :-)

01

Rose-storage package

What is rose-storage package?

Why and when to use it?

What is rose storage package?

- one of „extra“ packages
- aims for improved disk management, officially „adds additional enterprise data center functionality“
 - RAID (JBOD, RAID 0,1,4,5,6, 1+0)
 - iSCSI (both target and initiator)
 - file services over NFS (client and server, only v4)
 - NVMe over TCP
 - encryption (SED and dmccrypt)
 - (SMB - moved to main package since 7.14)
 - (RAM disk)

Current status of rose-storage

- added in 7.8
- supported architectures are ARM, ARM64, x86 and TILE
- rose-storage support page
 - missing descriptions
 - not updated since 21st March 2023
- only for additional storage? (i.e. excluding system disk)

Dashboard / RouterOS / Extended features

ROSE-storage

- Summary
- General interface properties
- Partitions
- RAID
 - RAID levels
 - RAID 0
 - RAID 1
 - RAID 4
 - RAID 5
 - RAID 6
 - Linear
 - Nested RAID

<https://help.mikrotik.com/docs/display/ROS/ROSE-storage>

Why to use rose-storage?

- RAID - e.g. mirror for Dude storage, but... ...devices with multiple drive options are rare
- NFS - e.g. good to have common „repository“, but... ...limited to supported architectures
- encryption - up to your own paranoia...
- iSCSI/NVMe over TCP - block devices over network?
- (r)sync - nice add-on feature for small setups, probably coupled with SMB/NFS

What is appropriate hardware?

- for „client“ applications
 - whatever of supported architecture
 - with enough processing power
- for „server“ applications (despite of supported architectures)
 - NVMe slot
 - SATA port
 - USB (preferably 3+) port
- hardware matrix

MIKROTIK

Products

Product filters Hide

Show products

Current

New

Archived

Option range **+**

Use filters

Reset filter

Useful links

- Product catalog (PDF)
- **Product matrix**
- What is RouterOS?
- Wireless link calculator
- Selection guide for PTP links

Result?

- good idea
- not much in focus
- poor hardware choice

02

Real world application

Vision versus reality

Just an idea...

- Premise
 - there are lots of devices out there
 - they run anyway
 - consumption of attached storage is close to negligible
- Is there any way how to benefit/take advantage of it?

Vision

Many devices with even more storage devices connected creating one big pool of high available, resilient, geo-distributed storage...



Vision versus reality

Pros

- distribution is nature
geo-“backup“
- no additional storage hardware
- reduced electricity cost ?

Cons

- worse reliability/availability
- MikroTik hardware is not primarily built for storage
- worse/storage replacement
- speed limited to network

What else to think of...

- If some kind of proxy/controller is needed, or end device should connect to storage service in RouterOS directly?
- If controller, where to locate it?
controller can handle redundancy/replicas
 - another RouterOS?
 - virtual machine?
 - docker?
- Where to do the encryption
- Serviceability

Adoption to real world

- due to speed of network and storage used - cold storage is safe bet
- number of replicas may be a challenge (ceph uses 3 by default)
- find appropriate hardware
 - RB1100AHx4 Dude edition (4x SATA)
 - some boards with 1 USB3 slot (RBM33G is not supported architecture)
 - CCR2116 (1x NVMe slot)

Disadvantages from the start

- it is complicated/not possible to do all operations with disks as we are used to from full-featured OSes
- monitoring of disk health may be a problem
- no native rose-storage bandwidth management (but yes, it is RouterOS, so there are plenty of options)
- advanced features missing (like authentication)
- mature enough? (7.14)


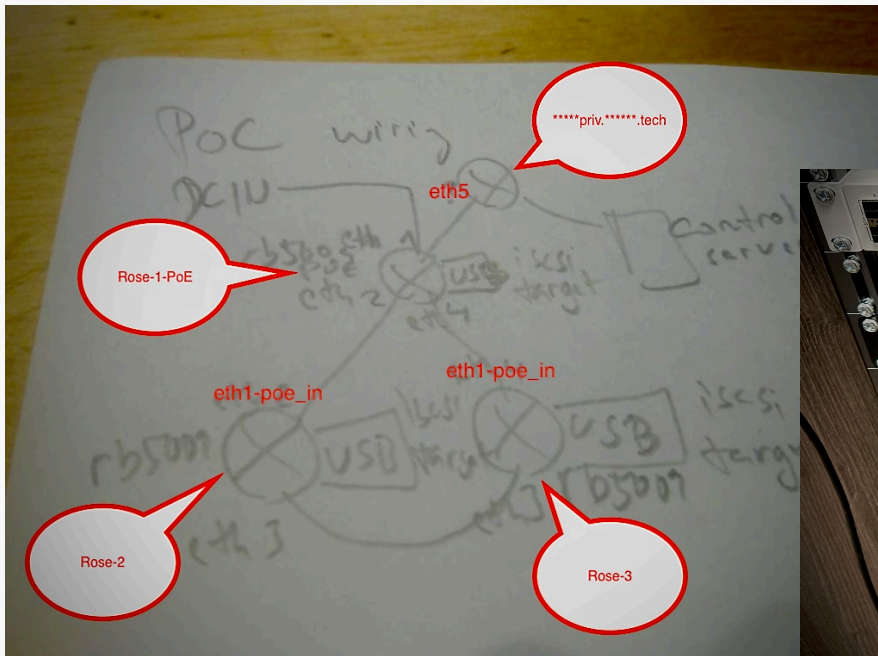
03

PoC

Three routers and three storage devices in the LAB.

Setup

- PoC wiring - from analog to digital.:-)



Outcome

- attach iSCSI succeeded
- attach NVMe over TCP (with USB) succeeded
- create RAID locally in controller succeeded
- export created array succeeded
- attach imported disks as CEPH OSD - work in progress
(succeeded for bluestore)

Possible enhancements

- tunnels used to connect to controller
- anycast addresses used to adapt to network changes (both for storage nodes and for controller)
- final redundancy option should be CEPH (but not really created for network attached storage at first place) or something similar in features (redundancy, self-healing, scrubbing)
- something else? I'm sure we could find more, but time is running out and therefore....



Thanks for your attention.



MIKROTİK
professionals conference

20240308 <https://mtpc.world/>

pinam/sshmik©2024