


MARS Manual

Multiversion Asynchronous Replicated Storage



Thomas Schöbel-Theuer (tst@1und1.de)

Version 0.1-36

Copyright (C) 2013-16 Thomas Schöbel-Theuer

Copyright (C) 2013-16 1&1 Internet AG (see <http://www.1und1.de> shortly called 1&1 in the following).


Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “[GNU Free Documentation License](#)”.

Abstract

MARS is a block-level storage replication system for long distances / flaky networks under GPL. It runs as a Linux kernel module. The sysadmin interface is similar to DRBD¹, but its internal engine is completely different from DRBD: it works with **transaction logging**, similar to some database systems.

Therefore, MARS can provide stronger **consistency guarantees**. Even in case of network bottlenecks / problems / failures, the secondaries may become outdated (reflect an elder state), but never become inconsistent. In contrast to DRBD, MARS preserves the **order of write operations** even when the network is flaky (**Anytime Consistency**).

The current version of MARS supports $k > 2$ replicas and works **asynchronously**. Therefore, application performance is completely decoupled from any network problems. Future versions are planned to also support synchronous or near-synchronous modes.



¹Registered trademarks are the property of their respective owner.

Contents

1. Why You should Replicate Big Data at Block Layer	7
1.1. Cost Arguments from Architecture	7
1.2. Cost Arguments from Technology	11
1.3. Performance Arguments from Architecture	11
2. Use Cases for MARS vs DRBD	14
2.1. Network Bottlenecks	14
2.1.1. Behaviour of DRBD	14
2.1.2. Behaviour of MARS	17
2.2. Long Distances / High Latencies	20
2.3. Higher Consistency Guarantees vs Actuality	20
3. Quick Start Guide	22
3.1. Preparation: What you Need	22
3.2. Setup Primary and Secondary Cluster Nodes	23
3.2.1. Kernel and MARS Module	23
3.2.2. Setup your Cluster Nodes	24
3.3. Creating and Maintaining Resources	24
3.4. Keeping Resources Operational	25
3.4.1. Logfile Rotation / Deletion	25
3.4.2. Switch Primary / Secondary Roles	26
3.4.2.1. Intended Switching / Planned Handover	27
3.4.2.2. Forced Switching	28
3.4.3. Split Brain Resolution	30
3.4.4. Final Destruction of a Damaged Node	31
3.4.5. Online Resizing during Operation	32
3.5. The State of MARS	32
3.6. Inspecting the State of MARS	34
4. Basic Working Principle	35
4.1. The Transaction Logger	35
4.2. The Lamport Clock	37
4.3. The Symlink Tree	38
4.4. Defending Overflow of /mars/	40
4.4.1. Countermeasures	40
4.4.1.1. Dimensioning of /mars/	40
4.4.1.2. Monitoring	41
4.4.1.3. Throttling	42
4.4.2. Emergency Mode and its Resolution	43
5. The Macro Processor	45
5.1. Predefined Macros	45
5.1.1. Predefined Complex and High-Level Macros	45
5.1.2. Predefined Primitive Macros	50
5.1.2.1. Intended for Humans	50
5.1.2.2. Intended for Scripting	52
5.2. Creating your own Macros	55
5.2.1. General Macro Syntax	55
5.2.2. Calling Builtin / Primitive Macros	56
5.2.3. Predefined Variables	60
5.3. Scripting HOWTO	61

6. The Sysadmin Interface (marsadm and /proc/sys/mars/)	63
6.1. Cluster Operations	64
6.2. Resource Operations	65
6.2.1. Resource Creation / Deletion / Modification	66
6.2.2. Operation of the Resource	68
6.2.3. Logfile Operations	73
6.2.4. Consistency Operations	74
6.3. Further Operations	74
6.3.1. Inspection Commands	74
6.3.2. Setting Parameters	75
6.3.2.1. Per-Resource Parameters	75
6.3.2.2. Global Parameters	75
6.3.3. Waiting	75
6.3.4. Low-Level Expert Commands	76
6.3.5. Senseless Commands (from DRBD)	76
6.3.6. Forbidden Commands (from DRBD)	77
6.4. The /proc/sys/mars/ and other Expert Tweaks	77
6.4.1. Syslogging	77
6.4.1.1. Logging to Files	77
6.4.1.2. Logging to Syslog	78
6.4.1.3. Tuning Verbosity of Logging	78
6.4.2. Tuning the Sync	79
7. Tips and Tricks	80
7.1. Avoiding Inappropriate Clustermanager Types for Medium and Long-Distance Replication	80
7.1.1. General Cluster Models	80
7.1.2. Handover / Failover Reasons and Scenarios	81
7.1.3. Granularity and Layering Hierarchy for Long Distances	81
7.1.4. Methods and their Appropriateness	82
7.1.4.1. Failover Methods	82
7.1.4.2. Handover Methods	88
7.1.4.3. Hybrid Methods	88
7.1.5. Special Requirements for Long Distances	88
7.2. Creating Backups via Pseudo Snapshots	89
8. MARS for Developers	90
8.1. Motivation / Politics	90
8.2. Architecture Overview	92
8.3. Some Architectural Details	92
8.3.1. MARS Architecture	92
8.3.2. MARS Full Architecture (planned)	93
8.4. Documentation of the Symlink Trees	93
8.4.1. Documentation of the MARS Symlink Tree	94
8.5. XIO Worker Bricks	94
8.6. StrategY Worker Bricks	94
8.7. The XIO Brick Personality	94
8.8. The Generic Brick Infrastructure Layer	94
8.9. The Generic Object and Aspect Infrastructure	94
A. Technical Data MARS	95
B. Handout for Midnight Problem Solving	96
B.1. Inspecting the State of MARS	96
B.2. Replication is Stuck	96
B.3. Resolution of Emergency Mode	97
B.4. Resolution of Split Brain and of Emergency Mode	98
B.5. Handover of Primary Role	99
B.6. Emergency Switching of Primary Role	99

Contents


C. Alternative Methods for Split Brain Resolution	101
D. Alternative De- and Reconstruction of a Damaged Resource	102
E. Cleanup in case of Complicated Cascading Failures	103
F. Experts only: Special Trick Switching and Rebuild	105
G. GNU Free Documentation License	107

1. Why You should Replicate Big Data at Block Layer

1.1. Cost Arguments from Architecture

Datacenters aren't usually operated for fun or for hobby. Costs are therefore a very important argument.

Many enterprise system architects are starting with a particular architecture in mind, called "big cluster". There is a common belief that otherwise **scalability** could not be achieved:



MARS Presentation by Thomas Schöbel-Theuer

The crucial point is the storage network here: n frontend servers are interconnected with $m = O(n)$ storage servers, in order to achieve properties like scalability, failure tolerance, etc.

Since *any* of the n frontends must be able to access *any* of the m storages in realtime, the storage network must be dimensioned for $O(n \cdot m) = O(n^2)$ network connections running in parallel. Even if the total network throughput would be scaling only with $O(n)$, the network has to *switch* the packets from n sources to m destinations (and their opposite way back) in **realtime**.


This **cross-bar functionality** in realtime makes the storage network expensive. Some further factors are increasing the costs of storage networks:

- In order to limit error propagation from other networks, the storage network is often built as a *physically separate / dedicated* network.
- Because storage networks are heavily reacting to high latencies and packet loss, they often need to be dimensioned for the **worst case** (load peaks, packet storms, etc), needing one of the best = most expensive components for reducing latency and increasing throughput. Dimensioning to the worst case instead of an average case plus some safety margins is nothing but an expensive **overdimensioning / over-engineering**.

1. Why You should Replicate Big Data at Block Layer

- When multipathing is required for improving fault tolerance of the storage network itself, these efforts will even **double**.
- When geo-redundancy is required, the whole mess may easily more than double another time because in cases of disasters like terrorist attacks the backup datacenter must be prepared for taking over for multiple days or weeks.

Fortunately, there is an alternative called “sharding architecture” which does not need a storage network at all, at least when built and dimensioned properly. Instead, it *should have* (but not always needs) a so-called replication network which can, when present, be dimensioned much smaller because it does neither need realtime operations, nor scalability to $O(n^2)$:



Sharding architectures are extremely well suited when both the input traffic and the data is **already partitioned**. For example, when several thousands or even millions of customers are operating on disjoint data sets, like in web hosting where each webspace is residing in its own home directory, or when each of millions of MySQL database instances has to be isolated from its neighbour.

Even in cases when any customer may potentially access any of the data items residing in the whole storage pool (e.g. like in a search engine), sharding can be often applied. The trick is to create some relatively simple content-based dynamic switching or redirect mechanism in the input network traffic, similar to HTTP load balancers or redirectors.

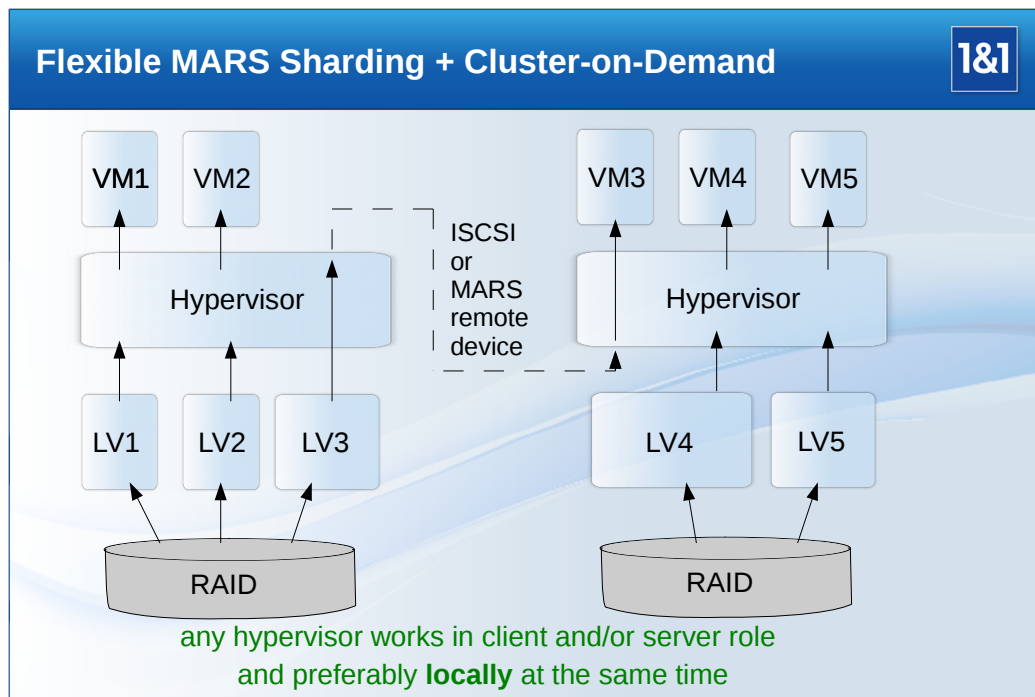
Only when partitioning of input traffic plus data is not possible in a reasonable way, big cluster architectures as implemented for example in Ceph or Swift (and partly even possible with MARS when restricted to the block layer) have their **usecase**. Only under such a precondition they are really needed.

When sharding is possible, it is the preferred model due to cost and performance reasons.



Notice that MARS' new remote device feature from the 0.2 branch series (which is a replacement for iSCSI) *could* be used for implementing the “big cluster” model at block layer.

Nevertheless, this sub-variant is not the preferred model. Following is the a super-model which combines both the “big cluster” and sharding model at block layer in a very flexible way. The following example shows only two servers from a pool consisting of hundreds or thousands of servers:



MARS Presentation by Thomas Schibel-Theuer

The idea is to use iSCSI or the MARS remote device *only where necessary*. Preferably, local storage is divided into multiple Logical Volumes (LVs) via LVM, which are *directly used locally* by Virtual Machines (VMs), such as KVM or filesystem-based variants like LXC containers.

In the above example, the left machine has relatively less CPU power or RAM than storage capacity. Therefore, not *all* LVs could be instantiated locally at the same time without causing operational problems, but *some* of them can be run locally. The example solution is to *exceptionally(!)* export LV3 to the right server, which has some otherwise unused CPU and RAM capacity.

Notice that locally running VMs doesn't produce any storage network traffic at all. Therefore, this is the preferred runtime configuration.

Only in cases of resource imbalance, such as (transient) CPU or RAM peaks (e.g. caused by DDOS attacks), *some* containers may be run somewhere else over the network. In a well-balanced and well-dimensioned system, this will be the **vast minority**, and should be only used for dealing with timely load peaks etc.

Running VMs directly on the same servers as their storage is a **major cost reducer**.

You simply don't need to buy and operate $n + m$ servers, but only about $\max(n, m) + m \cdot \epsilon$ servers, where ϵ corresponds to some relative small extra resources needed by MARS.

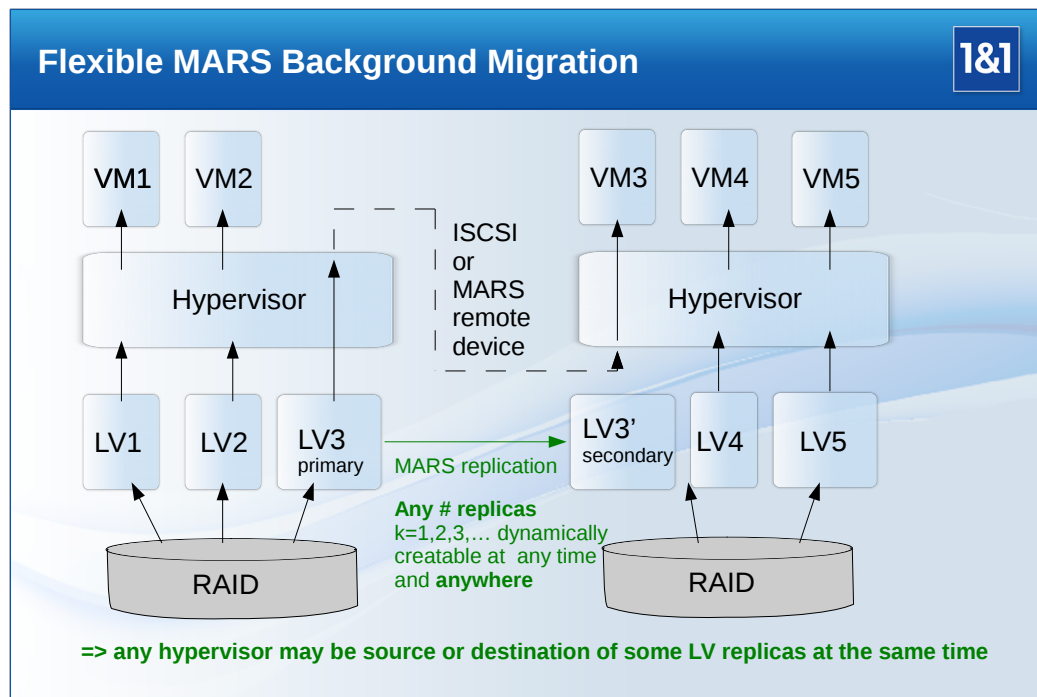


In addition to this and to reduced networking costs, there are further cost savings at power consumption, air conditioning, Height Units (HUs), number of HDDs, operating costs, etc as explained below in section 1.2.

The sharding model needs a different approach to load balancing of storage space than the big cluster model. There are several possibilities at different layers:

- Dynamically growing the sizes of LVs via `lvresize` followed by `marsadm resize` followed by `xfs_growfs` or similar operations.
- Moving customer data at filesystem or database level via `rsync` or `mysqldump` or similar.
- Moving whole LVs via MARS, as shown in the following example:

1. Why You should Replicate Big Data at Block Layer



MARS Presentation by Thomas Schöbel-Theuer

The idea is to dynamically create *additional* LV replicas for the sake of background migration. Examples:

- In case you had no redundancy at LV level before, you have $k = 1$ replicas during ordinary operation. If not yet done, you should transparently introduce MARS into your LVM-based stack by using the so-called “standalone mode” of MARS. When necessary, create the first MARS replica with `marsadm create-resource` on your already-existing LV data, which is retained unmodified, and restart your application again. Now, for the sake of migration, you just create an additional replica at another server via `marsadm join-resource` there and wait until the second mirror has been fully **synced** in background, while your application is running and while the contents of the LV is modified *in parallel* by your ordinary applications. Then you do a primary **handover** to your mirror. This is usually a matter of minutes, or even seconds. Once the application runs again at the new location, you can delete the old replica via `marsadm leave-resource` and `lvremove`. Finally, you may re-use the freed-up space for something else (e.g. `lvresize` of *another* LV followed by `marsadm resize` followed by `xfs_growfs` or similar). For the sake of some hardware lifecycle, you may run a different strategy: evacuate the original source server completely via the above MARS migration method, and eventually decommission it.
- In case you already have a redundant LV copy somewhere, you should run a similar procedure, but starting with $k = 2$ replicas, and temporarily increasing the number of replicas to either $k' = 3$ when moving each replica step-by-step, or you may even directly go up to $k' = 4$ when moving pairs at once.
- When already starting with $k > 2$ LV replicas in the starting position, you can do the same analogously, or you may then use a lesser variant. For example, we have some mission-critical servers at 1&1 which are running $k = 4$ replicas all the time on relatively small but important LVs for extremely increased safety. Only in such a case, you may have the freedom to temporarily decrease from $k = 4$ to $k' = 3$ and then going up to $k'' = 4$ again. This has the advantage of requiring less temporary storage space for *swapping* some LVs.

1.2. Cost Arguments from Technology

A common pre-judgement is that “big cluster” is the cheapest scaling storage technology when built on so-called “commodity hardware”. While this is very often true for the “commodity hardware” part, it is often not true for the “big cluster” part. But let us first look at the “commodity” part.

Here are some rough market prices for basic storage as determined around end of 2016 / start of 2017:

Technology	Enterprise-Grade	Price in € / TB
Consumer SATA disks via on-board SATA controllers	no (small-scale)	< 30 possible
SAS disks via SAS HBAs (e.g. in external 14" shelves)	halfways	< 80
SAS disks via hardware RAID + LVM (+DRBD/MARS)	yes	80 to 150
Commercial storage appliances via iSCSI	yes	around 1000
Cloud storage, S3 over 5 years lifetime	yes	3000 to 8000

You can see that any self-built and self-administered storage (whose price varies with slower high-capacity versus faster low-capacity disks) is much cheaper than any commercial offering by about a factor of 10 or even more. If you need to operate several petabytes of data, self-built storage is always cheaper than commercial one, even if additional manpower would be needed for commissioning and operating. Here we just assume that the storage is needed permanently for at least 5 years, as is the case in web hosting, databases, backup / archival systems, and many other application areas.

Cloud storage is way too much hyped. From a commercial perspective it usually pays off only when your storage demands are *extremely* varying over time, and when you need some *extra* capacity only *temporarily* for a *very* short time.

In addition to basic storage prices, many further factors come into play when roughly comparing big clusters versus sharding ($\times 2$ means with geo-redundancy):

	BC	SHA	BC $\times 2$	SHA $\times 2$
# of Disks	>200%	<120%	>400%	<240%
# of Servers	$\approx \times 2$	$\approx \times 1.1$ possible	$\approx \times 4$	$\approx \times 2.2$
Power Consumption	$\approx \times 2$	dito	dito	dito
HU Consumption	$\approx \times 2$	dito	dito	dito

The crucial point is not only the number of extra servers needed for dedicated storage boxes, but also the total number of HDDs. While big cluster implementations like Ceph or Swift can *theoretically* use some erasure encoding for avoiding full object replicas, their *practice* as seen in our internal 1&1 Ceph clusters is similar to RAID-10, but just on objects instead of block-based sectors.

Therefore a big cluster typically needs >200% disks to reach the same net capacity as a sharded cluster, where typically hardware RAID-60 with a significantly smaller overhead is sufficient for providing sufficient failure tolerance at disk level.

There is a surprising consequence from this: geo-redundancy is not as expensive as many people are believing. It just needs to be built with the proper architecture. A sharded geo-redundant pool based on hardware RAID-60 costs roughly about the same as (or when taking $O(n^2)$ storage networks into account it is possibly even cheaper than) a big cluster with full replicas without geo-redundancy. A geo-redundant sharded pool provides even better failure compensation.

Notice that geo-redundancy implies by definition that an unforeseeable **full datacenter loss** (e.g. caused by **disasters** like a terrorist attack or an earthquake) must be compensated for **several days or weeks**. Therefore it is *not* sufficient to take a big cluster and just spread it to two different locations.

In any case, a MARS-based geo-redundant sharding pool is cheaper than using commercial storage appliances which are much more expensive by their nature.

1.3. Performance Arguments from Architecture

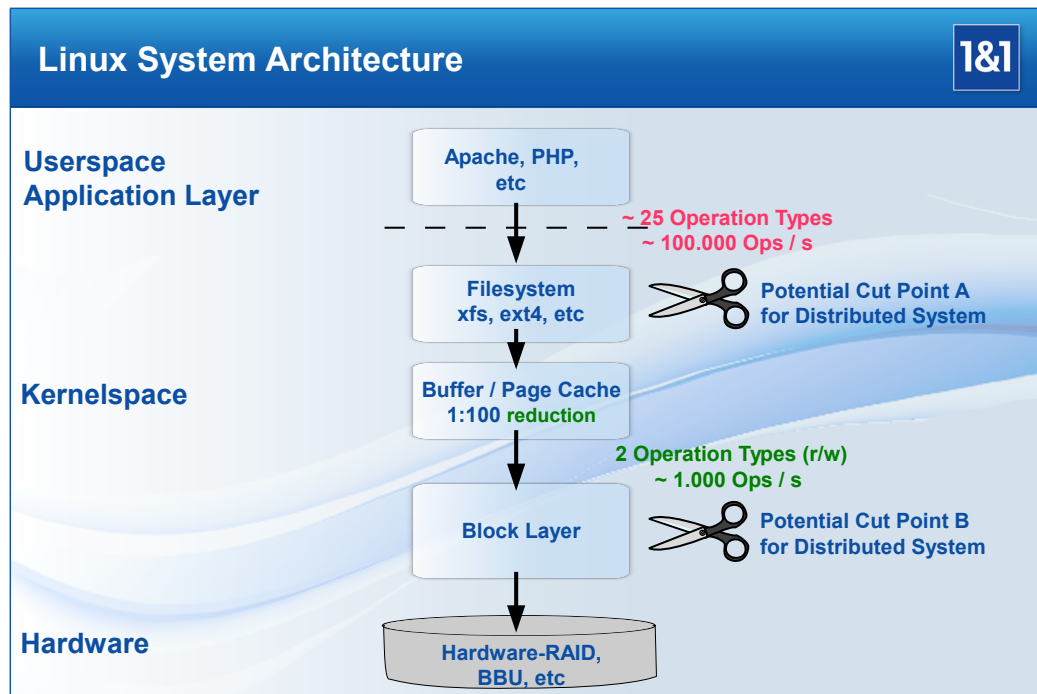
Some people think that replication is easily done at filesystem layer. There exist lots of cluster filesystems and other filesystem-layer solutions which claim to be able to replicate your data,

1. Why You should Replicate Big Data at Block Layer

sometimes even over long distances.

Trying to replicate several petabytes of data, or some billions of inodes, is however a much bigger challenge than many people can imagine.

Choosing the wrong layer for **mass data replication** may get you into trouble. Here is an explanation why replication at the block layer is more easy and less error prone:



Thomas Schöbel-Theuer

The picture shows the main components of a standalone Unix / Linux system. In the late 1970s / early 1980s, a so-called **Buffer Cache** had been introduced into the architecture of Unix. Today's Linux has refined the concept to various internal caches such as the Page Cache and the Dentry Cache.

All these caches serve only one purpose: they are reducing the load onto the storage by exploitation of fast RAM. A well-tuned cache can yield high cache hit ratios, typically 99%. In some cases (as observed in practice) even more than 99.9%.

Now start distributing the system over long distances. There are two potential cut points A and B. Cutting at A means replication at filesystem level. B means replication at block level.

When replicating at A, you will notice that the caches are *below* your cut point. Thus you will have to re-implement **distributed caches**, and you will have to **maintain cache coherence**.

When replicating at B, the Linux caches are *above* your cut point. Thus you will receive much less traffic, typically already reduced by a factor of 100, or even more. This is much more easy to cope with. You will also profit from **journalling filesystems** like `ext4` or `xfs`. In contrast, *truly distributed*¹ journalling is typically not available with distributed cluster filesystems.

A *potential* drawback of block layer replication is that you are typically limited to active-passive replication. An active-active operation is not impossible at block layer (see combinations of DRBD with `ocfs2`), but less common, and less safe to operate.

This limitation isn't necessarily caused by the choice of layer. It is simply caused by the **laws of physics**: communication is always limited by the speed of light. A distributed filesystem is nothing else but a logically **distributed shared memory (DSM)**.

Some decades of research on DSM have shown that there exist applications / workloads where the DSM model is *inferior* to the direct communication paradigm. Even in short-distance / cluster scenarios. Long-distance DSM is extremely cumbersome.

Therefore: you simply shouldn't try to solve long-distance communication needs via communication over filesystems. Even simple producer-consumer scenarios (one-way communication)

¹In this context, "truly" means that the POSIX semantics would be always guaranteed cluster-wide, and even in case of partial failures. In practice, some distributed filesystems like NFS don't even obey the POSIX standard *locally* on 1 standalone client. We know of projects which have *failed* right because of this.

are less performant (e.g. when compared to plain TCP/IP) when it comes to distributed POSIX semantics. There is simply too much **synchronisation overhead at metadata level**.

If you have a need for mixed operations at different locations in parallel: just split your data set into disjoint filesystem instances (or database / VM instances, etc). All you need is careful thought about the *appropriate granularity* of your data sets (such as well-chosen *sets* of user homedirectory subtrees, or database sets logically belonging together, etc).

Replication at filesystem level is often at single-file granularity. If you have several millions or even billions of inodes, you may easily find yourself in a snakepit.

Conclusion: active-passive operation over long distances (such as between continents) is even an advantage. It keeps you from trying bad / almost impossible things.

2. Use Cases for MARS vs DRBD

DRBD has a long history of successfully providing HA features to many users of Linux. With the advent of MARS, many people are wondering what the difference is. They ask for recommendations. In which use cases should DRBD be recommended, and in which other cases is MARS the better choice?

The following table is a short guide to the most important cases where the decision is rather clear:

Use Case	Recommendation
server pairs, each directly connected via crossover cables	DRBD
active-active / dual-primary, e.g. <code>gfs2</code> , <code>ocfs2</code>	DRBD
distance > 50km	MARS
> 100 server pairs over a short-distance shared line	MARS
all else / intermediate cases	read the following details

There exist some use cases where DRBD is clearly better than MARS. 1&1 has a long history of experiences with DRBD where it works very fine, in particular coupling Linux devices rack-to-rack via crossover cables. DRBD is just *constructed* for that use case (RAID-1 over network). In such a scenario, DRBD is better than MARS because it uses up less disk space resources. In addition, newer DRBD versions can run over high-speed but short-distance interconnects like Infiniband (via the SDP protocol). Another use case for DRBD is active-active / dual-primary mode, e.g. `ocfs2`¹ over short² distances.

On the other hand, there exist other use cases where DRBD did not work as expected, leading to incidents and other operational problems. We analyzed them for our specific use cases. The later author of MARS came to the conclusion that they could only be resolved by fundamental changes in the overall architecture of DRBD. The development of MARS started at the personal initiative of the author, first in form of a personal project during holidays, but later picked up by 1&1 as an official project.

MARS and DRBD simply have **different application areas**.

In the following, we will discuss the pros and cons of each system in particular situations and contexts, and we shed some light at their conceptual and operational differences.

2.1. Network Bottlenecks

2.1.1. Behaviour of DRBD

In order to describe the most important problem we found when DRBD was used to couple whole datacenters (each encompassing thousands of servers) over metro distances, we strip down that complicated real-life scenario to a simplified laboratory scenario in order to demonstrate the effect with minimal means.


¹Notice that `ocfs2` is apparently not constructed for long distances. 1&1 has some experiences on a specific short distance cluster where the `ocfs2` / DRBD combination scaled a little bit better than NFS, but worse than `glusterfs` (using 2 clients in both cases – notice that `glusterfs` showed extremely bad performance when trying to enable active-active `glusterfs` replication between 2 server instances, therefore we ended up using active-passive DRBD replication below a single `glusterfs` server). Conclusion: NFS < `ocfs2` < `glusterfs` < sharding. We found that `glusterfs` on top of active-passive DRBD scalability was about 2 times better than NFS on top of active-passive DRBD, while `ocfs2` on top of DRBD in active-active mode was somewhere inbetween. All cluster comparisons with an increasing workload over time (measured as number of customers which could be safely operated). Each system was replaced by the next one when the respective scalability was at its respective end, each time leading to operational problems. The ultimate solution was to replace all of these clustering concepts by the general concept of **sharding**.

²Active-active won't work over long distances at all because of high network latencies (cf chapter 1). Probably, for replication of whole clusters over long distances DRBD and MARS could be stacked: using DRBD on top for MARS for active-active clustering of `gfs2` or `ocfs2`, and a MARS instance *below* for failover of *one* of the DRBD replicas over long distances.



Notice that the following DRBD effect does not appear at crossover cables. The following scenario covers a non-standard case of DRBD. DRBD works fine when no network bottleneck appears!

The following picture illustrates an effect which has been observed in 1&1 datacenters when running masses of DRBD instances through a single network bottleneck. In addition, the effect is also reproducible by an elder version of the MARS test suite³:



The simplified scenario is the following:

1. DRBD is loaded with a low to medium, but constant rate of write operations for the sake of simplicity of the scenario.
2. The network has some throughput bottleneck, depicted as a red line. For the sake of simplicity, we just linearly decrease it over time, starting from full throughput, down to zero. The decrease is very slowly over time (some minutes, or even hours).

What will happen in this scenario?

As long as the actual DRBD write throughput is lower than the network bandwidth (left part of the horizontal blue line), DRBD works as expected.

Once the maximum network throughput (red line) starts to fall short of the required application throughput (first blue dotted line), we get into trouble. By its very nature, DRBD works **synchronously**. Therefore, it *must* transfer all your application writes through the bottleneck, but now it is impossible⁴ due to the bottleneck. As a consequence, the application running on top of DRBD will see increasingly higher IO latencies and/or stalls / hangs. We found practical cases (at least with former versions of DRBD) where IO latencies exceeded practical monitoring limits such as 5 s by far, up to the range of *minutes*. As an experienced sysadmin, you know what happens next: your application will run into an incident, and your customers will be dissatisfied.

In order to deal with such situations, DRBD has lots of tuning parameters. In particular, the `timeout` parameter and/or the `ping-timeout` parameter will determine when DRBD will give up in such a situation and simply drop the network connection as an emergency measure. Dropping the network connection is roughly equivalent to an automatic disconnect, followed by an automatic re-connect attempt after `connect-int` seconds. During the dropped connection, the incident will appear as being resolved, but at some hidden cost⁵.

³The effect has been demonstrated some years ago with DRBD version 8.3.13. By construction, it is independent from any of the DRBD series 8.3.x, 8.4.x, or 9.0.x.

⁴This is independent from the DRBD protocols A through C, because it just depends on an information-theoretic argument independently from any protocol. We have a fundamental conflict between network capabilities and application demands here, which cannot be circumvented due to the **synchronous** nature of DRBD.

⁵By appropriately tuning various DRBD parameters, such as `timeout` and/or `ping-timeout`, you can keep the impact of the incident below some viable limit. However, the automatic disconnect will then happen

2. Use Cases for MARS vs DRBD

What happens next in our scenario? During the `disconnect`, DRBD will record all positions of writes in its bitmap and/or in its activity log. As soon as the automatic re-connect succeeds after `connect-int` seconds, DRBD has to do a partial re-sync of those blocks which were marked dirty in the meantime. This leads to an *additional* bandwidth demand⁶ as indicated by the upper dotted blue box.

Of course, there is *absolutely no chance* to get the increased amount of data through our bottleneck, since not even the ordinary application load (lower dotted lines) could be transferred.

Therefore, you run at a **very high risk** that the re-sync cannot finish before the next `timeout` / `ping-timeout` cycle will drop the network connection again.

What will be the final result when that risk becomes true? Simply, your secondary site will be *permanently* in state `inconsistent`. This means, you have lost your redundancy. In our scenario, there is no chance at all to become consistent again, because the network bottleneck declines more and more, slowly. It is simply *hopeless*, by construction.

In case you lose your primary site now, you are lost at all.

Some people may argue that the probability for a similar scenario were low. We don't agree on such an argumentation. Not only because it really happens in practice, and it may even last some days until problems are fixed. In case of **rolling disasters**, the network is very likely to become flaky and/or overloaded shortly before the final damage. Even in other cases, you can easily end up with inconsistent secondaries. It occurs not only in the lab, but also in practice if you operate some hundreds or even thousands of DRBD instances.

The point is that you can produce an ill behaviour *systematically* just by overloading the network a bit for some sufficient duration.



When coupling whole datacenters via some thousands of DRBD connections, any (short) network loss will almost certainly increase the re-sync network load each time the outage appears to be over. As a consequence, overload may be *provoked* by the re-sync repair attempts. This may easily lead to self-amplifying **throughput storms** in some resonance frequency (similar to self-destruction of a bridge when an army is marching over it in lockstep).

The only way for reliable prevention of loss of secondaries is to start any re-connect *only* in such situations where you can *predict in advance* that the re-sync is *guaranteed* to finish before any network bottleneck / loss will cause an automatic disconnect again. We don't know of any method which can reliably predict the future behaviour of a complex network.



Conclusion: in the presence of network bottlenecks, you run a considerable risk that your DRBD mirrors get destroyed just in that moment when you desperately need them.





Notice that crossover cables usually never show a behaviour like depicted by the red line. Crossover cables are *passive components* which normally⁷ either work, or not. The binary connect / disconnect behaviour of DRBD has no problems to cope with that.


earlier and more often in practice. Flaky or overloaded networks may easily lead to an enormous number of automatic disconnects.

⁶DRBD parameters `sync-rate` resp `resync-rate` may be used to tune the height of the additional demand. In addition, the newer parameters `c-plan-ahead`, `c-fill-target`, `c-delay-target`, `c-min-rate`, `c-max-rate` and friends may be used to dynamically adapt to *some* situations where the application throughput *could* fit through the bottleneck. These newer parameters were developed in a cooperation between 1&1 and Linbit, the maker of DRBD.

Please note that lowering / dynamically adapting the resync rates may help in lowering the *probability* of occurrences of the above problems in practical scenarios where the bottleneck would recover to viable limits after some time. However, lowering the rates will also increase the *duration* of re-sync operations accordingly. The *total amount of re-sync data* simply does not decrease when lowering `resync-rate`; it even tends to increase over time when new requests arrive. Therefore, the *expectancy value* of problems caused by *strong* network bottlenecks (i.e. when not even the ordinary application rate is fitting through) is *not* improved by lowering or adapting `resync-rate`, but rather the expectancy value mostly depends on the *relation* between the amount of holdback data versus the amount of application write data, both measured for the duration of some given strong bottleneck.

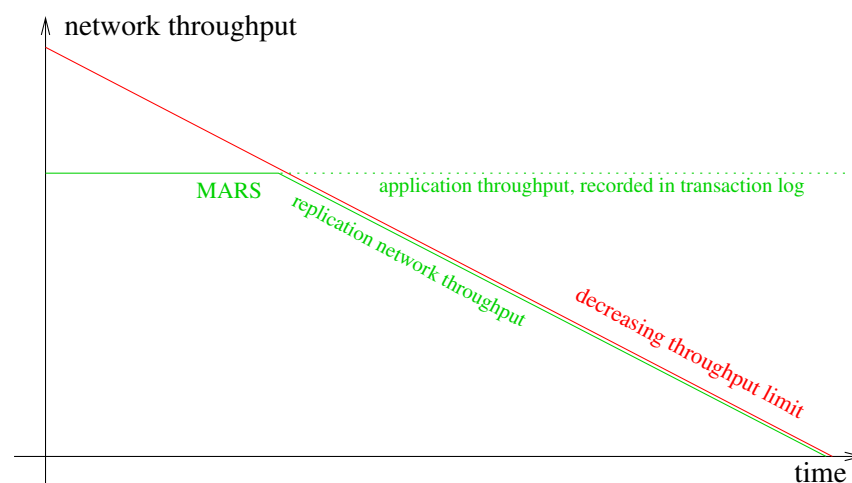
⁷Exceptions might be mechanical jiggling of plugs, or electro-magnetical interferences. We never noticed any of them.

 or  Linbit recommends a **workaround** for the inconsistencies during re-sync: LVM snapshots. We tried it, but found a *performance penalty* which made it prohibitive for our concrete application. A problem seems to be the cost of destroying snapshots. LVM uses by default a BOW strategy (Backup On Write, which is the counterpart of COW = Copy On Write). BOW increases IO latencies during ordinary operation. Retaining snapshots is cheap, but reverting them may be very costly, depending on workload. We didn't fully investigate that effect, and our experience is a few years old. You might come to a different conclusion for a different workload, for newer versions of system software, or for a different strategy if you carefully investigate the field.


 DRBD problems usually arise *only* when the network throughput shows some “awkward” analog behaviour, such as overload, or as occasionally produced by various switches / routers / transmitters, or other potential sources of packet loss.

2.1.2. Behaviour of MARS

The behaviour of MARS in the above scenario:

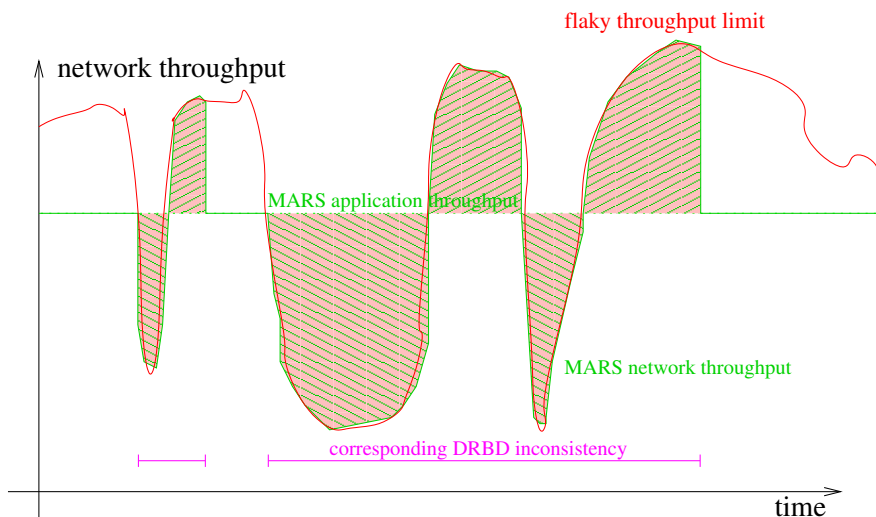


When the network is restrained, an asynchronous system like MARS will continue to serve the user IO requests (dotted green line) without any impact / incident while the actual network throughput (solid green line) follows the red line. In the meantime, all changes to the block device are recorded at the transaction logfiles.

 Here is one point in favour of DRBD: MARS stores its transaction logs on the filesystem `/mars/`. When the network bottleneck is lasting very long (some days or even some weeks), the filesystem will eventually run out of space some day. Section 4.4 discusses countermeasures against that in detail. In contrast to MARS, DRBD allocates its bitmap *statically* at resource creation time. It uses up less space, and you don't have to monitor it for (potential) overflows. The space for transaction logs is the price you have to pay if you want or need anytime consistency, or asynchronous replication in general.

In order to really grasp the *heart* of the difference between synchronous and asynchronous replication, we look at the following modified scenario:

2. Use Cases for MARS vs DRBD



This time, the network throughput (red line) is varying⁸ in some unpredictable way. As before, the application throughput served by MARS is assumed to be constant (dotted green line, often superseded by the solid green line). The actual replication network throughput is depicted by the solid green line.

As you can see, a network dropdown undershooting the application demand has no impact on the application throughput, but only on the replication network throughput. Whenever the network throughput is held back due to the flaky network, it simply catches up as soon as possible by overshooting the application throughput. The amount of lag-behind is visualized as shaded area: downward shading (below the application throughput) means an increase of the lag-behind, while the upwards shaded areas (beyond the application throughput) indicate a decrease of the lag-behind (catch-up). Once the lag-behind has been fully caught up, the network throughput suddenly jumps back to the application throughput (here visible in two cases).



Note that the existence of lag-behind areas is roughly corresponding to DRBD disconnect states, and in turn to DRBD inconsistent states of the secondary as long as the lag-behind has not been fully caught up. The very rough⁹ duration of the corresponding DRBD inconsistency phase is visualized as magenta line at the time scale.



MARS utilizes the existing network bandwidth as best as possible in order to pipe through as much data as possible, provided that there exists some data requiring expedition. Conceptually, there exists no better way due to information theoretic limits (besides data compression).



Note that *in average* during a longer period of time, the network must have enough capacity for transporting all of your data. MARS cannot magically break through information-theoretic limits. It cannot magically transport gigabytes of data over modem lines. Only *relatively short* network problems / packet loss can be compensated.



In case of lag-behind, the version of the data replicated to the secondary site corresponds

⁸In real life, many long-distance lines or even some heavily used metro lines usually show fluctuations of their network bandwidth by an order of magnitude, or even higher. We have measured them. The overall behaviour can be characterized as “**chaotic**”.

⁹Of course, this visualization is not exact. On one hand, the DRBD inconsistency phase may start later as depicted here, because it only starts *after* the first automatic disconnect, upon the first automatic re-connect. In addition, the amount of resync data may be smaller than the amount of corresponding MARS transaction logfile data, because the DRBD bitmap will coalesce multiple writes to the same block into one single transfer. On the other hand, DRBD will transfer no data at all during its disconnected state, while MARS continues its best. This leads to a prolongation of the DRBD inconsistent phase. Depending on properties of the workload and of the network, the real duration of the inconsistency phase may be both shorter or longer.

to some time in the past. Since the data is always transferred in the same order as originally submitted at the primary site, the secondary never gets inconsistent. Your mirror always remains usable. Your only potential problem could be the outdated state, corresponding to some state in the past. However, the “as-best-as-possible” approach to the network transfer ensures that your version is always *as up-to-date as possible* even under ill-behaving network bottlenecks. **There is simply no better way to do it.** In presence of temporary network bottlenecks such as network congestion, there exists no better method than prescribed by the information theoretic limit (red line, neglecting data compression).



In order to get all of your data through the line, somewhen the network must be healthy again. Otherwise, data will be recorded until the capacity of the `/mars/` filesystem is exhausted, leading to an emergency mode (see section B.3).



MARS’ property of never sacrificing local data consistency (at the possible cost of actuality, as long as you have enough capacity in `/mars/`) is called **Anytime Consistency**.




Even when the capacity of `/mars/` is exhausted and when emergency mode is entered, the replicas will not become inconsistent by themselves. However, when the emergency mode is later *cleaned up* for a replica, it will become temporarily inconsistent during the fast full sync. Details are in section B.3.



Conclusion: you can even use **traffic shaping** on MARS’ TCP connections in order to globally balance your network throughput (of course at the cost of actuality, but without sacrificing local data consistency). If you would try to do the same with DRBD, you could easily provoke a disaster. MARS simply tolerates any network problems, provided that there is enough disk space for transaction logfiles. Even in case of completely filling up your disk with transaction logfiles after some days or weeks, you will not lose local consistency anywhere (see section 4.4).

Finally, here is yet another scenario where MARS can cope with the situation:



This time, the network throughput limit (solid red line) is assumed to be constant. However, the application workload (dotted green line) shows some heavy peaks. We know from our 1&1 datacenters that such an application behaviour is very common (e.g. in case of certain kinds of DDOS attacks etc).

When the peaks are exceeding the network capabilities for some time, the replication network throughput (solid green line) will be limited for a short time, stay a little bit longer at the limit, and finally drop down again to the normal workload. In other words, you get a flexible buffering behaviour, coping with the peaks.

2. Use Cases for MARS vs DRBD

Similar scenarios (where both the application workload has peaks and the network is flaky to some degree) are rather common. If you would use DRBD there, you were likely to run into regular application performance problems and/or frequent automatic disconnect cycles, depending on the height and on the duration of the peaks, and on network resources.

2.2. Long Distances / High Latencies

In general and in some theories, latencies are conceptually independent from throughput, at least to some degree. There exist all 4 possible combinations:

1. There exist communication lines with high latencies but also high throughput. Examples are raw fibre cables at the ground of the Atlantic.
2. High latencies on low-throughput lines is very easy to achieve. If you never saw it, you never ran interactive `vi` over `ssh` in parallel to downloads on your old-fashioned modem line.
3. Low latencies need not be incompatible with high throughput. See Myrinet, InfiniBand or high-speed point-to-point interconnects, such as modern RAM busses.
4. Low latency combined with low throughput is also possible: in an ATM system (or another pre-reservation system for bandwidth), just increase the multiplex factor on low-capacity but short lines, which is only possible at the cost of assigned bandwidth.

In the *internet* practice, however, it is very likely that high latencies will also lead to worse throughput, because of the *congestion control algorithms* running all over the world.

We have experimented with extremely large TCP send/receive buffers plus various window sizes and congestion control algorithms over long-distance lines between the USA and Europe. Yes, it is possible to improve the behaviour to some degree. But magic does not happen. Natural laws will always hold. You simply cannot travel faster than the speed of light.

Our experience leads to the following rule of thumb, not formally proven by anything, but just observed in practice:

In general¹⁰, synchronous data replication (not limited to applications of DRBD) works reliably only over distances < 50 km, or sometimes even less.

There may be some exceptions, e.g. when dealing with low-end workstation loads. But when you are responsible for a whole datacenter and/or some centralized storage units, don't waste your time by trying (almost) impossible things. We recommend to use MARS in such use cases.

2.3. Higher Consistency Guarantees vs Actuality

We already saw in section 2.1 that certain types of network bottlenecks can easily (and reproducibly) destroy the consistency of your DRBD secondary, while MARS will preserve local consistency at the cost of actuality (**anytime consistency**).

Some people, often located at database operations, are obtrusively arguing that actuality is such a high good that it must not be sacrificed under any circumstances.

Anyone arguing this way has at least the following choices (list may be incomplete):

1. None of the above use cases for MARS apply. For instance, short distance replication over crossover cables is sufficient (which occurs very often), or the network is reliable enough such that bottlenecks can never occur (e.g. because the total load is extremely low, or conversely the network is extremely overengineered / expensive), or the occurrence of bottlenecks can *provably* be taken into account. In such cases, DRBD is clearly the better solution than MARS, because it provides better actuality than the current version of MARS, and it uses up less disk resources.

¹⁰We have heard of cases where even less than 50 km were not working with DRBD. It depends on application workload, on properties of the line, and on congestion caused by other traffic. Some other people told us that according to *their* experience, much lesser distances should be considered operable, only in the range of a few single kilometers. However, they agree that DRBD is rock stable when used on crossover cables.

2. In the presence of network bottlenecks, people didn't notice and/or didn't understand and/or did under-estimate the risk of accidental invalidation of their DRBD secondaries. They should carefully check that risk. They should convince themselves that the risk is *really* bearable. Once they are hit by a systematic chain of events which *reproducibly* provoke the bad effect, it is too late¹¹.
3. In the presence of network bottlenecks, people found a solution such that DRBD does not automatically re-connect after the connection has been dropped due to network problems (c.f. `ko-count` parameter). So the risk of inconsistency *appears* to have vanished. In some cases, people did not notice that the risk has *not completely*¹² vanished, and/or they did not notice that now the actuality produced by DRBD is even drastically worse than that of MARS (in the same situation). It is true that DRBD provides better actuality in `connected` state, but for a full picture the actuality in `disconnected` state should not be neglected¹³. So they didn't notice that their argumentation on the importance of actuality may be fundamentally wrong. A possible way to overcome that may be re-reading section 2.1.2 and comparing its outcome with the corresponding outcome of DRBD in the same situation.
4. People are stuck in contradictive requirements because the current version of MARS does not yet support synchronous or pseudo-synchronous operation modes. This should be resolved some day.



A common misunderstanding is about the actuality guarantees provided by filesystems. The buffer cache / page cache uses by default a **writeback strategy** for performance reasons. Even modern journaling filesystems will (by default) provide only consistency guarantees, but no strong actuality guarantee. In case of power loss, some transactions may be even *rolled back* in order to restore consistency. According to POSIX¹⁴ and other standards, the only *reliable* way to achieve actuality is usage of system calls like `sync()`, `fsync()`, `fdatasync()`, flags like `O_DIRECT`, or similar. For performance reasons, the *vast majority of applications* don't use them at all, or use them only sparingly!



It makes no sense to require strong actuality guarantees from any block layer replication (whether DRBD or future versions of MARS) while higher layers such as filesystems or even applications are already sacrificing them!



In summary, the **anytime consistency** provided by MARS is an argument you should consider, even if you need an extra hard disk for transaction logfiles.

¹¹Some people seem to need a bad experience before they get the difference between risk caused by reproducible effects and inverted luck.

¹²Hint: what's the *conceptual* difference between an automatic and a manual re-connect? Yes, you can try to *lower* the risk in some cases by transferring risks to human analysis and human decisions, but did you take into account the possibility of human errors?

¹³Hint: a potential hurdle may be the fact that the current format of `/proc/drbd` does neither display the timestamp of the first *relevant* network drop nor the total amount of lag-behind user data (which is *not* the same as the number of dirty bits in the bitmap), while `marsadm view` can display it. So it is difficult to judge the risks. Possibly a chance is inspection of DRBD messages in the syslog, but quantification could remain hard.

¹⁴The above argumentation also applies to Windows filesystems in analogous way.

3. Quick Start Guide

This chapter is for impatient but experienced sysadmins who already know DRBD. For more complete information, refer to chapter [The Sysadmin Interface \(marsadm and /proc/sys/mars/\)](#).

3.1. Preparation: What you Need

Typically, you will use MARS at servers in a datacenter for replication of big masses of data.

Typically, you will use MARS for replication *between* multiple datacenters, when the distances are greater than ≈ 50 km. Many other solutions, even from commercial storage vendors, will not work reliably over large distances when your network is not *extremely* reliable, or when you try to push huge masses of data from high-performance applications through a network bottleneck. If you ever encountered suchlike problems (or try to avoid them in advance), MARS is for you.

You can use MARS both at dedicated storage servers (e.g. for serving Windows clients), or at standalone Linux servers where CPU and storage are not separated.

In order to protect your data from low-level disk failures, you should use a hardware RAID controller with BBU. Software RAID is explicitly *not* recommended, because it generally provides worse performance due to the lack of a hardware BBU (for some benchmark comparisons with/out BBU, see <https://github.com/schoebel/blkreplay/raw/master/doc/blkreplay.pdf>).

Typically, you will need more than one RAID set¹ for big masses of data. Therefore, use of LVM is also recommended² for your data.

MARS' tolerance of networking problems comes with some cost. You will need some extra space for the transaction logfiles of MARS, residing at the `/mars/` filesystem.

The exact space requirements for `/mars/` depend on the *average write rate* of your application, not on the size of your data. We found that only few applications are writing more than 1 TB per day. Most are writing even less than 100 GB per day. Usually, you want to dimension `/mars/` such that you can survive a network loss lasting 3 days / about one weekend. This can be achieved with current technology rather easily: as a simple rule of thumb, just use one **dedicated disk** having a capacity of 4 TB or more. Typically, that will provide you with plenty of headroom even for bigger networking incidents.

Dedicated disks for `/mars/` have another advantage: their mechanical head movement is completely independent from your data head movements. For best performance, attach that dedicated disk to your hardware RAID controller with BBU, building a separate RAID set (even if it consists only of a single disk – notice that the **hardware BBU** is the crucial point).

If you are concerned about reliability, use two disks switched together as a relatively small RAID-1 set. For extremely high performance demands, you may consider (and check) RAID-10.

Since the transaction logfiles are highly sequential in their access pattern, a cheap but high-capacity SATA disk (or nearline-SAS disk) is usually sufficient. At the time of this writing, standard SATA SSDs have shown to be *not* (yet) preferable. Although they offer high random IOPS rate, their sequential throughput is worse, and their long-term stability is questioned by many people at the time of this writing. However, as technology evolves and becomes more mature, this could change in future.

¹For low-cost storage, RAID-5 is no longer regarded safe for today's typical storage sizes, because the error rate is regarded too high. Therefore, use RAID-6. If you need more than 15 disks in total, create multiple RAID sets (each having at most 15 disks, better about 12 disks) and stripe them via LVM (or via your hardware RAID controller if it supports RAID-60).

²You may also combine MARS with commercial storage boxes connected via Fibrechannel or iSCSI, but we have not yet operational experiences at 1&1 with such setups.

Use `ext4` for `/mars/`. Avoid `ext3`, and don't use `xfs`³ at all.



Notice that the filesystem `/mars/` has nothing to do with an ordinary filesystem. It is completely reserved for MARS internal purposes, namely as a **storage container** for MARS' persistent data. It does not obey any userspace rules like FHS (filesystem hierarchy standard), and it should not be accessed by any userspace tool except the official `marsadm` tool. Its internal data format should be regarded as a **blackbox** by you. The internal data format may change in future, or the complete `/mars/` filesystem may be even replaced by a totally different container format, while the official `marsadm` interface is supposed to remain stable.



That said, you might look into its contents *by hand* for curiosity or for *debugging purposes*, and only as root. But don't program any tools / monitoring scripts / etc bypassing the official `marsadm` tool.



Like DRBD, the current version of MARS has **no security** built in. MARS assumes that it is running in a **trusted network**. Anyone who can connect to the MARS ports (default 7777 to 7779) can potentially breach in and become root! Therefore, you **must** protect your network by appropriate means, such as firewalling and/or encrypted VPN.

Currently, MARS provides no shared secret like DRBD, because a simple shared secret is way too weak to provide any real security (potentially misleading people about the real level of security). Future versions of MARS should provide at least 2-factor authorization, and encryption via dynamic session keys. Until that is implemented, use a secured VPN instead! And don't forget to *audit* it for security holes!

3.2. Setup Primary and Secondary Cluster Nodes

If you already use DRBD, you may migrate to MARS (or even back from MARS to DRBD) if you use *external*⁴ DRBD metadata (which is not touched by MARS).

3.2.1. Kernel and MARS Module

At the time of this writing, a small pre-patch for the Linux kernel is needed. It is trivial and consists mostly of `EXPORT_SYMBOL()` statements. The pre-patch must be applied to the kernel source tree before building your (custom) kernel. Future versions of MARS are planned to require no pre-patch anymore.

The MARS kernel module can be built in two different ways:

1. inplace in the kernel source tree: `cd block/ && git clone git://github.com/schoebel/mars`
2. as a separate kernel module, only for experienced⁵ sysadmins: see file `Makefile.dist` (tested with some older versions of Debian; may need some extra work with other distros).

Further / more accurate / latest instructions can be found in `README` and in `INSTALL`. You must not only install the kernel and the `mars.ko` kernel module to all of your cluster nodes, but also the `marsadm` userspace tool.

³It seems that the late internal resource allocation strategy of `xfs` (or another currently unknown reason) could be the reason for some resource deadlocks which appear only with `xfs` and only under *extremely* high IO load in combination with high memory pressure.

⁴*Internal* DRBD metadata should also work as long as the filesystem inside your block device / disk already exists and is not re-created. The latter would destroy the DRBD metadata, but even that will not hurt you really: you can always switch back to DRBD using *external* metadata, as long as you have some small spare space somewhere.

⁵You should be familiar with the problems arising from orthogonal combination of different kernel versions with different MARS module versions and with different `marsadm` userspace tool versions at the package management level. Hint: `modinfo` is your friend.

3. Quick Start Guide

3.2.2. Setup your Cluster Nodes

For your cluster, you need at least two nodes. In the following, they will be called A and B. In the beginning, A will have the **primary** role, while B will be your initial **secondary**. The roles may change later.

1. You must be **root**.
2. On each of A and B, create the `/mars/` mountpoint.
3. On each node, create an **ext4** filesystem on your separate disk / RAID set via `mkfs.ext4` (for requirements on size etc see section [Preparation: What you Need](#)).
4. On each node, mount that filesystem to `/mars/`. It is advisable to add an entry to `/etc/fstab`.
5. For security reasons, execute `chmod 0700 /mars` everywhere after `/mars/` has been mounted. If you forget this step, any following `marsadm` command will drop you a warning, but will fix the problem for you.
6. On node A, say `marsadm create-cluster`.
This must be done *exactly once*, on exactly one node of your cluster. Never do this twice or on different nodes, because that would create two different clusters which would have nothing to do with each other. The `marsadm` tool protects you against accidentally joining / merging two different clusters. If you accidentally created two different clusters, just unmount that `/mars/` partition and start over with step 3 at that node.
7. On node B, you must have a working **ssh** connection to node A (as **root**). Test it by saying `ssh A w` on node B. It should work without entering a password (otherwise, use `ssh-agent` to achieve that). In addition, `rsync` must be installed.
8. On node B, say `marsadm join-cluster A`
9. Only *after*⁶ that, do `modprobe mars` on each node.

3.3. Creating and Maintaining Resources

In the following example session, a block device `/dev/lv-x/mydata` (shortly called *disk*) must already exist on both nodes A and B, respectively, having the same⁷ size. For the sake of simplicity, the disk (underlying block device) as well as its later logical resource name as well as its later virtual device name will all be named uniformly by the same suffix `mydata`. In general, you might name each of them differently, but that is not recommended since it may easily lead to confusion in larger installations.

You may have already some data inside your disk `/dev/lv-x/mydata` at the initially primary side A. Before using it for MARS, it must be unused for any other purpose (such as being mounted, or used by DRBD, etc). MARS will require **exclusive access** to it.

1. On node A, say `marsadm create-resource mydata /dev/lv-x/mydata`.
As a result, a directory `/mars/resource-mydata/` will be created on node A, containing some symlinks. Node A will automatically start in the primary role for this resource. Therefore, a new pseudo-device `/dev/mars/mydata` will also appear after a few seconds. Note that the initial contents of `/dev/mars/mydata` will be exactly the same as in your pre-existing disk `/dev/lv-x/mydata`.
If you like, you may already use `/dev/mars/mydata` for mounting your already pre-existing data, or for creating a fresh filesystem, or for exporting via iSCSI, and so on. You may even do so before any other cluster node has joined the resource (so-called “standalone mode”). But you can also do so later after setup of (one or many) secondaries.

⁶In fact, you may already `modprobe mars` at node A after the `marsadm create-cluster`. Just don't do any of the `*-cluster` operations when the kernel module is loaded. All other operations should have no such restriction.

⁷Actually, the disk at the initially secondary side may be larger than that at the initially primary side. This will waste space and is therefore not recommended.

2. Wait a few seconds until the directory `/mars/resource-mydata/` and its symlink contents also appears on cluster node B. The command `marsadm wait-cluster` may be helpful.
3. On node B, say `marsadm join-resource mydata /dev/lv-x/mydata`. As a result, the initial full-sync from node A to node B should start automatically.



Of course, your old contents of your disk `/dev/lv-x/mydata` at side B (and *only* there!) is overwritten by the version from side A. Since you are an experienced sysadmin, you knew that, and it was just the effect you deliberately wanted to achieve. If you didn't check that your old contents didn't contain any valuable data (or if you accidentally provided a wrong disk device argument), it is too late now. The `marsadm` command checks that the disk device argument is really a block device, and that exclusive access to it is possible (as well as some further safety checks, e.g. matching sizes). However, MARS cannot know the *purpose* of your generic block device. MARS (as well as DRBD) is completely ignorant of the *contents* of a generic block device; it does not interpret it in any way. Therefore, you may use MARS (as well as DRBD) for mirroring Windows filesystems, or raw devices from databases, or virtual machines, or whatever.



Hint: by default, MARS uses the so-called “fast fullsync” algorithm. It works similar to `rsync`, first reading the data on both sides and computing an md5 checksum for each block. Heavy-weight data is only transferred over the long-distance network upon checksum mismatch. This is extremely fast if your data is already (almost) identical on both sides. Conversely, if you know in advance that your initial data is completely different on both sides, you may choose to switch off the fast fullsync algorithm via `echo 0 > /proc/sys/mars/do_fast_fullsync` in order to save the additional IO overhead and network latencies introduced by the separate checksum comparison steps.

4. Optionally, only for experienced sysadmins who *really* know what they are doing: if you will create a *new* filesystem on `/dev/mars/mydata` *after(!)* having created the MARS resource as well as *after* having already joined it on every replica, you may abandon the fast fullsync phase *before* creating the fresh filesystem, because the old content of `/dev/mars/mydata` will then be just garbage not used by the freshly created filesystem⁸. Then, and only then, you may say `marsadm fake-sync mydata` in order to abort the sync operation.



Never do a `fake-sync` unless you are **absolutely sure** that you really don't need to sync the data! Otherwise, you are *guaranteed* to have produced harmful inconsistencies. If you accidentally issued `fake-sync`, you may startover the fast full sync at your secondary side by saying `marsadm invalidate mydata` (analogously to the corresponding DRBD command).

3.4. Keeping Resources Operational

3.4.1. Logfile Rotation / Deletion

As explained in section [The Transaction Logger](#), all changes to your resource data are recorded in transaction logfiles residing on the `/mars/` filesystem. These files are always growing over time. In order to avoid filesystem overflow, the following must be done in regular time intervals:

⁸It is *vital* that the transaction logfile contents created by `mkfs` is *fully* propagated to the secondaries and then replayed there.

Analogously, another exception is also possible, but at your own risk (be careful, really!): when migrating your data from DRBD to MARS, and you have ensured that (1) at the end of using DRBD both your replicas were really equal (you should have checked that), and (2) before and after setting up any side of MARS (`create-resource` as well as `join-resource`) nothing has been written at all to it (i.e. no usage, neither of `/dev/lv/mydata` nor of `/dev/mars/mydata` has occurred in any way), the first transaction logfile `/mars/resource-mydata/log-00000001-$primary` created by MARS will be empty. Check whether this is really true! Then, and only then, you may also issue a `fake-sync`.

3. Quick Start Guide

1. `marsadm log-rotate all`

This starts appending to a new logfile on all of your resources. The logfiles are automatically numbered by an increasing 9-digit logfile number. This will suffice for many centuries even if you would logrotate once a minute. Practical frequencies for logfile rotation are more like once an hour, or every 10 minutes when having highly-loaded storage servers.

2. `marsadm log-delete-all all`

This determines all logfiles from all resources which are no longer needed (i.e. which are *fully* replayed, on *all* relevant secondaries). All superfluous logfiles are then deleted, including all copies on all secondaries.



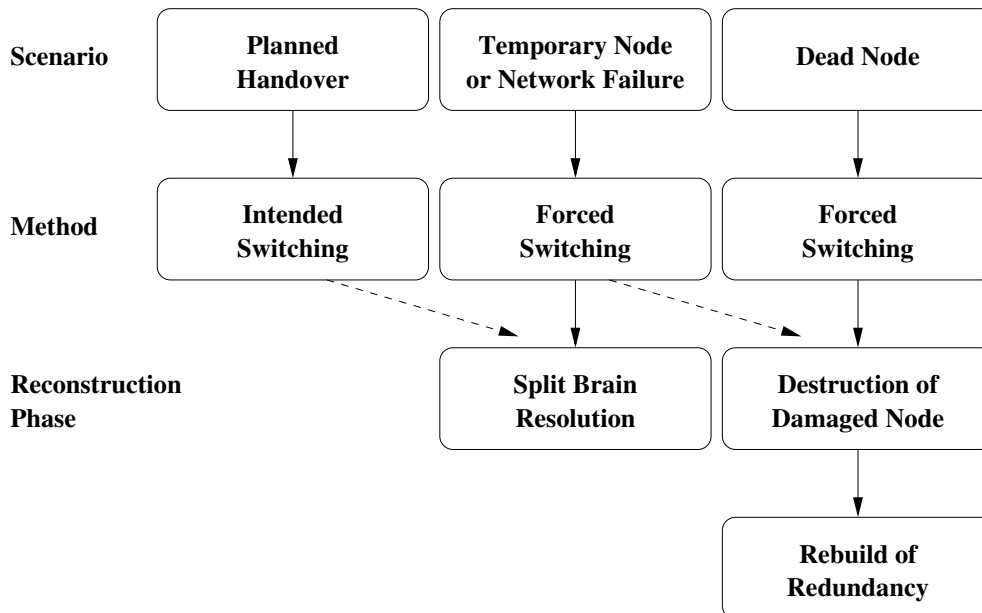
The current version of MARS deletes either *all* replicas of a logfile everywhere, or *none* of the replicas. This is a simple rule, but has the drawback that one node may hinder other nodes from freeing space in `/mars/`. In particular, the command `marsadm pause-replay $res` (as well as `marsadm disconnect $res`) will freeze the space reclamation in the whole cluster when the pause is lasting very long.



Best practice is to do both `log-rotate` and `log-delete-all` in a cron job. In addition, you should establish some regular monitoring of the free space present in the `/mars/` filesystem.

More detailed information about about avoidance of `/mars/` overflow is in section 4.4.

3.4.2. Switch Primary / Secondary Roles



In contrast to DRBD, MARS distinguishes between *intended* and *forced* switching. This distinction is necessary due to differences in the communication architecture (asynchronous communication vs synchronous communication, see sections 4.2 and 4.3).

Asynchronous communication means that (in worst case) a message may take (almost) arbitrary time in a distorted network to propagate to another node. As a consequence, the risk for accidentally creating an (unintended) split brain is increased (compared to a synchronous system like DRBD).

In order to minimize this risk, MARS has invested a lot of effort into an internal handover protocol when you start an *intended* primary switch.

3.4.2.1. Intended Switching / Planned Handover

Before starting a planned handover from your old primary A to a new primary B, you should check the replication of the resource. As a human, use `marsadm view mydata`. For scripting, use the macros from section 5.1.2 (see also section 5.3; an example can be found in `contrib/example-scripts/check-mars-switchable.sh`). The network should be OK, and the amount of replication delay should be as low as possible. Otherwise, handover may take a very long time.



Best practice is to **prepare a planned handover** by the following steps:

1. Check the network and the replication lag. It should be low (a few hundred megabytes, or a low number of gigabytes - see also the rough time forecast shown by `marsadm view mydata` when there is a larger replication delay, or directly access the forecast by `marsadm view-replinfo`).
2. Stop your application, then `umount /dev/mars/mydata` on host A.
3. When scripting, or when typing extremely fast, or for better safety, say `marsadm wait-umount mydata host B`. When your network is OK, the propagation of the device usage state⁹ should take only a few seconds. Otherwise, check for any network problems or any other problems.
4. On host B, wait until `marsadm view mydata` (or `view-diskstate`) shows `UpToDate`. It is possible to omit this step, but then you have no control on the duration of the handover, and in case of any transfer problems, disk space problems, etc you are potentially risking to produce a split brain (although `marsadm` will do its best to avoid it). Doing the wait by yourself, *before* starting `marsadm primary`, has a big advantage: you can abort the handover cycle at any time, just by re-mounting the device `/dev/mars/mydata` at the old primary A again, and by re-starting your application. Once you have started `marsadm primary` on host B, you might have to switch back, or possibly even via `primary --force` (see sections 3.4.2.2 and 3.4.3).

Switching the roles is very similar to DRBD: just issue the command

- `marsadm primary mydata`

on your formerly secondary node B.



The most important difference to DRBD: don't use an intermediate `marsadm secondary mydata` anywhere. Although it would be possible, it has some *disadvantages*. Always switch *directly*!



In contrast to DRBD, MARS remembers the designated primary, even when your system crashes and reboots. While in case of a crash you have to re-setup DRBD with commands like `drbdadm up ...; drbdadm primary ...`, MARS will automatically resume its former roles just by saying `modprobe mars`.



Another fundamental difference to DRBD: when the network is healthy, there can only exist *one* designated primary at a time (modulo some communication delays caused by the

⁹Notice that the usage check for `/dev/mars/mydata` on host B is based on the *open count* transferred from *another* node A. Since MARS is operating asynchronously (in contrast to DRBD), it may take some time until our node B knows that the device is no longer used at A. This can lead to a race condition if you automate an intended takeover with a script like `ssh root@A "umount /dev/mars/mydata"; ssh root@B "marsadm primary mydata"` because your second ssh command may be faster than the internal MARS symlink tree propagation (cf section 4.3). In order to prevent such races, you are strongly advised to use the command

- `marsadm wait-umount mydata`

on node B before trying to become primary. See also section 5.3.

3. Quick Start Guide

“eventually consistent” communication model, see section 4.2). By saying `marsadm primary mydata` on host B, **all other** hosts (including A) will **automatically go into secondary role** after a while!



You simply *don't need* an intermediate `marsadm secondary mydata` for planned hand-over!

Precondition for `marsadm primary` is that you are up, that means in attached and connected state (cf. `marsadm up`), and that any old primary (in this case A) does not use its `/dev/mars/mydata` device any longer, and that the network is healthy. If some (parts of) logfiles are not yet (fully) transferred to the new primary, you will need enough space on `/mars/` at the target side. If one of the preconditions described in section 6.2.2 is violated, `marsadm primary` may refuse to start.

The preconditions try to protect you from doing silly things, such as accidentally provoking a split brain error state. We try to avoid split brain as best as we can. Therefore, we distinguish between *intended* and *emergency* switching. Intended switching will try to avoid split brain *as best as it can*.



Don't *rely* on split brain avoidance, in particular when scripting any higher-level applications such as cluster managers (cf. section 5.3). `marsadm` does its best, but at least in case of (unnoticed) network outages / partitions (or *extremely, really extremely* slow / overloaded networks), an attempt to become `UpToDate` may fail. If you want to *ensure* that no split brain can result from intended primary switching, please obey the the best practices from above, and please give the `primary` command only after your secondary is *known*¹⁰ to be *really* `UpToDate` (see `marsadm wait-cluster` and `marsadm view` and other macros described in section 3.6).



A *very rough* estimation of the time to become `UpToDate` is displayed by `marsadm view mydata` or other macros (e.g. `view-replinfo`). However, on very flaky networks, the estimation may not only flicker much, but also be inaccurate.

3.4.2.2. Forced Switching

In case the connection to the old primary is lost for whatever reason, we just don't know anything about its *current* state (which may deviate from its *last known* state). The following command sequence will skip many checks and tell your node to become primary forcefully:

- `marsadm pause-fetch mydata`




- notice that this is similar to `drbdadm disconnect mydata` as you are probably used from DRBD. For better compatibility with DRBD, you may use the alternate syntax `marsadm disconnect mydata` instead. However, there is a subtle difference to DRBD: DRBD will drop *both* sides of its single bi-directional connection and no longer try to re-connect from any of both sides, while `pause-fetch` is equivalent to `pause-fetch-local`, which instructs only the *local* host to stop fetching logfiles. Other members of the cluster, including the former primary, are *not* instructed to do so. They may continue fetching logfiles over their own private TCP connections, potentially using many connections in parallel, and potentially even from any *other* member of the resource, if they think they can get the data from there. In order to instruct¹¹ *all* members of the resource to stop fetching logfiles, you may use `marsadm pause-fetch-global mydata` instead (cf section 6.2.2).

- `marsadm primary mydata --force`

¹⁰As noted in many places in this manual, checking this cannot be done by looking at the local state of a single cluster node. You have to check several nodes. `marsadm` can only check the *local* node reliably!

¹¹Notice that not all such instructions may arrive at all sites when the network is interrupted (or extremely slow).

-  this is the forceful switchover. Use `--force` only if you know what you are doing!

- `marsadm resume-fetch mydata`

- As such, the new primary does not really need this, because primaries are producing their own logfiles without need for fetching. This is only to undo the previous `pause-fetch`, in order to avoid future surprises when the new primary will somewhen change to secondary mode again (in the far-distant future), and you have forgotten to remember the fact that fetching had been switched off.

When using `--force`, many precondition checks and other internal checks are skipped, and in particular the internal handover protocol for split brain avoidance.

Therefore, use of `--force` is *likely* to **provoke a split brain**.



Split brain is always an **erroneous state** which should be never entered deliberately! Once you have entered it accidentally, you **must** resolve it ASAP (see section 3.4.3), otherwise you cannot operate your resource in the long term.

In order to impede you from giving an accidental `--force`, the precondition is different: `--force` works only in *locally disconnected* state. This is similar to DRBD.

Remember: `marsadm primary` without `--force` tries to prevent split brain as best as it can. Use of the `--force` option will almost *certainly* provoke a split brain, at least if the old primary continues to operate on its local `/dev/mars/mydata` device. Therefore, you are **strongly advised** to do this **only** after

1. `marsadm primary` without `--force` has failed *for no good reason*¹², and
2. You are sure you *really* want to switch, even when that eventually leads to a split brain. You also declare that you are willing to do *manual* split-brain resolution as described in section 3.4.3, or even destruction / reconstruction of a damaged node as described in section 3.4.4.



Notice: in case of *connection loss* (e.g. networking problems / network partitions), you may not be able to reliably detect whether a split brain actually resulted, or not.

Some Background In contrast to DRBD, split brain situations are handled differently by MARS. When two primaries are accidentally active at the same time, each of them writes into different logfiles `/mars/resource-mydata/log-000000001-A` and `/mars/resource-mydata/log-000000001-B` where the *origin* host is always recorded in the filename. Therefore, both nodes *can theoretically* run in primary mode independently from each other, at least for some time. They *might* even `log-rotate` independently from each other. However, this is really no good idea. The replication to third nodes will likely get stuck, and your `/mars/` filesystem(s) will eventually run out of space. Any further secondary node (when having $k > 2$ replicas) will certainly get into serious problems: it simply does not know which split-brain version it should follow. Therefore, you will certainly lose the actuality of your redundancy.



`marsadm secondary` is *strongly discouraged*. It tells the whole cluster that *nobody* is designated as primary any more. *All* nodes should go into secondary mode, globally. In the current version of MARS, the secondaries will no longer fetch any logfiles, since they don't know which version is the "right" one. Syncing is also not possible. When the device `/dev/mars/mydata` is in use somewhere, it will remain in *actual* primary mode during that time. As soon as the local `/dev/mars/mydata` is released, the node will *actually* go into secondary mode if it is no longer designated as primary. You should avoid it in advance by always *directly* switching over from one primary to another one, without intermediate **secondary** command. This is different from DRBD.

¹²Most reasons will be displayed by `marsadm` when it is rejecting the switchover.

3. Quick Start Guide



Split brain situations are detected *passively* by secondaries. Whenever a secondary detects that somewhere a split brain has happened, it refuses to replay any logfiles behind the split point (and also to fetch them when possible), or anywhere where something appears suspect or ambiguous. This tries to keep its local disk state always being consistent, but outdated with respect to any of the split brain versions. As a consequence, becoming primary may be impossible, because it cannot always know which logfiles are the correct ones to replay before `/dev/mars/mydata` can appear. The ambiguity must be resolved first.



If you *really* need the local device `/dev/mars/mydata` to disappear *everywhere* in a split brain situation, you don't need a *strongly discouraged* `marsadm secondary` command for this. `marsadm detach` or `marsadm down` can do it also, without destroying knowledge about the former designated primary.

3.4.3. Split Brain Resolution

Split brain can naturally occur during a long-lasting network outage (aka network partition) when you (forcefully) switch primaries inbetween, or due to final loss of your old primary node (fatal node crash) when not all logfile data had been transferred immediately before the final crash.



Remember that split brain is an **erroneous state** which must be resolved as soon as possible!

Whenever split brain occurs for whatever reason, you have two choices for resolution: either destroy one of your versions, or retain it under a different resource name.

In any of both cases, do the following steps ASAP:

1. **Manually** check which (surviving) version is the “right” one. Any error is up to you: destroying the wrong version is *your* fault, not the fault of MARS.
2. If you did not already switch your primary to the final destination determined in the previous step, do it now (see description in section 3.4.2.2). Don't use an intermediate `marsadm secondary` command (as known from DRBD): *directly* switch to the new designated primary!
3. On each non-right version (which you don't want to retain) which had been primary before, umount your `/dev/mars/mydata` or otherwise stop using it (e.g. stop iSCSI or other users of the device). Wait until each of them has actually left primary state and until their local logfile(s) have been fully written back to the underlying disk.
4. Wait until the network works again. All your (surviving) cluster nodes *must*¹³ be able to communicate with each other. If that is not possible, or if it takes too long, you may fall back to the method described in section 3.4.4, but do this only as far as necessary.

The next steps are different for different use cases:

Destroying a Wrong Split Brain Version Continue with the following steps, each on those cluster node(s) where you do not want to retain its split-brain version. In preference, start with the old “wrong” primaries first (see advice at the end of this section):

5. `marsadm invalidate mydata`

When no split brain is reported anymore after that (via `marsadm view all`), you are done. You need to repeat this on other secondaries only when necessary.

In very rare cases when things are screwed up very heavily (e.g. a partly destroyed `/mars/` partition), you may try an alternate method described in appendix C.

¹³If you are a MARS expert and you really know what you are doing (in particular, you can anticipate the effects of the Lamport clock and of the symlink update protocol including the “eventually consistent” behaviour including the not-yet-consistent intermediate states, see sections 4.2 and 4.3), you may deviate from this requirement.

Keeping a Split Brain Version On those cluster node(s) where you want to retain the version (e.g. for inspection purposes):

5. `marsadm leave-resource mydata`
6. After having done this on *all* those cluster nodes, check that the split brain is gone (e.g. by saying `marsadm view mydata`), as documented above. In very rare cases, you might also need a `log-purge-all` (see page 73).
7. Check that each underlying local disk `/dev/lv-x/mydata` is really usable afterwards, e.g. by test-mounting it (or `fsck` if you can afford it). If all is OK, don't forget to unmount it before proceeding with the next step.
8. Create a completely new MARS resource out of the underlying disk `/dev/lv-x/mydata` having a different name, such as `mynewdata` (see description in section 3.3 on page 24).

Keeping a Good Version When you had a secondary which did not participate in the split brain, but just got confused and therefore stopped replaying logfiles immediately before the split-brain point, it may very well happen¹⁴ that you don't need to do any action for it. When all wrong versions have disappeared from the cluster (by `invalidate` or `leave-resource` as described before), the confusion should be over, and the secondary should automatically resume tracking of the new unique version.

Please check that *all* of your secondaries are no longer stuck. You need to execute split brain resolution only for *stuck* nodes.



Hint / advice for $k > 2$ replicas: it is a good idea to start split brain resolution *first* with those (few) nodes which had been (accidentally) primary before, but are not the new designated primary. Usually, you had 2 primaries during split brain, so this will apply only to *one* of them. Leave the other one intact, by not unmounting `/dev/mars/mydata` at all, and keeping your applications running. Even during emergency mode, see section 4.4.2. *First* resolve the problem of the “wrong” primary(s) via `invalidate` or `leave-resource`. Wait for a short while. Then check the rest of your secondaries, whether they now are already following the new (unique) primary, and finally check whether the split brain warning reported by `marsadm view all` is gone everywhere. This way, you can often skip unnecessary invalidations of replicas.

3.4.4. Final Destruction of a Damaged Node

When a node has eventually died, do the following steps ASAP:

1. *Physically* remove the dead node from your network. Unplug all network cables! Failing to do so might provoke a disaster in case it somehow resurrects in an uncontrolled manner, such as a partly-damaged `/mars/` filesystem, a half-defective kernel, RAM / kernel memory corruption, disk corruption, or whatever. Don't risk any such unpredictable behaviour!
2. **Manually** check which of the surviving versions will be the “right” one. Any error is up to you: resurrecting an unnecessarily old / outdated version and/or destroying the newest / best version is *your* fault, not the fault of MARS.
3. If you did not already switch your primary to the final destination determined in the previous step, do it now (see description in section 3.4.2.2).
4. On a surviving node, but preferably *not* the new designated primary, give the following commands:
 - a) `marsadm --host=your-damaged-host down mydata`
 - b) `marsadm --host=your-damaged-host leave-resource mydata`

¹⁴In general, such a “good” behaviour cannot be guaranteed for all secondaries. Race conditions in complex networks may asynchronously transfer “wrong” logfile data to a secondary much earlier than conflicting “good” logfile data which will be marked “good” only in the *future*. It is impossible to predict this in advance.

3. Quick Start Guide



Check for misspellings, in particular the hostname of the dead node, and check the command syntax before typing return! Otherwise, you may forcefully destroy the wrong node!

5. In case any of the previous commands should fail (which is rather likely), repeat it with an additional `--force` option. Don't use `--force` in the first place, always try first without it!
6. Repeat the same with *all* resources which were formerly present at `your-damaged-host`.
7. Finally, say `marsadm --host=your-damaged-host leave-cluster` (optionally augmented with `--force`).

Now your surviving nodes should *believe* that the old node `your-damaged-host` does no longer exist, and that it does no longer participate in any resource.



Even if your dead node comes to life again in some way: always ensure that the mars kernel module cannot run any more. *Never* do a `modprobe mars` on a node marked as dead this way!

Further instructions for complicated cases are in appendix [D](#) and [E](#).

3.4.5. Online Resizing during Operation

You should have LVM or some other means of increasing the physical size of your disk (e.g. via firmware of some RAID controllers). The network must be healthy. Do the following steps:

1. Increase your local disks (usually `/dev/vg/mydata`) *everywhere* in the whole cluster. In order to avoid wasting space, increase them *uniformly* to the same size (when possible). The `lvresize` tool is documented elsewhere.
2. Check that all MARS switches are on. If not, say `marsadm up mydata` everywhere.
3. At the primary: `marsadm resize mydata`
4. If you have intermediate layers such as iSCSI, you may need some `iscsiadm` update or other command.
5. Now you may increase your filesystem. This is specific for the filesystem type and documented elsewhere.



Hint: the secondaries will start syncing the increased new part of the underlying primary disk. In many cases, this is not really needed, because the new junk data just does not care. If you are sure and if you know what you are doing, you may use `marsadm fake-sync mydata` to abort such unnecessary traffic.

3.5. The State of MARS

In general, MARS tries to *hide* any network failures from you as best as it can. After a network problem, any internal low-level socket connections are *transparently* tried to re-open ASAP, without need for sysadmin intervention. In difference to DRBD, network failures will *not* automatically alter the state of MARS, such as switching to `disconnected` after a `ko_timeout` or similar. From a high-level sysadmin viewpoint, communication may just take a very long time to succeed.

When the behaviour of MARS is different from DRBD, it is usually intended as a feature.

MARS is not only an **asynchronous** system at block IO level, but also **at control level**.


This is *necessary* because in a widely distributed long-distance system running on slow or even temporarily failing networks, actions may take a long time, and there may be many actions **started in parallel**.



Synchronous concepts are generally not sufficient for expressing that. Because of inherent asynchronicity and of dynamic creation / joining of resources, it is neither possible to comprehensively depict a complex distributed MARS system, nor a comprehensive standalone snippet of MARS, as a finite state transition diagram¹⁵.

Although MARS tries to *approximate* / *emulate* the synchronous control behaviour of DRBD at the interface level (`marsadm`) in many situations as best as it can, the *internal* control model is necessarily asynchronous. As an experienced sysadmin, you will be curious how it works in principle. When you know something about it, you will no longer be surprised when some (detail) behaviour is different from DRBD.

The general principle is an asynchronous 2-edge handshake protocol, which is used almost everywhere in MARS:



We have a binary todo switch, which can be either in state “on” or “off”. In addition, we have an actual response indicator, which is similar to an LED indicating the actual status. In our example, we imagine that both are used for controlling a big ventilator, having a huge inert mass. Imagine a big machine from a power plant, which is as tall as a human.

We start in a situation where the binary switch is off, and the ventilator is stopped. At point 1, we turn on the switch. At that moment, a big contactor will sound like “zonggg”, and a big motor will start to hum. At first you won’t hear anything else. It will take a while, say 1 minute, until the big wheel will have reached its final operating RPM, due to the huge inert mass. During that spin-up, the lights in your room will become slightly darker. When having reached the full RPM at point 2, your workplace will then be noisier, but in exchange your room lights will be back at ordinary strength, and the actual response LED will start to lit in order to indicate that the big fan is now operational.

Assume we want to turn the system off. When turning the todo switch to “off” at point 3, first nothing will seem to happen at all. The big wheel will keep spinning due to its heavy inert mass, and the RPM as well as the sound will go down only slowly. During spin-down, the actual response LED will stay illuminated, in order to warn you that you should not touch the wheel, otherwise you may get injured¹⁶. The LED will only go off after, say, 2 minutes, when the wheel has actually stopped at point 4. After that, the cycle may potentially start over again.

As you can see, all four possible cartesian product combinations between two boolean values are occurring in the diagram.

The same handshake protocol is used in MARS for communication between userspace and kernelspace, as well as for communication in the widely distributed system.

¹⁵Probably it could be possible to formally model MARS as a Petri net. However, complete Petri nets are tending to become very complex, and to describe lots of low-level details. Expressing hierarchy, in a top-down fashion, is cumbersome. We find no clue in trying to do so.

¹⁶Notice that it is only safe to access the wheel when *both* the switch and the LED are off. Conversely, if at least one of them is on, something is going on inside the machine. Transferred to MARS: always look at *both* the todo switch and the corresponding actual indicator in order to not miss something.

3.6. Inspecting the State of MARS

The main command for viewing the current state of MARS is

- `marsadm view mydata`

or its more specialized variant

- `marsadm view-$macroname mydata`

where *\$macroname* is one of the macros described in chapter 5, or a macro which has been written by yourself.

As always, you may replace the resource name `mydata` with the special keyword `all` in order to get the state of all locally joined resources, as well as a list of all those resources.



When using the variant `marsadm view all`, additionally the global communication status will be displayed. This helps humans in diagnosing problems.




Hint: use the compound command `watch marsadm view all` for continuous display of the current state of MARS. When starting this side-by-side in `ssh` terminal windows for all your cluster nodes, you can easily watch what's going on in the whole cluster.

4. Basic Working Principle

Even if you are impatient, please read this chapter. At the *surface*, MARS appears to be very similar to DRBD. It looks like almost being a drop-in replacement for DRBD.

When taking this naïvely, you could easily step into some trivial pitfalls, because the internal working principle of MARS is totally different from DRBD. Please forget (almost) anything you already know about the internal working principles of DRBD, and look at the very different working principles of MARS.

4.1. The Transaction Logger



MARS LCA2014 Presentation by Thomas Schöbel-Theuer

The basic idea of MARS is to record all changes made to your block device in a so-called **transaction logfile**. *Any* write request is treated like a transaction which changes the contents of your block device.

This is similar in concept to some database systems, but there exists no separate “commit” operation: *any* write request is acting like a commit.

The picture shows the flow of write requests. Let’s start with the primary node.

Upon submission of a write request on `/dev/mars/mydata`, it is first buffered in a *temporary* memory buffer.

The temporary memory buffer serves multiple purposes:

- It keeps track of the order of write operations.
- Additionally, it keeps track of the positions in the underlying disk `/dev/lv-x/mydata`. In particular, it detects when the same block is overwritten multiple times.
- During pending write operation, any concurrent reads are served from the memory buffer.

4. Basic Working Principle

After the write has been buffered in the temporary memory buffer, the main logger thread of the transaction logger creates a so-called *log entry* and starts an “append” operation on the transaction logfile. The log entry contains vital information such as the logical block number in the underlying disk, the length of the data, a timestamp, some header magic in order to detect corruption, the log entry sequence number, of course the data itself, and optional information like a checksum or compression information.

Once the log entry has been written through to the `/mars/` filesystem via `fsync()`, the application waiting for the write operation at `/dev/mars/mydata` is signalled that the write was successful.

This may happen even *before* the writeback to the underlying disk `/dev/lv-x/mydata` has started. Even when you power off the system right now, the information is not lost: it is present in the logfile, and can be reconstructed from there.

Notice that the order of log records present in the transaction log defines a total order among the write requests which is *compatible* to the partial order of write requests issued on `/dev/mars/mydata`.

Also notice that despite its sequential nature, the transaction logfile is typically *not* the performance bottleneck of the system: since appending to a logfile is almost purely sequential IO, it runs much faster than random IO on typical datacenter workloads.

In order to reclaim the temporary memory buffer, its content must be written back to the underlying disk `/dev/lv-x/mydata` somewhen. After writeback, the temporary space is freed. The writeback can do the following optimizations:

1. writeback may be in *any* order; in particular, it may be *sorted* according to ascending sector numbers. This will reduce the average seek distances of magnetic disks in general.
2. when the same sector is overwritten multiple times, only the “last” version need to be written back, skipping some intermediate versions.

In case the primary node crashes during writeback, it suffices to replay the log entries from some point in the past until the end of the transaction logfile. It does no harm if you accidentally replay some log entries twice or even more often: since the replay is in the original total order, any temporary inconsistency is *healed* by the logfile application.



In mathematics, the property that you can apply your logfile twice to your data (or even as often as you want), is called **idempotence**. This is a very desirable property: it ensures that nothing goes wrong when replaying “too much” / starting your replay “too early”. Idempotence is even more beneficial: in case anything should go wrong with your data on your disk (e.g. IO errors), replaying your logfile once more often may¹ even **heal** some defects. Good news for desperate sysadmins forced to work with flaky hardware!

The basic idea of the asynchronous replication of MARS is rather simple: just transfer the logfiles to your secondary nodes, and replay them onto their copy of the disk data (also called *mirror*) in the same order as the total order defined by the primary.

Therefore, a mirror of your data on any secondary may be outdated, but it always corresponds to some version which was valid in the past. This property is called **anytime consistency**².



As you can see in the picture, the process of transferring the logfiles is *independent* from the process which replays the logfiles onto the data at some secondary site. Both processes can be switched on / off separately (see commands `marsadm {dis,}connect` and `marsadm {pause,resume}-replay` in section 6.2.2). This may be *exploited*: for example, you may replicate your logfiles as soon as possible (to protect against catastrophic failures), but deliberately

¹Miracles cannot be guaranteed, but *higher chances* and *improvements* can be expected (e.g. better chances for `fsck`).

²Your secondary nodes are always consistent in themselves. Notice that this kind of consistency is a *local* consistency model. There exists no global consistency in MARS. Global consistency would be practically impossible in long-distance replication where Einstein’s law of the speed of light is limiting global consistency. The front-cover pictures showing the planets Earth and Mars tries to lead your imagination away from global consistency models as used in “DRBD Think(tm)”, and try to prepare you mentally for local consistency as in “MARS Think(tm)”.

wait one hour until it is replayed (under regular circumstances). If your data inside your filesystem `/mydata/` at the primary site is accidentally destroyed by `rm -rf /mydata/`, you have an old copy at the secondary site. This way, you can substitute *some parts*³ of conventional backup functionality by MARS. In case you need the actual version, just replay in “fast-forward” mode (similar to old-fashioned video tapes).



Future versions of MARS Full are planned to also allow “fast-backward” rewinding, of course at some cost.

4.2. The Lamport Clock

MARS is always *asynchronously* communicating in the distributed system on *any* topics, even strategic decisions.

If there were a *strict* global consistency model, which would be roughly equivalent to a standalone model, we would need *locking* in order to serialize conflicting requests. It is known for many decades that *distributed locks* do not only suffer from performance problems, but they are also cumbersome to get them working reliably in scenarios where nodes or network links may fail at any time.

Therefore, MARS uses a very different consistency model: **Eventually Consistent**.



Notice that the network bottleneck problems described in section 2.1 are *demanding* an “eventually consistent” model. You have **no chance** against natural laws, like Einstein’s laws. In order to cope with the problem area, you have to *invest some additional effort*. Unfortunately, asynchronous communication models are more tricky to program and to debug than simple strictly consistent models. In particular, you *have to cope with* additional **race conditions inherent** to the “eventually consistent” model. In the face of the laws of the universe, motivate yourself by looking at the graphics at the cover page: the planets are a *symbol* for what you have to do!



Example: the asynchronous communication protocol of MARS leads to a different behaviour from DRBD in case of **network partitions** (temporary interruption of communication between some cluster nodes), because MARS *remembers* the old state of remote nodes over long periods of time, while DRBD knows absolutely nothing about its peers in disconnected state. Sysadmins familiar with DRBD might find the following behaviour unusual:

Event	DRBD Behaviour	MARS Behaviour
1. the network partitions	automatic disconnect	nothing happens, but replication lags behind
2. on A: <code>umount \$device</code>	works	works
3. on A: <code>{drbd,mars}adm secondary</code>	works	works
4. on B: <code>{drbd,mars}adm primary</code>	works, split brain happens	refused because B believes that A is primary
5. the network resumes	automatic connect attempt fails	communication automatically resumes

If you intentionally want to switch over (and to produce a split brain as a side effect), the following variant must be used with MARS:

Event	DRBD Behaviour	MARS Behaviour
1. the network partitions	automatic disconnect	nothing happens, but replication lags behind
2. on A: <code>umount \$device</code>	works	works
3. on A: <code>{drbd,mars}adm secondary</code>	works	works (but <i>not recommended!</i>)
4. on B: <code>{drbd,mars}adm primary</code>	split brain, but nobody knows	refused because B believes that A is primary
5. on B: <code>marsadm disconnect</code>	-	works, nothing happens
6. on B: <code>marsadm primary --force</code>	-	works, split brain happens on B, but A doesn’t know
7. on B: <code>marsadm connect</code>	-	works, nothing happens
8. the network resumes	automatic connect attempt fails	communication resumes, A now detects the split brain

³Please note that MARS cannot *fully* substitute a backup system, because it can keep only *physical* copies, and does not create logical copies.

4. Basic Working Principle

In order to implement the consistency model “eventually consistent”, MARS uses a so-called Lamport⁴ clock. MARS uses a special variant called “physical Lamport clock”.

The physical Lamport clock is another almost-realtime clock which *can* run independently from the Linux kernel system clock. However, the Lamport clock tries to remain as near as possible to the system clock.

Both clocks can be queried at any time via `cat /proc/sys/mars/lamport_clock`. The result will show both clocks in parallel, in units of seconds since the Unix epoch, with nanosecond resolution.

When there are no network messages at all, both the system clock and the Lamport clock will show almost the same time (except some minor differences of a few nanoseconds resulting from the finite processor clock speed).

The physical Lamport clock works rather simple: *any* message on the network is augmented with a Lamport time stamp telling when the message was *sent* according to the local Lamport clock of the sender. Whenever that message is received by some receiver, it checks whether the time ordering relation would be violated: whenever the Lamport timestamp in the message would claim that the sender had sent it *after* it arrived at the receiver (according to drifts in their respective local clocks), something must be wrong. In this case, the local Lamport clock of the *receiver* is advanced shortly after the sender Lamport timestamp, such that the time ordering relation is no longer violated.

As a consequence, any local Lamport clock may precede the corresponding local system clock. In order to avoid accumulation of deltas between the Lamport and the system clock, the Lamport clock will run slower after that, possibly until it reaches the system clock again (if no other message arrives which sets it forward again). After having reached the system clock, the Lamport clock will continue with “normal” speed.

MARS uses the local Lamport clock for anything where other systems would use the local system clock: for example, timestamp generation in the `/mars/` filesystem. Even symlinks created there are timestamped according to the Lamport clock. Both the kernel module and the userspace tool `marsadm` are always operating in the timescale of the Lamport clock. Most importantly, all timestamp comparisons are always carried out with respect to Lamport time.



Bigger differences between the Lamport and the system clock can be annoying from a human point of view: when typing `ls -l /mars/resource-mydata/` many timestamps may appear as if they were created in the “future”, because the `ls` command compares the output formatting against the system clock (it does not even know of the existence of the MARS Lamport clock).



Always use `ntp` (or another clock synchronization service) in order to pre-synchronize your system clocks as close as possible. Bigger differences are not only annoying, but may lead some people to wrong conclusions and therefore even lead to bad human decisions!

In a professional datacenter, you should use `ntp` anyway, and you should monitor its effectiveness anyway.



Hint: many internal logfiles produced by the MARS kernel module contain Lamport timestamps written as numerical values. In order to convert them into human-readable form, use the command `marsadm cat /mars/5.total.status` or similar.

4.3. The Symlink Tree



The symlink tree as described here will be replaced by another representation in future versions of MARS. Therefore, don’t do any scripting by directly accessing symlinks! Use the primitive macros described in section 5.1.2.

The current `/mars/` filesystem container format contains not only transaction logfiles, but also acts as a generic storage for (persistent) state information. Both configuration information

⁴Published in the late 1970s by Leslie Lamport, also known as inventor of L^AT_EX.

and runtime state information are currently stored in symlinks. Symlinks are “misused⁵” in order to represent some `key -> value` pairs.



It is not yet clear / decided, but there is a *chance* that the *concept* of `key -> value` pairs will be retained in future versions of MARS. Instead of being represented by symlinks, another representation will be used, such that hopefully the `key` part will remain in the form of a pathname, even if there were no longer a physical representation in an actual filesystem.



A fundamentally different behaviour than DRBD: when your DRBD primary crashed some time ago, and now comes up again, you have to setup DRBD again by a sequence of commands like `modprobe drbd; drbdadm up all; drbdadm primary all` or similar. In contrast, MARS needs only `modprobe mars` (after `/mars/` has been mounted by `/etc/fstab`). The *persistence* of the symlinks residing in `/mars/` will automatically remember your previous state, even if some your resources were primary while others were secondary (mixed operations). You don’t need to do any actions in order to “restore” a previous state, no matter how “complex” it was.

(Almost) all symlinks appearing in the `/mars/` directory tree are automatically replicated throughout the whole cluster, provided that the cluster `uuids` are equal⁶ at all sites. Thus the `/mars/` directory forms some kind of *global namespace*.

In order to avoid name clashes, each pathname created at node A follows a convention: the node name A should be a suffix of the pathname. Typically, internal MARS names follow the scheme `/mars/something/myname-A`. When using the expert command `marsadm {get,set}-link` (which will likely be replaced by something else in future MARS releases), you should follow the best practice of systematically using pathnames like `/mars/userspace/myname-A` or similar. As a result, each node will automatically get informed about the state at any other node, like B when the corresponding information is recorded on node B under the name `/mars/userspace/myname-B` (context-dependent names).



Experts only: the symlink replication works generically. You might use the `/mars/userspace/` directory in order to place your own symlink there (for whatever purpose, which need not have to do with MARS). However, the symlinks are likely to disappear. Use `marsadm {get,set}-link` instead. There is a chance that these abstract commands (or variants thereof) will be retained, by acting on the new data representation in future, even if the old symlink format will vanish some day.



Important: the convention of placing the **creator host name** inside your pathnames should be used wherever possible. The name part is a kind of “ownership indicator”. It is crucial that no other host writes any symlink not “belonging” to him. Other hosts may read foreign information as often as they want, but never modify them. This way, your cluster nodes are able to *communicate* with each other via symlink / information updates.

Although experts might create (and change) the current symlinks with userspace tools like `ln -s`, you should use the following `marsadm` commands instead:

- `marsadm set-link myvalue /mars/userspace/mykey-A`
- `marsadm delete-file /mars/userspace/mykey-A`

There are many reasons for this: first, the `marsadm set-link` command will automatically use the Lamport clock for symlink creation, and therefore will avoid any errors resulting from a “wrong” system clock (as in `ln -s`). Second, the `marsadm delete-file` (which also deletes

⁵This means, the symlink targets need not be other files or directories, but just any values like integers or strings.

⁶This is protection against accidental “merging” of two unrelated clusters which had been created at different times with different `uuids`.

4. Basic Working Principle

symlinks) works on the *whole cluster*. And finally, there is a chance that this will work in future versions of MARS even after the symlinks have vanished.

What's the difference? If you would try to remove your symlink locally by hand via `rm -f`, you will be surprised: since the symlink has been replicated to the other cluster nodes, it will be re-transferred from there and will be resurrected locally after some short time. This way, you cannot delete any object reliably, because your whole cluster (which may consist of many nodes) remembers all your state information and will "correct" it whenever "necessary".

In order to solve the deletion problem, MARS uses some internal deletion protocol using auxiliary symlinks residing in `/mars/todo-global/`. The deletion protocol ensures that all replicas get deleted in the whole cluster, and only thereafter the auxiliary symlinks in `/mars/todo-global/` are also deleted eventually.

You may update your already existing symlink via `marsadm set-link some-other-value /mars/userspace/mykey-A`. The new value will be propagated throughout the cluster according to a **timestamp comparison protocol**: whenever node B notices that A has a *newer* version of some symlink (according to the Lamport timestamp), it will replace its elder version by the newer one. The opposite does *not* work: if B notices that A has an elder version, just nothing happens. This way, the timestamps of symlinks can only progress in forward direction, but never backwards in time.

As a consequence, symlink updates made "by hand" via `ln -sf` may get lost when the local system clock is much more earlier than the Lamport clock.

When your cluster is fully connected by the network, the last timestamp will finally win everywhere. Only in case of network outages leading to *network partitions*, some information may be *temporarily inconsistent*, but only for the duration of the network outage. The timestamp comparison protocol in combination with the Lamport clock and with the persistence of the `/mars/` filesystem will automatically heal any temporary inconsistencies as soon as possible, even in case of temporary node shutdown.

The meaning of some internal MARS symlinks residing in `/mars/` will be hopefully documented in section 8.4 some day.

4.4. Defending Overflow of `/mars/`

This section describes an important difference to DRBD. The metadata of DRBD is allocated *statically* at *creation time* of the resource. In contrast, the MARS transaction logfiles are allocated *dynamically* at *runtime*.

This leads to a potential risk from the perspective of a sysadmin: what happens if the `/mars/` filesystem runs out of space?

No risk, no fun. If you want a system which survives long-lasting network outages while keeping your replicas always consistent (anytime consistency), you *need* dynamic memory for that. It is *impossible* to solve that problem using static memory⁷.

Therefore, DRBD and MARS have different application areas. If you just want a simple system for mirroring your data over short distances like a crossover cable, DRBD will be a suitable choice. However, if you need to replicate over longer distances, or if you need higher levels of reliability even when multiple failures may accumulate (such as network loss during a *resync* of DRBD), the transaction logs of MARS can solve that, but at some *cost*.

4.4.1. Countermeasures

4.4.1.1. Dimensioning of `/mars/`

The first (and most important) measure against overflow of `/mars/` is simply to dimension it large enough to survive longer-lasting problems, at least one weekend.

Recommended size is at least one dedicated disk, residing at a hardware RAID controller with BBU (see section 3.1). During normal operation, that size is needed only for a small fraction, typically a few percent or even less than one percent. However, it is your **safety margin**. Keep it high enough!

⁷The bitmaps used by DRBD don't preserve the *order* of write operations. They cannot do that, because their space is $O(k)$ for some constant k . In contrast, MARS preserves the order. Preserving the order as such (even when only *facts* about the order were recorded without recording the actual data contents) requires $O(n)$ space where n is infinitely growing over time.

4.4.1.2. Monitoring

The next (equally important) measure is **monitoring in userspace**.

Following is a list of countermeasures both in userspace and in kernelspace, in the order of “defensive walling”:

1. Regular userspace monitoring must throw an INFO if a certain freespace limit l_1 of /mars/ is undershot. Typical values for l_1 are 30%. Typical actions are automated calls of `marsadm log-rotate all` followed by `marsadm log-delete-all all`. You have to implement that yourself in sysadmin space.
2. Regular userspace monitoring must throw a WARNING if a certain freespace limit l_2 of /mars/ is undershot. Typical values for l_2 are 20%. Typical actions are (in addition to `log-rotate` and `log-delete-all`) alarming human supervisors via SMS and/or further stronger automated actions.



Frequently large space is occupied by files stemming from debugging output, or from other programs or processes. A hot candidate is “forgotten” removal of debugging output to /mars/. Sometimes, an `rm -rf $(find /mars/ -name “*.log”)` can work miracles.



Another source of space hogging is a “forgotten” `pause-sync` or `disconnect`. Therefore, a simple `marsadm connect-global all` followed by `marsadm resume-replay-global all` may also work miracles (if you didn’t want to freeze some mirror deliberately).



If you just wanted to freeze a mirror at an outdated state for a very long time, you simply *cannot* do that without causing infinite growth of space consumption in /mars/. Therefore, a `marsadm leave-resource $res` at *exactly that(!)* secondary site where the mirror is frozen, can also work miracles. If you want to automate this in userspace, be careful. It is easy to get unintended effects when choosing the wrong site for `leave-resource`.



Hint: you can / should start some of these measures even earlier at the INFO level (see item 1), or even earlier.

3. Regular userspace monitoring must throw an ERROR if a certain freespace limit l_3 of /mars/ is undershot. Typical values for l_3 are 10%. Typical actions are alarming the CEO via SMS and/or even stronger automated actions. For example, you may choose to automatically call `marsadm leave-resource $res` on some or all secondary nodes, such that the primary will be left alone and now has a chance to really delete its logfiles because no one else is any longer potentially needing it.
4. First-level kernelspace action, automatically executed when `/proc/sys/mars/required_free_space_4_gb + /proc/sys/mars/required_free_space_3_gb + /proc/sys/mars/required_free_space_2_gb + /proc/sys/mars/required_free_space_1_gb` is undershot:
a warning will be issued.
5. Second-level kernelspace action, automatically executed when `/proc/sys/mars/required_free_space_3_gb + /proc/sys/mars/required_free_space_2_gb + /proc/sys/mars/required_free_space_1_gb` is undershot:
all locally secondary resources will delete local copies of transaction logfiles which are no longer needed locally. This is a desperate action of the kernel module.
6. Third-level kernelspace action, automatically executed when `/proc/sys/mars/required_free_space_2_gb + /proc/sys/mars/required_free_space_1_gb` is undershot:

4. Basic Working Principle

all locally secondary resources will stop fetching transaction logfiles. This is a more desperate action of the kernel module. You don't want to get there (except for testing).

7. Last desperate kernelspace action when all else has failed and `/proc/sys/mars/required_free_space_1_gb` is undershot:

all locally primary resources will enter **emergency mode** (see description below in section 4.4.2). This is the most desperate action of the kernel module. You don't want to get there (except for testing).

In addition, the kernel module obeys a general global limit `/proc/sys/mars/required_total_space_0_gb` + the sum of all of the above limits. When the *total size* of `/mars/` undershots that sum, the kernel module refuses to start at all, because it assumes that it is senseless to try to operate MARS on a system with such low memory resources.



The current level of emergency kernel actions may be viewed at any time via `/proc/sys/mars/mars_emergency_mode`.

4.4.1.3. Throttling

The last measure for defense of overflow is **throttling your performance pigs**.

Motivation: in rare cases, some users with `ssh` access can do *very* silly things. For example, some of them are creating their own backups via user-cron jobs, and they do it every 5 minutes. Some example guy created a zip archive (almost 1GB) by regularly copying his old zip archive into a new one, then appending deltas to the new one, and finally deleting the old archive. Every 5 minutes. Yes, every 5 minutes, although almost never any new files were added to the archive. Essentially, he copied over his archive, for nothing. This led to massive bulk write requests, for ridiculous reasons.

In general, your hard disks (or even RAID systems) allow much higher write IO rates than you can ever transport over a standard TCP network from your primary site to your secondary, at least over longer distances (see use cases for MARS in chapter 2). Therefore, it is easy to create a such a high write load that it will be *impossible* to replicate it over the network, *by construction*.

Therefore, we *need* some mechanism for throttling bulk writers whenever the network is weaker than your IO subsystem.



Notice that DRBD will *always* throttle your writes whenever the network forms a bottleneck, due to its synchronous operation mode. In contrast, MARS allows for buffering of performance peaks in the transaction logfiles. *Only when* your buffer in `/mars/` runs short (cf subsection 4.4.1.1), MARS will start to throttle your application writes.

There are a lot of screws named `/proc/sys/mars/write_throttle_*` with the following meaning:

write_throttle_start_percent Whenever the used space in `/mars/` is below this threshold, no throttling will occur at all. Only when this threshold is exceeded, throttling will start *slowly*. Typical values for this are 60%.

write_throttle_end_percent Maximum throttling will occur once this space threshold is reached, i.e. the throttling is now at its maximum effect. Typical values for this are 90%. When the actual space in `/mars/` lies between `write_throttle_start_percent` and `write_throttle_end_percent`, the strength of throttling will be interpolated linearly between the extremes. In practice, this should lead to an equilibrium between new input flow into `/mars/` and output flow over the network to secondaries.

write_throttle_size_threshold_kb (readonly) This parameter shows the internal strength calculation of the throttling. Only write⁸ requests exceeding this size (in KB) are throttled at all. Typically, this will hurt the bulk performance pigs first, while leaving ordinary users (issuing small requests) unaffected.

⁸Read requests are never throttled at all.

`write_throttle_ratelimit_kb` Set the global IO rate in KB/s for those write requests which are throttled. In case of strongest⁹ throttling, this parameter determines the input flow into `/mars/`. The default value is 5.000 KB/s. Please adjust this value to your application needs and to your environment.

`write_throttle_rate_kb` (readonly) Shows the current rate of exactly those requests which are actually throttled (in contrast to *all* requests).

`write_throttle_cumul_kb` (logically readonly) Same as before, but the cumulative sum of all throttled requests since startup / reset. This value can be reset from userspace in order to prevent integer overflow.

`write_throttle_count_ops` (logically readonly) Shows the cumulative number of throttled requests. This value can be reset from userspace in order to prevent integer overflow.

`write_throttle_maxdelay_ms` Each request is delayed at most for this timespan. Smaller values will improve the responsiveness of your userspace application, but at the cost of potentially retarding the requests not sufficiently.

`write_throttle_minwindow_ms` Set the minimum length of the measuring window. The measuring window is the timespan for which the average (throughput) rate is computed (see `write_throttle_rate_kb`). Lower values can increase the responsiveness of the controller algorithm, but at the cost of accuracy.

`write_throttle_maxwindow_ms` This parameter must be set sufficiently much greater than `write_throttle_minwindow_ms`. In case the flow of throttled operations pauses for some natural reason (e.g. switched off, low load, etc), this parameter determines when a completely new rate calculation should be started over¹⁰.

4.4.2. Emergency Mode and its Resolution

When `/mars/` is almost full and there is really absolutely no chance of getting rid of any local transaction logfile (or free some space in any other way), there is only one exit strategy: stop creating new logfile data.

This means that the ability for replication gets lost.

When entering emergency mode, the kernel module will execute the following steps for all resources where the affected host is acting as a primary:

1. Do a kind of “logrotate”, but create a *hole* in the sequence of transaction logfile numbers. The “new” logfile is left empty, i.e. no data is written to it (for now). The hole in the numbering will prevent any secondaries from replaying any logfiles behind the hole (should they ever contain some data, e.g. because the emergency mode has been left again). This works because the secondaries are regularly checking the logfile numbers for contiguity, and they will refuse to replay anything which is not contiguous. As a result, the secondaries will be left in a consistent, but outdated state (at least if they already were consistent before that).
2. The kernel module writes back all data present in the temporary memory buffer (see figure in section 4.1). This may lead to a (short) delay of user write requests until that has finished (typically fractions of a second or a few seconds). The reason is that the temporary memory buffer must not be increased in parallel during this phase (race conditions).
3. After the temporary memory buffer is empty, all local IO requests (whether reads or writes) are directly going to the underlying disk. This has the same effect as if MARS would not be present anymore. Transaction logging does no longer take place.

⁹In case of lighter throttling, the input flow into `/mars/` may be higher because small requests are not throttled.

¹⁰Motivation: if requests would pause for one hour, the measuring window could become also an hour. Of course, that would lead to completely meaningless results. Two requests in one hour is “incorrect” from a human point of view: we just have to ensure that averages are computed with respect to a reasonable maximum time window in the magnitude of 10s.

4. Basic Working Principle

4. Any sync from any secondary is stopped ASAP. In case they are resuming their sync somewhen later, they will start over from the beginning (position 0).

In order to leave emergency mode, the sysadmin should do the following steps:

1. Free enough space. For example, delete any foreign files on `/mars/` which have nothing to do with MARS, or resize the `/mars/` filesystem, or whatever.
2. If `/proc/sys/mars/mars_reset_emergency` is not set, now it is time to set it. Normally, it should be already set.
3. Notice: as long as not enough space has been freed, a message containing “EMEGENCY MODE HYSTERESIS” (or similar) will be displayed by `marsadm view all`. As a consequence, any sync will be automatically halted. This applies to freshly invoked syncs also, for example created by `invalidate` or `join-resource`.
4. On the secondaries, use `marsadm invalidate $res` in order to request updating your outdated mirrors.
5. On the primary: `marsadm log-delete-all all`
6. As soon as enough space has been freed everywhere to leave the `EMEGENCY MODE HYSTERESIS`, sync should really start. Until that it had been halted.

Alternatively, there is another method by roughly following the instructions from appendix C, but in a slightly different order. In this case, do `leave-resource` everywhere on *all* secondaries, but *don't* start the `join-resource` phase *for now*. Then cleanup all your secondaries via `log-purge-all`, and finally `log-delete-all all` at the primary, and wait until the emergency has vanished everywhere. Only after that, re-`join-resource` your secondaries.



Expert advice for $k = 2$ replicas: this means you had only 1 mirror per resource before the overflow happened. Provided that you have enough space on your LVMs and on `/mars/`, and provided that transaction logging has automatically restarted after `leave-resource` and `log-purge-all`, you can recover redundancy by creating a *new* replica via `marsadm join-resource $res` on a *third* node. Only after the initial full sync has finished there, run `join-resource` at your original mirror. This way, you will always retain at least one **consistent mirror** somewhere. After all is up-to-date, you can delete the superfluous mirror by `marsadm leave-resource $res` and reclaim the disk space from its underlying LVM disk.



If you already have $k > 2$ replicas in total, it may be a wise idea to prefer the `leave-resource ; log-purge-all ; join-resource` method in front of `invalidate` because it does not invalidate *all* your replicas at the same time (when handled properly in the right order).

5. The Macro Processor

`marsadm` comes with a customizable macro processor. It can be used for high-level complex display of the state of MARS (so-called *complex macros*), as well as for low-level display of lots of individual state values (so-called *primitive macros*).

From the commandline, any macro can be called via `marsadm view-$macroname mydata`. The short form `marsadm view mydata` is equivalent to `marsadm view-default mydata`.



In general, the command `marsadm view-$macroname all` will first call the macro *\$macroname* in a loop for *all* resources we are a *member locally*. Finally, a trailing macro *\$macroname-global* will be called with an empty `%{res}` argument, provided that such a macro is defined. This way, you can produce per-resource output followed by global output which does not depend on a particular resource.

5.1. Predefined Macros

The macro processor is a very flexible and versatile tool for **customizing**. You can create your own macros, but probably the rich set of predefined macros is already sufficient for your needs.

5.1.1. Predefined Complex and High-Level Macros

The following predefined complex macros try to address the information needs of humans. Use them only in scripts when you are prepared about the fact that the output format may change during development of MARS.

Notice: the definitions of predefined complex macros may be updated in the course of the MARS project. However, the primitive macros recursively called by the complex ones will be hopefully rather stable in future (with the exception of bugfixes). If you want to retain an old / outdated version of a complex macro, just check it out from git, follow the instructions in section 5.2, and preferably give it a different name in order to avoid confusion with the newer version. In general, it should be possible to use old macros with newer versions of `marsadm`¹.

default This is equivalent to `marsadm view mydata` without *-macroname* suffix. It shows a one-line status summary for each resource, optionally followed by informational lines such as progress bars whenever a sync or a fetch of logfiles is currently running. The status line has the following fields:

```
%{res}    resource name.
%include{diskstate} see diskstate macro below.
%include{replstate} see replstate macro below.
%include{flags}    see flags macro below.
%include{role}    see role macro below.
%include{primarynode} see primarynode macro below.
%include{commstate} see commstate macro below.
```

After that, optional lines such as progress bars are appearing only when something unusual is happening. These lines are subject to future changes. For examples, wasted disk space due to missing `resize` is reported when `%{threshold}` is exceeded.

¹You might need to check out also old versions of further macros and adapt their names, whenever complex macros call each other.

5. The Macro Processor

`1and1` or `default-1and1` A variant of `default` for internal use by 1&1 Internet AG. You may call this complex macro by saying `marsadm view-1and1 all`.



Note: the `marsadm view-1and1` command has been intensely tested in Spring 2014 to produce exactly the same output than the 1&1 internal² tool `marsview`³



Customization via your own macros (see section 5.2) is explicitly encouraged by the developer. It would be nice if a vibrant user community would emerge, helping each other by exchange of macros.



Hint: in order to produce your own customized inspection / monitoring tools, you may ask the author for an official reservation of a macro sub-namespace such as `*-yourcompanyname`. You will be fully responsible for your own reserved namespace and can do with it whatever you want. The official MARS release will guarantee that *no name clashes* with your reserved sub-namespace will occur in future.

`default-global` Currently, this just calls `comminfo` (see below). May be extended in future.

`diskstate` Shows the status of the underlying disk device, in the following order of precedence⁴:

- `NotJoined` (cf `%get-disk{}`) No underlying disk device is configured.
- `NotPresent` (cf `%disk-present{}`) The underlying disk device (as configured, see `marsadm view-get-disk`) does not exist or the device node is not accessible. Therefore MARS cannot work. Check that LVM or other software is properly configured and running.
- `Detached` (cf `InConsistent`, `NeedsReplay`, `%todo-attach{}`, `%is-attach{}`) The underlying disk is willingly switched off (see `marsadm detach`), and it actually is no longer opened by MARS.
- `Detaching` (cf `%todo-attach{}` and `%is-attach{}`) Access to the underlying disk is switched off, but actually not yet `close()`d by MARS. This can happen for a long time on a primary when other secondaries are accessing the disk remotely for syncing.
- `DefectiveLog[description-text]` (cf `%replay-code{}`) Typically this indicates an md5 checksum error in a transaction logfile, or another (hardware / filesystem) defect. This occurs extremely rarely in practice, but has been observed more frequently during a massive failure of air conditioning in a datacenter, when disk temperatures raised to more than 80° Celsius. Notice that a secondary **refuses** to apply any knowingly defective logfile data to the disk. Although this message is *not directly* referring to the underlying disk, it is mentioned here because of its superior **relevance** for the `diskstate`. A damaged transaction logfile will always affect the *actuality* of the disk, but not its *integrity* (by itself). What to do in such a case?

²In addition to allow for customization, the macro processor is also meant as an exit strategy for removing dependencies from non-free software. **Please put your future macros also under GPL!**

³There are some subtle differences: numbers are displayed in a different precision, some bug fixes in the macro version (which might have occurred *in the meantime*) may lead to different output as a side effect from bug fixes in *predefined* macros, because the original `marsview` command is currently not actively maintained. Documentation of `marsview` can be found in the corresponding manpage, see `man marsview`. By construction, this is also the (unmaintained) documentation of `marsadm view-1and1` and other `-1and1` macros. Notice that all `*-1and1` macros are not officially supported by the developer of MARS, and they may disappear in a future major release. However, they could be useful for your own customization macros.

⁴When an earlier list item is displayed, no combinations with following items are possible. This kind of “hiding effect” can lead to an *information loss*. In order to get a non-lossy picture from the state of your system, please look at the `flags` which are able to display cartesian combinations of more detailed internal states.

1. When the damage is only at one of your secondaries, you should first ensure that the primary has a good logfile after a `marsadm log-rotate`, then try `marsadm invalidate` at the damaged secondary. It is crucial that the primary has a fresh correct logfile behind the error position, and that it is continuing to operate correctly.
2. When *all* of your secondaries are reporting `DefectiveLog`, the primary could have *produced* a damaged logfile (e.g. in RAM, in a DMA channel, etc) while continuing to operate, and all of your secondaries got that defective logfile. After `marsadm log-delete-all all`, you can check this by comparing the `md5sum` of the first primary logfile (having the lowest serial number) with the versions on your replicas. The problem is that you don't know whether the primary side has a silent corruption on any of its disks, or not. You will need to take an operational decision whether to switchover to a secondary via `primary --force`, or whether to continue operation at the primary and `invalidate` your secondaries.
3. When the original primary is affected in a very bad way, such that it crashed badly and afterwards even recovery of the *primary* is impossible⁵ due to this error (which typically occurs extremely rarely, observed two times during 7 millions of operating hours on defective hardware), you need to take an operational decision between the following alternatives:
 - a) switchover to a former secondary via `primary --force`, producing a split brain, and producing some (typically small) data loss. However, integrity is more important than actuality in such an extreme case.
 - b) deconstruction of the resource at *all* replicas via `leave-resource --force`, running `fsck` or similar tools by hand at the underlying disks, selecting the best replica out of them, and finally reconstructing the resource again.
 - c) restore your backup.

NoAttach (cf `%is-attach{}`) The underlying disk is currently not opened by MARS. Reasons may be that the kernel module is not loaded, or an exclusive `open()` is currently not possible because somebody else has already opened it.

InConsistent (cf `%is-consistent{}`) A logfile replay and/or sync is known to be needed / or to complete (e.g. after `invalidate` has started) in order to restore local consistency (for details, look at `flags`).



Hint: in the current implementation of MARS, this will never happen on secondaries during ordinary replay (but only when either sync has not yet finished, or when the *initial* logfile replay after the sync has not yet finished), because the ordinary logfile replay always maintains anytime consistency once a consistent state had been reached.



Only in case of a primary node crash, and *only* after attempts have failed to become primary again (e.g. IO errors, etc), this *can* (but need not) mean that something went wrong. Even in such an extremely unlikely event, chances are high that `fsck` can fix any remaining problems (and, of course, you can also switchover to a former secondary).

⁵In such a rare case, the *original primary* (but not any other host) **refuses** to come up during recovery with *his own* logfile originally produced by *himself*. This is not a bug, but saves you from incorrectly assuming that your original primary disk were consistent - it is *known* to be inconsistent, but recovery is impossible due to the damaged logfile. Thus *this one* replica is trapped by defective hardware. The other replicas shouldn't.



When this message appears, simply start MARS again (e.g. `modprobe mars; marsadm up all`), in whatever role you are intending. This will *automatically* try to replay any necessary transaction logfile(s) in order to fix the inconsistency. Only if the automatic fix fails and this message persists for a long time without progress, you *might* have a problem. Typically, as observed at a large installation at 1&1, this happens extremely rarely, and then typically indicates that your hardware is likely to be defective.

OutDated[FR] (cf `%work-reached{}`) Only at secondaries. Tells whether it is *currently known* that the disk has any lag-behind when compared to the *currently known* state of the current designated primary (if there exists one). Only meaningful if a current designated primary exists. Notice that this kind of status display is subject to *natural races*, for example when new logfile data has been produced in parallel, or network propagation is very slow. Additional information is in brackets:

[F] Fetch is known to be needed.
 [R] Replay is known to be needed.
 [FR] Both are known to be needed.

WriteBack (cf `%is-primary{}`) Appears only at actual primaries (whether designated or not), when the writeback from the RAM buffer is active (see section 4.1)

Recovery (cf `%todo-primary{}`) Appears only at the designated primary before it actually has become primary. Similar to database recovery, this indicates the recovery phase after a crash⁶.

EmergencyMode (cf `%is-emergency{}`) A current designated primary exists, and it is known that this host has entered emergency mode. See section 4.4.2.

UpToDate Displayed when none of the above has been detected.

diskstate-1and1 A variant for internal use by 1&1 Internet AG. See above note.

replstate Shows the status of the replication in the following order of precedence:

ModuleNotLoaded (cf `%is-module-loaded{}`) No kernel module is loaded, and as a consequence no `/proc/sys/mars/` does exist.

UnResponsive (cf `%is-alive{%{host}}`) The main thread `mars_light` did not do any noticeable work for more than `%{window}` (default 30) seconds. Notice that this may happen when deleting *extremely* large logfiles (up to hundreds of gigabytes or terabytes). If this happens for a *very* long time, you should check whether you might need a reboot in order to fix the hang. The time window may be changed by `--window=$seconds`.

NotJoined (cf `%get-disk{}`) No underlying disk device is configured for this resource.

NotStarted (cf `%todo-attach{}`) Replication has not been started.

- When the current host is designated as a primary, the rest of the precedence list looks as follows:

EmergencyMode (cf. `%is-emergency{}`) See section 4.4.2.

Replicating (cf. `%is-primary{}`) Primary mode has been entered.

NotYetPrimary (catchall) This means the current host *should* act as a primary (see `marsadm primary` or `marsadm primary --force`), but currently doesn't (yet). This happens during logfile replay, before primary mode is actually entered. Notice that replay of very big logfiles may take a long time.

⁶In some cases, `primary --force` may also trigger this message.

- When the current host is *not* designated as a primary:
 - PausedSync (cf. `%sync-rest{}` and `%todo-sync{}`) Some data needs to be synced, but sync is currently switched off. See `marsadm {pause,resume}-sync`.
 - Syncing (cf. `%is-sync{}`) Sync is currently running.
 - PausedFetch (cf. `%todo{fetch}`) Fetch is currently switched off. See `marsadm {pause,resume}-fetch`.
 - PausedReplay (cf. `%todo{replay}`) Replay is currently switched off. See `marsadm {pause,resume}-replay`.
 - NoPrimaryDesignated (cf. `%get-primary{}`) A **secondary** command has been given somewhere in the cluster. Thus no designated primary exists. All resource members are in state **Secondary** or try to approach it. Sync and other operations are not possible. This state is therefore not recommended.
 - PrimaryUnreachable (cf. `%is-alive{}`) A current designated primary has been set, but this host has not been remotely updated for more than 30 seconds (see also `--window=$seconds`).
 - Replaying (catchall) None of the previous conditions have triggered.

`replstate-1and1` A variant for internal use by 1&1 Internet AG. See above note.

flags For each of disk, consistency, attach, sync, fetch, and replay, show exactly one character. Each character is either a capital one, or the corresponding lowercase one, or a dash. The meaning is as follows:

disk/device: **D** = the device `/dev/mars/mydata` is present, **d** = only the underlying disk `/dev/lv-x/mydata` is present, **-** = none present / configured.

consistency: this relates to the *underlying disk*, not to `/dev/mars/mydata`! **C** = locally consistent, **c** = maybe inconsistent (no guarantee), **-** = cannot determine. Notice: this does not tell anything about *actuality*. Notice: like the other flags, this flag is subject to races and therefore should be relied on only in *detached* state! See also description of macro `is-consistent` below.

attach: **A** = attached, **a** = currently trying to attach/detach but not yet ready (intermediate state), **-** = attach is switched off.

sync: **S** = sync finished, **s** = currently syncing, **-** = sync is switched off.

fetch: **F** = according to knowlege, fetched logfiles are up-to-date, **f** = currently fetching (some parts of) a logfile, **-** = fetch is switched off.

replay: **R** = all fetched logfiles are replayed, **r** = currently replaying, **-** = replay is switched off.

`flags-1and1` A variant for internal use by 1&1 Internet AG.

`todo-role` Shows the *designated* state: **None**, **Primary** or **Secondary**.

role Shows the *actual* state: **None**, **NotYetPrimary**, **Primary**, **RemainsPrimary**, or **Secondary**. Any differences to the designated state are indicated by a prefix to the keyword **Primary**: **NotYet** means that it *should* become primary, but actually hasn't. Vice versa, **Remains** means that it *should* leave primary state in order to become secondary, but actually cannot do that because the `/dev/mars/mydata` device is currently in use

	<code>%todo-primary{}</code> == 0	<code>%todo-primary{}</code> == 1
<code>%is-primary{}</code> == 0	None / Secondary	NotYetPrimary
<code>%is-primary{}</code> == 1	RemainsPrimary	Primary

`role-1and1` A variant for internal use by 1&1 Internet AG.

`primarynode` Display (none) or the hostname of the designated primary.

5. The Macro Processor

primarynode-1and1 A variant for internal use by 1&1 Internet AG.

commstate When the last metadata communication to the designated primary is longer ago than $\${window}$ (see also `--window=seconds` option), display that age in human readable form. See also primitive macro `%alive-age{}`.

syncinfo Shows an informational progress bar when sync is running. Intended for humans. Scripts should not rely on any details from this. Scripts may use this only as an *approximate* means for detecting progress (when comparing the *full* output text to a prior version and finding *any* difference, they may conclude that some progress has happened, how small whatsoever).

syncinfo-1and1 A variant for internal use by 1&1 Internet AG.

replinfo Shows an informational progress bar when fetch is running. This should not be used for scripting at all, because it contains realtime information in human-readable form.

replinfo-1and1 A variant for internal use by 1&1 Internet AG.

fetch-line Additional details, called by `replinfo`. Shows the amount of data to be fetched, as well as the current transfer rate and a very rough estimation of the future duration. When primitive macros `%fetch-age{}` or `%fetch-lag{}` exceed $\${window}$, their values are also displayed for human informational purposes. See description of these primitive macros.

replay-line Additional details, called by `replinfo`. Shows the amount of data to be replayed, as well as the current replay rate and a very rough estimation of the future duration. When primitive macro `%replay-age{}` exceeds $\${window}$, it is also displayed for human informational purposes.

comminfo When the network communication is in an unusual condition, display it. Otherwise, don't produce any output.

5.1.2. Predefined Primitive Macros

5.1.2.1. Intended for Humans

In the following, shell glob notation `{a,b}` is used to document similar variants of similar macros in a single place. When you actually call the macro, you must choose one of the possible variants (excluding the braces).

the-err-msg Show reported errors for a resource. When the resource argument is missing or empty, show global error information.

all-err-msg Like before, but show all information including those which are OK. This way, you get a list⁷ of *all* potential error information present in the system.

{all,the}-wrn-msg Show all / reported warnings in the system.

{all,the}-inf-msg Show all / reported informational messages in the system.

{all,the}-msg Show all / reported messages regardless of its classification.

{all,the}-global-msg Show global messages not associated with any resource (the resource argument of the `marsadm` command is ignored in this case).

{all,the}-global-`{inf,wrn,err}`-msg Dito, but more specific.

{all,the}-pretty-`{global-,}{inf-,wrn-,err-,}`msg Dito, but show numerical timestamps in a human readable form.

{all,the}-`{global-,}{inf-,wrn-,err-,}`count Instead of showing the messages, show their count (number of lines).

⁷The list may be extended in future versions of MARS.

errno-text This macro takes 1 argument, which must represent a Linux `errno` number, and converts it to human readable form (similar to the C `strerror()` function).

todo-`{attach, sync, fetch, replay, primary}` Shows a boolean value (0 or 1) indicating the current state of the corresponding todo switch (whether on or off). The meaning of todo switches is illustrated in section 3.5.

get-resource-`{fat, err, wrn}` Access to the internal error status files. This is not an official interface and may thus change at any time without notice. Use this only for human inspection, not for scripting!



These macros, as well as the error status files, are likely to disappear in future versions of MARS. They should be used for debugging only. At least when merging into the upstream Linux kernel, only the `*-msg` macros will likely survive.

get-resource-`{fat, err, wrn}`-count Dito, but get the number of lines instead of the text.

replay-code Indicate the current state of logfile replay / recovery:

(empty)	Unknown.
0	No replay is currently running.
1	Replay is currently running.
2	Replay has successfully stopped.
<0	See Linux <code>errno</code> code. Typically this indicates a damaged logfile, or another filesystem error at <code>/mars</code> .

is-`{attach, sync, fetch, replay, primary, module-loaded}` Shows a boolean value (0 or 1) indicating the *actual* state, whether the corresponding action has been actually carried out, or not (yet). Notice that the values indicated by `is-*` may differ from the `todo-*` values when something is not (yet) working. More explanations can be found in section 3.5.

is-split-brain Shows whether split brain (see section 3.4.3) has been detected, or not.

is-consistent Shows whether the *underlying disk* is in a locally consistent state, i.e. whether it *could* be (potentially) detached and then used for read-only test-mounting⁸. Don't confuse this with the consistency of `/dev/mars/mydata`, which is by construction *always* locally consistent once it has appeared⁹. By construction of MARS, the disk of secondaries will *always* remain in a locally consistent state once the initial sync has finished as well as the initial logfile replay. Notice that local consistency does not necessarily imply actuality (see high-level explanation in section 2.1.2).

is-emergency Shows whether emergency mode (see section 4.4.2) has been entered for the named resource, or not.

rest-space (global, no resource argument necessary) Shows the *logically* available space in `/mars/`, which may deviate from the physically available space as indicated by the `df` command.

get-`{disk, device}` Show the name of the underlying disk, or of the `/dev/mars/mydata` device (if it is available).

⁸Notice that the *writeback* at the primary side is out-of-order by default, for performance reasons. Therefore, the underlying disk is only guaranteed to be consistent when there is no data left to be written back. Notice that this condition is racy by construction. When your primary node crashes during writeback and then comes up again, you must do a `modprobe mars` first in order to automatically replay the transaction logfiles, which will automatically heal such temporary inconsistencies.

⁹Exceptions are possible when using `marsadm fake-sync`. Even in split brain situations, `marsadm primary --force` tries to prevent any further potential exception as best as it can, by not letting `/dev/mars/mydata` to appear and by insisting on split brain resolution first. In future implementations, this might change if more pressure is put on the developer to sacrifice consistency in preference to not waiting for a full logfile replay.

5. The Macro Processor

{disk,device}-present Show (as a boolean value) whether the underlying disk, or the `/dev/mars/mydata` device, is available.

device-opened Show (as a number) how often `/dev/mars/mydata` has been actually opened, e.g. by `mount` or by some processes like `dd`, or by iSCSI, etc.

5.1.2.2. Intended for Scripting

While complex macros may output a whole bunch of information, the following primitive macros are outputting exactly one value. They are intended for script use (cf. section 5.3). Of course, curious humans may also try them :)

In the following, shell glob notation `{a,b}` is used to document similar variants of similar macros in a single place. When you actually call the macro, you must choose one of the possible variants (excluding the braces).

Name Querying

cluster-members Show a newline-separated list of all host names participating in the cluster.

resource-members Show a newline-separated list of all host names participating in the particular resource `{res}`. Notice that this may be a subset of `%cluster-members{}`.

{my,all}-resources Show a newline-separated list of either all resource names existing in the cluster, or only those where the current host `{host}` is member. Optionally, you may specify the hostname as a parameter, e.g. `%my-resources{otherhost}`.

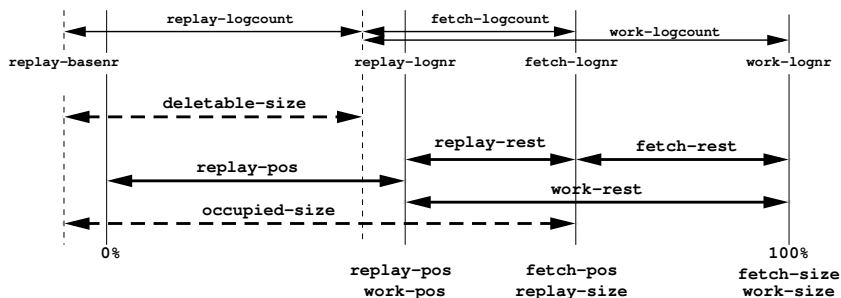


Figure 5.1.: overview on amounts / cursors

Amounts of Data Inquiry The following macros are meaningful for both primary and secondary nodes:

deletable-size Show the total amount of *locally present* logfile data which *could* be deleted by `marsadm log-delete-all mydata`. This differs almost always from both `replay-pos` and `occupied-size` due to granularity reasons (only whole logfiles can be deleted). Units are *bytes*, not kilobytes.

occupied-size Show the total amount of *locally present* logfile data (sum of all file sizes). This is often roughly approximate to `fetch-pos`, but it may differ vastly (in both directions) when logfiles are not completely transferred, when some are damaged, during split brain, after a `join-resource` / `invalidate`, or when the resource is in emergency mode (see section 4.4.2).

disk-size Show the size of the underlying local disk in bytes.

resource-size Show the logical size of the resource in bytes. When this value is lower than `disk-size`, you are wasting space.

device-size At a primary node, this may differ from `resource-size` only for a very short time during the `resize` operation. At secondaries, there will be no difference.

The following macros are only meaningful for secondary nodes. By information theoretic limits, they can only tell what is *locally known*. They **cannot** reflect the “true (global) state”¹⁰ of a cluster, in particular during network partitions.

`{sync,fetch,replay,work}-size` Show the total amount of data which is / was to be processed by either sync, fetch, or replay. `work-size` is equivalent to `fetch-size`. `replay-size` is equivalent to `fetch-pos` (see below). Units are *bytes*, not kilobytes.

`{sync,fetch,replay,work}-pos` Show the total amount of data which is already processed (current “cursor” position). `work-pos` is equivalent to `replay-pos`.



The 0% point is the *locally contiguous* amount of data since the last `create-resource`, `join-resource`, or `invalidate`, or since the last emergency mode, but possibly shortened by `log-deletes`. Notice that the 0% point may be different on different cluster nodes, because their resource history may be different or non-contiguous during split brain, or after a `join-resource`, or after `invalidate`, or during / after emergency mode.

`{sync,fetch,replay,work}-rest` Shows the difference between `*-size` and `*-pos` (amount of work to do). `work-rest` is therefore the difference between `fetch-size` and `replay-pos`, which is the *total* amount of work to do (regardless whether to be fetched and/or to be replayed).

`{sync,fetch,replay,work}-reached` Boolean value indicating whether `*-rest` dropped down to zero¹¹.

`{fetch,replay,work}-threshold-reached` Boolean value indicating whether `*-rest` dropped down to `%{threshold}`, which is pre-settable by the `--threshold=size` command line option (default is 10 MiB). In asynchronous use cases of MARS, this should be preferred over `*-reached` for *human display*, because it produces less flickering by the inevitable replication delay.

`{fetch,replay,work}-almost-reached` Boolean value indicating whether `*-rest` *almost / approximately* dropped down to zero. The default is that at least 990 permille are reached. In asynchronous use cases of MARS, this can be preferred over `*-reached` for *human display* only, because it produces less flickering by the inevitable replication delay. However, don’t base any decisions on this!

`{sync,fetch,replay,work}-percent` The cursor position `*-pos` as a percentage of `*-size`.

`{sync,fetch,replay,work}-permille` The cursor position `*-pos` as permille of `*-size`.

¹⁰Notice that according to Einstein’s law, and according to observations by Lamport, the concept of “true state” does not exist at all in a distributed system. Anything you can know in a distributed system is always local knowledge, which races with other (remote) knowledge, and may be outdated at *any* time.

¹¹Recall from chapter 2 that MARS (in its current stage of development) does only guarantee local consistency, but cannot guarantee actuality in all imaginable situations. Notice that a general notion of “actuality” is *undefinable* in a widely distributed system at all, according to Einstein’s laws.

Let’s look at an example. In case of a node crash, and after the node is up again, a `modprobe mars` has to occur, in order to replay the transaction logs of MARS again. However, at the recovery phase before, the journaling `ext4` filesystem `/mars/` may have rolled back some internal symlink updates which have occurred immediately before the crash. MARS is relying on the fact that journaling filesystems like `ext4` should do their recovery in a consistent way, possibly by sacrificing actuality a little bit. Therefore, the above macros cannot guarantee to deliver true information about what is persisted at the moment.

Notice that there are further potential caveats.

In case of `{sync,fetch}-reached`, MARS uses `bio` callbacks resp. `fdatsync()` by default, thus the underlying storage layer has *told* us that it *believes* it has committed the data in a reboot-safe way. Whether this is *really* true does not depend on MARS, but on the lower layers of the storage hierarchy. There exists hardware where this claim is known to be wrong under certain circumstances, such as certain hard disk drives in certain modes of operation. Please check the hardware for any violations of storage semantics under certain circumstances such as power loss, and check information sources like magazines about the problem area. Please notice that such a problem, if it exists at all, is independent from MARS. It would also exist if you wouldn’t use MARS on the same system.

5. The Macro Processor

`{sync,fetch,replay,work}-rate` Show the current throughput in bytes¹² per second. `work-rate` is the *maximum* of `fetch-rate` and `replay-rate`.

`{sync,fetch,replay,work}-remain` Show the *estimated* remaining time for completion of the respective operation. This is just a very raw guess. Units are seconds.

`summary-vector` Show the colon-separated CSV value `%replay-pos{ }:%fetch-pos{ }:%fetch-size{ }`.

`replay-basennr` Get currently first reachable logfile number (see figure 5.1 on page 52). Only for curious humans or for debugging / monitoring - don't base any decisions on this. Use the `*-{pos,size}` macros instead.

`{replay,fetch,work}-lognr` Get current logfile number of replay or fetch position, or of the currently known last reachable number (see figure 5.1 on page 52). Only for curious humans or for debugging / monitoring - don't base any decisions on this. Use the `*-{pos,size}` macros instead.

`{replay,fetch,work}-logcount` Get current number of logfiles which are already replayed, or are already fetched, or are to be applied in total (see figure 5.1 on page 52). Only for curious humans or for debugging / monitoring - don't base any decisions on this. Use the `*-{rest}` macros instead.

`alive-timestamp` Tell the Lamport Unix timestamp (seconds since 1970) of the last metadata communication to the designated primary (or to any other host given by the first argument). Returns `-1` if no such host exists.

`{fetch,replay,work}-timestamp` Tell the Lamport Unix timestamp (seconds since 1970) when the last progress has been made. When no such action exists, `-1` is returned. `%work-timestamp{hostname}` is the maximum of `%fetch-timestamp{hostname}` and `%replay-timestamp{hostname}`. When the parameter `hostname` is empty, the local host will be reported (default). Example usage: `marsadm view all --macro='%replay-timestamp{t` shows the timestamp of the last reported¹³ writeback action at the designated primary.

`{alive,fetch,replay,work}-age` Tell the number of seconds since the last respective action, or `-1` if none exists.

`{alive,fetch,replay,work}-lag` Report the time difference (in seconds) between the last *known* action at the local host and at the designated primary (or between any other hosts when 2 parameters are given). Returns `-1` if no such action exists at any of the two hosts. Attention! This need not reflect the *actual* state in case of networking problems. Don't draw wrong conclusions from a high `{fetch,replay}-lag` value: it could also mean that simply no write operation at all has occurred at the primary side for a long time. Conversely, a low lag value does not imply that the replication is recent: it may refer to *different* write operations at each of the hosts; therefore it only tells that *some* progress has been made, but says nothing about the amount of the progress.

Misc Informational Status

`get-primary` Return the name of the current designated primary node as locally known.

`actual-primary` (deprecated) try to determine the name of the node which *appears* to be the actual primary. This only a *guess*, because it is not generally unique in split brain situations! Don't use this macro. Instead, use `is-primary` on those nodes you are interested in. The explanations from section 3.5 also apply to `get-primary` versus `actual-primary` analogously.

¹²Notice that the internal granularity reported by the kernel may be coarser, such as KiB. This interfaces abstracts away from kernel internals and thus presents everything in byte units.

¹³Updates of this information are occurring with lower frequency than actual writebacks, for performance reasons. The metadata network update protocol will add further delays. Therefore, the accuracy is only in the range of minutes.

- `is-alive` Boolean value indicating whether all other nodes participating in `mydata` are reachable / healthy.
- `uuid` (global) Show the unique identifier created by `create-cluster` or by `create-uuid`. Hint: this is immutable, and it is firmly bound to the `/mars/` filesystem. It can only be destroyed by deleting the whole filesystem (see section 6.2).
- `tree` (global) Indicate symlink tree version (see section 4.3).

Experts Only The following is for hackers who know what they are doing. The following is not officially supported.

`wait-{is, todo}-{attach, sync, fetch, replay, primary}-{on, off}` This may be used to program some useful waiting conditions in advanced macro scripts. Use at your own risk!

5.2. Creating your own Macros

In order to create your own macros, you could start writing them from scratch with your favorite ASCII text editor. However, it is much easier to take an existing macro and to customize it to your needs. In addition, you can learn something about macro programming by looking at the existing macro code.

Go to a new empty directory and say

- `marsadm dump-macros`

in order to get the most interesting complex macros, or say

- `marsadm dump-all-macros`

in order to additionally get some primitive macros which could be customized if needed. This will write lots of files `*.tpl` into your current working directory.

Any modified or new macro file should be placed either into the current working directory `./`, or into `$HOME/.marsadm/`, or into `/etc/marsadm/`. They will be searched in this order, and the first match will win. When no macro file is found, the built-in version will be used if it exists. This way, you may override builtin macros.

Example: if you have a file `./mymacro.tpl` you just need to say `marsadm view-mymacro mydata` in order to invoke it in the resource context `mydata`.

5.2.1. General Macro Syntax

Macros are simple ASCII text, enriched with calls to other macros.

ASCII text outside of comments are copied to the output verbatim. Comments are skipped. Comments may have one of the following well-known forms:

- `#` skipped text until / including next newline character
- `//` skipped text until / including next newline character
- `/*` skipped text including any newline characters `*/`
- denoted as Perl regex: `\\n\s*` (single backslash directly followed by a newline character, and eating up any whitespace characters at the beginning of the next line) Hint: this may be fruitfully used to structure macros in a more readable form / indentation.

Special characters are always initiated by a backslash. The following pre-defined special character sequences are recognized:

- `\n` newline
- `\r` return (useful for DOS compatibility)
- `\t` tab
- `\f` formfeed

5. The Macro Processor

- `\b` backspace
- `\a` alarm (bell)
- `\e` escape (e.g. for generating ANSI escape sequences)
- `\` followed by anything else: assure that the next character is taken verbatim. Although possible, please don't use this for escaping letters, because further escape sequences might be pre-defined in future. Best practice is to use this only for escaping the backslash itself, or for escaping the percent sign when you don't want to call a macro (protect against evaluation), or to escape a brace directly after a macro call (verbatim brace not to be interpreted as a macro parameter).
- All other characters stand for their own. If you like, you should be able to produce XML, HTML, JSON and other ASCII-based output formats this way.

Macro calls have the following syntax:

- `%macroname {arg1}{arg2}{argn}`
- Of course, arguments may be empty, denoted as `{}`
- It is possible to supply more arguments than required. These are simply ignored.
- There must be always at least 1 argument, even for parameterless macros. In such a case, it is good style to leave it empty (even if it is actually ignored). Just write `%parameterlessmacro{}` in such a case.
- `%{varname}` syntax: As a special case, the macro name may be empty, but then the first argument must denote a previously defined variable (such as assigned via `%let{varname}{myvalue}`), or a pre-defined standard variable like `%{res}` for the current resource name, see later paragraph 5.2.3).
- Of course, parameter calls may be (almost) arbitrarily nested.
- Of course, the *correctness* of nesting of braces must be generally obeyed, as usual in any other macro processor language. General rule: for each opening brace, there must be exactly one closing brace somewhere afterwards.

These rules are hopefully simple and intuitive. There are currently no exceptions. In particular, there is no special infix operator syntax for arithmetic expressions, and therefore no operator precedence rules are necessary. You have to write nested arithmetic expressions always in the above prefix syntax, like `#{7}{%+{2}{3}}` (similar to non-inverse polish notation).



When deeply nesting macros and their braces, you may easily find yourself in a feeling like in the good old days of Lisp. Use the above backslash-newline syntax to indent your macros in a readable and structured way. Fortunately, modern text editors like (x)emacs or vim have modes for dealing with the correctness of nested braces.

5.2.2. Calling Builtin / Primitive Macros

Primitive macros can be called in two alternate forms:

- `%primitive-macroname{something}`
- `%macroname{something}`

When using the `%primitive-*` form, you *explicitly disallow* interception of the call by a `*.tpl` file. Otherwise, you may override the standard definition even of primitive macros by your own template files.



Notice that `%call{}` conventions are used in such a case. The parameters are passed via `{0} ... {n}` variables (see description below).

Standard MARS State Inspection Macros These are already described in section 5.1.2. When calling one of them, the call will simply expand to the corresponding value.

Example: `%get-primary{}` will expand to the hostname of the current designated primary node.

Further MARS State Inspection Macros

Variable Access Macros

- `%let{varname}{expression}` Evaluates both *varname* and the *expression*. The *expression* is then assigned to *varname*.
- `%let{varname}{expression}` Evaluates both *varname* and the *expression*. The *expression* is then appended to *varname* (concatenation).
- `%{varname}` Evaluates *varname*, and outputs the value of the corresponding variable. When the variable does not exist, the empty string is returned.
- `%{++}{varname}` or `%{varname}{++}` Has the obvious well-known side effect e.g. from C or Java. You may also use `--` instead of `++`. This is handy for programming loops (see below).
- `%dump-vars{}` Writes all currently defined variables (from the currently active scope) to `stderr`. This is handy for debugging.

CSV Array Macros

- `%{varname}{delimiter}{index}` Evaluates all arguments. The contents of *varname* is interpreted as a comma-separated list, delimited by *delimiter*. The *index*'th list element is returned.
- `%set{varname}{delimiter}{index}{expression}` Evaluates all arguments. The contents of the old *varname* is interpreted as a comma-separated list, delimited by *delimiter*. The *index*'th list element is the assigned to, or substituted by, *expression*.

Arithmetic Expression Macros The following macros can also take more than two arguments, carrying out the corresponding arithmetic operation in sequence (it depends on the operator whether this accords to the associative law).

- `%+{arg1}{arg2}` Evaluates the arguments, interprets them as numbers, and adds them together.
- `%-{arg1}{arg2}` Subtraction.
- `%*{arg1}{arg2}` Multiplication.
- `%/{arg1}{arg2}` Division.
- `%%{arg1}{arg2}` Modulus.
- `%&{arg1}{arg2}` Bitwise Binary And.
- `%|{arg1}{arg2}` Bitwise Binary Or.
- `%^{arg1}{arg2}` Bitwise Binary Exclusive Or.
- `%<<{arg1}{arg2}` Binary Shift Left.
- `%>>{arg1}{arg2}` Binary Shift Right.
- `%min{arg1}{arg2}` Compute the arithmetic minimum of the arguments.
- `%max{arg1}{arg2}` Compute the arithmetic maximum of the arguments.

Boolean Condition Macros

- `%=={arg1}{arg2}` Numeral Equality.
- `%!={arg1}{arg2}` Numeral Inequality.
- `%<{arg1}{arg2}` Numeral Less Than.
- `%<={arg1}{arg2}` Numeral Less or Equal.
- `%>{arg1}{arg2}` Numeral Greater Than.
- `%>={arg1}{arg2}` Numeral Greater or Equal.
- `%eq{arg1}{arg2}` String Equality.
- `%ne{arg1}{arg2}` String Inequality.
- `%lt{arg1}{arg2}` String Less Than.
- `%le{arg1}{arg2}` String Less or Equal.
- `%gt{arg1}{arg2}` String Greater Than.
- `%ge{arg1}{arg2}` String Greater or Equal.
- `%~{string}{regex}{opts}` or `%match{string}{regex}{opts}` Checks whether *string* matches the Perl regular expression *regex*. Modifiers can be given via *opts*.

Shortcut Evaluation Operators The following operators evaluate their arguments only when needed (like in C).

- `%&&{arg1}{arg2}` Logical And.
- `%and{arg1}{arg2}` Alias for `%&&{}`.
- `%||{arg1}{arg2}` Logical Or.
- `%or{arg1}{arg2}` Alias for `%||{}`.

Unary Operators

- `%!{arg}` Logical Not.
- `%not{arg}` Alias for `%!{}`.
- `%~{arg}` Bitwise Negation.

String Functions

- `%length{string}` Return the number of ASCII characters present in *string*.
- `%toupper{string}` Return all ASCII characters converted to uppercase.
- `%tolower{string}` Return all ASCII characters converted to lowercase.
- `%append{varname}{string}` Equivalent to `%let{varname}{%{varname}string}`.
- `%subst{string}{regex}{subst}{opts}` Perl regex substitution.
- `%sprintf{fmt}{arg1}{arg2}{argn}` Perl `sprintf()` operator. Details see Perl manual.
- `%human-number{unit}{delim}{unit-sep}{number1}{number2}...` Convert a number or a list of numbers into human-readable B, KiB, MiB, GiB, TiB, as given by *unit*. When *unit* is empty, a reasonable unit will be guessed automatically from the maximum of all given numbers. A single result string is produced, where multiple numbers are separated by *delim* when necessary. When *delim* is empty, the slash symbol / is used by default (the most obvious use case is result strings like “17/32 KiB”). The final unit text is separated from the previous number(s) by *unit-sep*. When *unit-sep* is empty, a single blank is used by default.
- `%human-seconds{number}` Convert the given number of seconds into `hh:mm:ss` format.

Complex Helper Macros

- `%progress{20}` Return a string containing a progress bar showing the values from `%summary-vector{}`. The default width is 20 characters plus two braces.
- `%progress{20}{minvalue}{midvalue}{maxvalue}` Instead of taking the values from `%summary-vector{}`, use the supplied values. `minvalue` and `midvalue` indicate two different intermediate points, while `maxvalue` will determine the 100% point.

Control Flow Macros

- `%if{expression}{then-part}` or `%if{expression}{then-part}{else-part}` Like in any other macro or programming language, this evaluates the `expression` once, not copying its outcome to the output. If the result is non-empty and is not a string denoting the number 0, the `then-part` is evaluated and copied to the output. Otherwise, the `else-part` is evaluated and copied, provided that one exists.
- `%unless{expression}{then-part}` or `%unless{expression}{then-part}{else-part}` Like `%if{}`, but the expression is logically negated. Essentially, this is a shorthand for `%if{!expression}{...}` or similar.
- `%elsif{expr1}{then1}{expr2}{then2}... or %elsif{expr1}{then1}{expr2}{then2}...{odd-else-p}` This is for simplification of boring if-else-if chains. The classical if-syntax (as shown above) has the drawback that inner if-parts need to be nested into outer else-parts, so rather deep nestings may occur when you are programming longer chains. This is an alternate syntax for avoidance of deep nesting. When giving an odd number of arguments, the last argument is taken as final else-part.
- `%elsunless...` Like `%elsif`, but *all* conditions are negated.
- `%while{expression}{body}` Evaluates the `expression` in a while loop, like in any other macro or programming language. The `body` is evaluated exactly as many times as the `expression` holds. Notice that endless loops can be only avoided by a calling a non-pure macro inspecting external state information, or by creating (and checking) another side effect somewhere, like assigning to a variable somewhere.
- `%until{expression}{body}` Like `%while{expression}{body}`, but negate the expression.
- `%for{expr1}{expr2}{expr3}{body}` As you will expect from the corresponding C, Perl, Java, or (add your favorite language) construct. Only the syntactic sugar is a little bit different.
- `%foreach{varname}{CSV-delimited-string}{delimiter}{body}` As you can expect from similar `foreach` constructs in other languages like Perl. Currently, the macro processor has no arrays, but can use comma-separated strings as a substitute.
- `%eval{count}{body}` Evaluates the `body` exactly as many times as indicated by the numeric argument `count`. This may be used to re-evaluate the output of other macros once again.
- `%protect{body}` Equivalent to `%eval{0}{body}`, which means that the body is not evaluated at all, but copied to the output verbatim¹⁴.
- `%eval-down{body}` Evaluates the `body` in a loop until the result does not change any more¹⁵.
- `%tmp{body}` Evaluates the `body` once in a temporary scope which is thrown away afterwards.

¹⁴TeX or L^AT_EX fans usually know what this is good for ;)

¹⁵Mathematicians knowing Banach's fixedpoint theorem will know what this is good for ;)

5. The Macro Processor

- `%call{macroname}{arg1}{arg2}{argn}` Like in many other macro languages, this evaluates the named macro in the a new scope. This means that any side effects produced by the called macro, such as variable assignments, will be reverted after the call, and therefore not influence the old scope. However notice that the arguments `arg1` to `argn` are evaluated in the *old* scope before the call actually happens (possibly producing side effects if they contain some), and their result is respectively assigned to `%{1}` until `%{n}` in the new scope, analogously to the Shell or to Perl. In addition, the new `%{0}` gets the *macroname*. Notice that the argument evaluation happens non-lazily in the old scope and therefore differs from other macro processors like T_EX.
- `%include{macroname}{arg1}{arg2}{argn}` Like `%call{}`, but evaluates the named macro in the *current* scope (similar to the `source` command of the bourne shell). This means that any side effects produced by the called macro, such as variable assignments, will *not* be reverted after the call. Even the `%{0}` until `%{n}` variables will continue to exist (and may lead to confusion if you aren't aware of that).
- `%callstack{}` Useful for debugging: show the current chain of macro invocations.

Time Handling Macros

- `%time{}` Return the current Lamport timestamp (see section 4.2), in units of seconds since the Unix epoch.
- `%sleep{seconds}` Pause the given number of seconds.
- `%timeout{seconds}` Like `%sleep{seconds}`, but abort the `marsadm` command after the total waiting time has exceeded the timeout given by the `--timeout=` parameter.

Misc Macros

- `%warn{text}` Show a WARNING:
- `%die{text}` Abort execution with an error message.

Experts Only - Risky The following macros are unstable and may change at any time without notice.

- `%get-msg{name}` Low-level access to system messages. You should not use this, since this is not extensible (you must know the name in advance).
- `%readlink{path}` Low-level access to symlinks. Don't misuse this for circumvention of the abstraction macros from the symlink tree!
- `%setlink{value}{path}` Low-level creation of symlinks. Don't misuse this for circumvention of the abstraction macros for the symlink tree!
- `%fetch-info{}` etc. Low-level access to internal symlink formats. Don't use this in scripts! Only for curious humans.
- `%is-almost-consistent{}` Whatever you guess what this could mean, don't use it, at least never in place of `%is-consistent{}` - it is risky to base decisions on this. Mostly for historical reasons.
- `%does{name}` Equivalent to `%is-name{}` (just more handy for computing the macro name). Use with care!

5.2.3. Predefined Variables

- `%{cmd}` The command argument of the invoked `marsadm` command.
- `%{res}` The resource name given to the `marsadm` command as a command line parameter (or, possibly expanded from `all`).

- `%{resdir}` The corresponding resource directory. The current version of MARS uses `/mars/resource-%{res}/`, but this may change in future. Normally, you should need this, since anything should be already abstracted for you. In case you *really* need low-level access to something, please prefer this variable over `%{mars}/resource-%{res}` because it is a bit more abstracted.
- `%{mars}` Currently the fixed string `/mars`. This may change in future, probably with the advent of MARS Full.
- `%{host}` The hostname of the local node.
- `%{ip}` The IP address of the local node.
- `%{timeout}` The value given by the `--timeout=` option, or the corresponding default value.
- `%{threshold}` The value given by the `--threshold=` option, or the corresponding default value.
- `%{window}` The value given by the `--window=` option, or the corresponding default value.
- `%{force}` The number of times the `--force` option has been given.
- `%{dry-run}` The number of times the `--dry-run` option has been given.
- `%{verbose}` The number of times the `--verbose` option has been given.
- `%{callstack}` Same as the `%callstack{}` macro. The latter gives you an opportunity for overriding, while the former is firmly built in.

5.3. Scripting HOWTO

Both the **asynchronous communication model** of MARS (cf section 4.2) including the Lamport clock, and the **state model** (cf section 3.5) is something you *definitely* should have in mind when you want to do some scripting. Here is some further concrete advice:

- Don't access anything on `/mars/` directly, except for debugging purposes. Use `marsadm`.
- Avoid running scripts in parallel, other than for inspection / monitoring purposes. When you give two `marsadm` commands in parallel (whether on the same host, or on different hosts belonging to the same cluster), it is very likely to produce a mess. `marsadm` has no internal locking. There is no cluster-wide locking at all. Unfortunately, some systems like Pacemaker are violating this in many cases (depending on their configuration). Best is if you have a dedicated / more or less centralized **control machine** which controls masses of your georedundant working servers. This reduces the risk of running interfering actions in parallel. Of course, you need backup machines for your control machines, and in different locations. Not obeying this advice can easily lead to problems such as complex races which are very difficult to solve in long-distance distributed systems, even in general (not limited to MARS).

- `marsadm wait-cluster` is your friend. Whenever your (near-)central script has to switch between different hosts A and B (of the same cluster), use it in the following way:
`ssh A "marsadm action1"; ssh B "marsadm wait-cluster; marsadm action2"`



Don't ignore this advice! Interference is almost *sure*! As a rule of thumb, precede almost any action command with some appropriate waiting command!

- Further friends are any `marsadm wait-*` commands, such as `wait-umount`.
- In some places, busy-wait loops might be needed, e.g. for waiting until a specific resource is `UpToDate` or matches some other condition. Examples of waiting conditions can be found under github.com/schoebel/test-suite in subdirectory `mars/modules/`, specifically `02_predicates.sh` or similar.

5. *The Macro Processor*




- In case of network problems, some command may hang (forever), if you don't set the `--timeout=` option. Don't forget to check the return state of any failed / timeouted commands, and to take appropriate measures!
- Test your scripts in failure scenarios!

6. The Sysadmin Interface (marsadm and /proc/sys/mars/)


In general, the term “after a while” means that other cluster nodes will take notice of your actions according to the “eventually consistent” propagation protocol described in sections 4.2 and 4.3. Please be aware that this “while” may last very long in case of network outages or bad firewall rules.

In the following tables, column “Cmp” means compatibility with DRBD. Please note that 100% exact compatibility is not possible, because of the asynchronous communication paradigm.

The following table documents common options which work with (almost) any command:




Option	Cmp	Description
<code>--dry-run</code>	no	<p>Run the command without actually creating symlinks or touching files or executing rsync. This option <i>should</i> be used first at any dangerous command, in order to check what would happen.</p> <p> Don't use in scripts! Only use by hand!</p> <p>This option does not change the waiting logic. Many commands are waiting until the desired effect has taken place. However, with <code>--dry-run</code> the desired effect will never happen, so the command may wait forever (or abort with a timeout).</p> <p>In addition, this option can lead to additional aborts of the commands due to unmet conditions, which cannot be met because the symlinks are not actually created / altered.</p> <p>Thus this option can give only a rough estimate of what would happen later!</p>
<code>--force</code>	almost	<p>Some preconditions are skipped, i.e. the command will / should work although some (more or less) vital preconditions are violated.</p> <p>Instead of giving <code>--force</code>, you may alternatively prefix your command with <code>force-</code></p> <p> THIS OPTION IS DANGEROUS!</p> <p>Use it only when you are absolutely sure that you know what you are doing!</p> <p>Use it only as a last resort if the same command without <code>--force</code> has failed <i>for no good reason!</i></p>
<code>--verbose</code>	no	Some (few) commands will become more speaky.
<code>--timeout=\$seconds</code>	no	<p>Some commands require response from either the local kernel module, or from other cluster nodes. In order to prevent infinite waiting in case of network outages or other problems, the command will fail after the given timeout has been reached.</p> <p>When <code>\$seconds</code> is -1, the command will wait forever.</p> <p>When <code>\$seconds</code> is 0, the command will not wait in case any precondition is not met, und abort without performing an action..</p> <p>The default timeout is 5s.</p>
<code>--window=\$seconds</code>	no	The time window for checking the aliveness of other nodes in the network. When no symlink updates have occurred during the last window, the node is considered dead. Default is 30s
<code>--threshold=\$size</code>	no	<p>The macros containing the substring <code>-threshold-</code> or <code>-almost-</code> are using this as a default value for approximation whether something has been approximately reached. Default is 10MiB.</p> <p>The <code>\$size</code> argument may be a number optionally followed by one the lowercase characters <code>k m g t p</code> for indicating kilo mega giga tera or peta bytes as multiples of 1000. When using the corresponding uppercase character, multiples of 1024 are formed instead.</p>
<code>--host=\$host</code>	no	<p>The command acts as if the command were executed on another host <code>\$host</code>. This option should not be used regularly, because the local information in the symlink tree may be outdated or even wrong. Additionally, some local information like remote sizes of physical devices (e.g. remote disks) is not present in the symlink tree at all, or is wrong (reflecting only the <i>local</i> state).</p> <p> THIS OPTION IS DANGEROUS!</p> <p>Use it only for final destruction of dead cluster nodes, see section 3.4.4.</p>
Option	Cmp	Description

6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)

Option	Cmp	Description
<code>--ip=\$ip</code>	no	By default, <i>marsadm</i> always uses the IP for <code>\$host</code> as stored in the symlink tree (directory <code>/mars/ips/</code>). When such an IP entry does not (yet) exist (e.g. <code>create-cluster</code> or <code>join-cluster</code>), all local network interfaces are automatically scanned for IPv4 addresses, and the first one is taken. This may lead to wrong decisions if you have multiple network interfaces. In order to override the automatic IP detection and to explicitly tell the IP address of your storage network, use this option.  Usually you will need this only at <code>{create,join}-cluster</code> .
<code>--verbose</code>	no	Some (few) commands will become more speaky.
Option	Cmp	Description

6.1. Cluster Operations

Command / Params	Cmp	Description
<code>create-cluster</code>	no	Precondition: the <code>/mars/</code> filesystem must be mounted and it must be empty (<code>mkfs.ext4</code> , see instructions in section 3.2.2). The kernel module must <i>not</i> be loaded. Postcondition: the initial symlink tree is created in <code>/mars/</code> . Additionally, the <code>/mars/uuid</code> symlink is created for later distribution in the cluster. It uniquely identifies the cluster in the world. This must be called exactly once at the initial primary. Hint: use the <code>--ip=</code> option if you have multiple interfaces.
<code>join-cluster</code> <code>\$host</code>	no	Precondition: the <code>/mars/</code> filesystem must be mounted and it must be empty (<code>mkfs.ext4</code> , see instructions in section 3.2.2). The kernel module must <i>not</i> be loaded. The cluster must have been already created at another node <code>\$host</code> . A working ssh connection to <code>\$host</code> as root must exist (without password). <code>rsync</code> must be installed at all cluster nodes. Postcondition: the initial symlink tree <code>/mars/</code> is replicated from the remote host <code>\$host</code> , and the local host has been added as another cluster member. This must be called exactly once at every initial secondary node. Hint: use the <code>--ip=</code> option if you have multiple interfaces.
Command / Params	Cmp	Description

Command / Params	Cmp	Description
<code>leave-cluster</code>	no	<p>Precondition: the <code>/mars/</code> filesystem must be mounted and it must contain a valid MARS symlink tree produced by the other <code>marsadm</code> commands. The local node must no longer be member of any resource (see <code>marsadm leave-resource</code>). The kernel module should be loaded and the network should be operating in order to also propagate the effect to the other nodes.</p> <p>Postcondition: the local node is removed from the replicated symlink tree <code>/mars/</code> such that other nodes will cease to communicate with it after a while. The converse is not true: the local node may continue passively fetching the symlink tree. In order to really stop all communication, the kernel module should be unloaded afterwards. The local <code>/mars/</code> filesystem may be manually destroyed after that (at least if you need to reuse it).</p> <p>In case of an eventual node loss (e.g. fire, water, ...) this command should be used on another node <code>\$helper</code> in order to finally remove <code>\$damaged</code> from the cluster via the command <code>marsadm leave-cluster --host=\$damaged --force</code>.</p> <p> In case you cannot use <code>leave-resource</code> for any reason, you may do the following: just destroy the <code>/mars/</code> filesystem on the host <code>\$deadhost</code> you want to remove (e.g. by <code>mkfs</code>), or take other measures to <i>ensure</i> that it cannot be accidentally re-used in any way (e.g. physical destruction of the underlying RAID, <code>lvremove</code>, etc). On all other hosts, do <code>rmmod mars</code>, then delete the symlink <code>/mars/ips/ip-\$deadhost</code> everywhere by hand, and finally <code>modprobe mars</code> again.</p> <p> Notice that the last <code>leave-resource</code> operation does not delete the cluster as such. It just creates an <i>empty</i> cluster which has no longer any members. In particular, the cluster ID <code>/mars/uuid</code> is <i>not</i> removed, deliberately^b.</p> <p> Before you can re-use <i>any</i> left-over <code>/mars/</code> filesystem for creating / joining a new / different cluster, you <i>must</i> obey the instructions in section 3.2.2 and use <code>mkfs.ext4</code> accordingly.</p> <p>^aReason: <code>leave-cluster</code> removes only its <i>own</i> IP address from <code>/mars/ips/</code>, but does not destroy the usual symmetry of the symlink tree by leaving the other IPs intact. Therefore, the local node will continue fetching updates from all nodes present in <code>/mars/ips/</code>. As an effect, the local node will <i>passively</i> mirror the symlinks of other cluster members, but not vice versa. There is no communication from the local node to the other ones, turning the local node into a whitness according to some terminology from Distributed Systems. This is a feature, not a bug. It could be used for post-mortem analysis, or for monitoring purposes. However, <i>deletions</i> of symlinks are not guaranteed to take place, so your whitness may <i>accumulate</i> thousands of old symlinks over a long time. If you want to eventually stop all communication to the local node, just run <code>rmmod</code>.</p> <p>^bThis is a feature, not a bug. The <code>uuid</code> is created once, but never altered anywhere. The only way to get rid of it is <i>external</i> deletion (not by <code>marsadm</code>) <i>together(!)</i> with all other contents of <code>/mars/</code>. This prevents you from accidentally merging half-dead remains which could have survived a disaster for any reason, such as snapshotting filesystems / VMs or whatever.</p>
<code>wait-cluster</code>	no	See section 6.3.3.
<code>create-uuid</code>	no	<p>Deprecated. Only for compatibility with old version <code>light0.1beta05</code> or earlier.</p> <p>Precondition: the <code>/mars/</code> filesystem must be mounted. A <code>uuid</code> (such as automatically created by recent versions of <code>marsadm create-cluster</code>) must not already exist; i.e. you have a very old and outdated symlink tree.</p> <p>Postcondition: the <code>/mars/uuid</code> symlink is created for later distribution in the cluster. It uniquely identifies the cluster in the world. This must be called at most once at the current primary.</p>
Command / Params	Cmp	Description

6.2. Resource Operations

Common precondition for all resource operations is that the `/mars/` filesystem is mounted, that it contains a valid MARS symlink tree produced by other `marsadm` commands (including

6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)

a unique `uuid`), that your current node is a valid member of the cluster, and that the kernel module is loaded. When communication is impossible due to network outages or bad firewall rules, most commands will succeed, but other cluster nodes may take a long time to notice your changes.




Instead of executing `marsadm` commands several times for each resource argument, you may give the special resource argument `all`. This work even when combined with `--force`, but be cautious when giving dangerous command combinations like `marsadm delete-resource --force all`.






Beware when combining this with `--host=somebody`. In some very rare cases, like final destruction of a whole datacenter after an earthquake, you might need a combination like `marsadm --host=defective delete-resource --force all`. Don't use such combinations if you don't need them *really*! You can easily shoot yourself in your head if you are not carefully operating such commands!

6.2.1. Resource Creation / Deletion / Modification

Command / Params	Cmp	Description
<code>create-resource</code> <code> \$res</code> <code> \$disk_dev</code> <code> [\$mars_name]</code> <code> [\$size]</code>	no	<p>Precondition: the resource argument <code>\$res</code> must not denote an already existing resource name in the cluster. The argument <code>\$disk_dev</code> must denote an absolute path to a usable local block device, its size must be greater zero. When the optional <code>\$mars_name</code> is given, that name must not already exist on the local node; when not given, <code>\$mars_name</code> defaults to <code>\$res</code>. When the optional <code>\$size</code> argument is given, it must be a number, optionally followed by a lowercase suffix <code>k</code>, <code>m</code>, <code>g</code>, <code>t</code>, or <code>p</code> (denoting size factors as multiples of 1000), or an uppercase suffix <code>K</code>, <code>M</code>, <code>G</code>, <code>T</code> or <code>P</code> (denoting size factors as multiples of 1024). The given size must not exceed the actual size of <code>\$disk_dev</code>. It will specify the future resource size as shown by <code>marsadm view-resource-size \$res</code>.</p> <p>Postcondition: the resource <code>\$res</code> is created, the initial role of the current node is primary. The corresponding symlink tree information is asynchronously distributed in the cluster (in the background). The device <code>/dev/mars/\$mars_name</code> should appear after a while.</p> <p>Notice: when <code>\$size</code> is strictly smaller than the size of <code>\$disk_dev</code>, you will unnecessarily waste some space..</p> <p>This must be called exactly once for any new resource.</p>
<code>join-resource</code> <code> \$res</code> <code> \$disk_dev</code> <code> [\$mars_name]</code>	no	<p>Precondition: the resource argument <code>\$res</code> must denote an already existing resource in the cluster (i.e. its symlink tree information must have been received). The resource must have a designated primary, and it must not be in emergency mode. There must not exist a split brain in the cluster. The local node must not be already member of that resource. The argument <code>\$disk_dev</code> must denote an absolute path to a usable (but currently unused) local block device, its size must be greater or equal to the logical size of the resource. When the optional <code>\$mars_name</code> is given, that name must not already exist on the local node; when not given, <code>\$mars_name</code> defaults to <code>\$res</code>.</p> <p>Postcondition: the current node becomes a member of resource <code>\$res</code>, the initial role is secondary. The initial full sync should start after a while.</p> <p>Notice: when the size of <code>\$disk_dev</code> is strictly greater than the size of the resource, you will unnecessarily waste some space..</p>
Command / Params	Cmp	Description

Command / Params	Cmp	Description
leave-resource \$res	no	<p>Precondition: the local node must be a member of the resource \$res; its current role must be secondary. Sync, fetch and replay must be paused (see commands pause-{sync,fetch,replay} or their abbreviation down). The disk must be detached (see commands detach or down). The kernel module should be loaded and the network should be operating in order to also propagate the effect to the other nodes.</p> <p>Postcondition: the local node is no longer a member of \$res.</p> <p>Notice: as a side effect for other nodes, their log-delete may now become possible, since the current node does no longer count as a candidate for logfile application. In addition, a split brain situation may be (partly) resolved by this.</p> <p> Please notice that this command <i>may</i> lead to (but does not guarantee) split-brain resolution.</p> <p> The contents of the disk is not changed by this command. Before issuing this command, check whether the disk appears to be locally consistent (see view-is-consistent)! After giving this command, any internal information indicating the consistency state will be gone, and you will no longer be able to guess consistency properties.</p> <p> When you are <i>sure</i> that the disk was consistent before (or is now by manually checking it), you may re-create a new resource out of it via create-resource.</p> <p>In case of an eventual node loss (e.g. fire, water, ...) this command may be used on another node \$helper in order to finally remove all the resources \$damaged from the cluster via the command marsadm leave-resource \$res --host=\$damaged --force.</p>
Command / Params	Cmp	Description

6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)








Command / Params	Cmp	Description
delete-resource \$res	no	<p>Precondition: the resource must be empty (i.e. all members must have left via <code>leave-resource</code>). This precondition is overridable by <code>--force</code>, increasing the danger to maximum! It is even possible to combine <code>--force</code> with an invalid resource argument and an invalid <code>--host=somebodyelse</code> argument in order to desperately try to destroy remains of incomplete or physically damaged hardware.</p> <p>Postcondition: all cluster members will somewhen be forcefully removed from \$res. In case of network interruptions, the forced removal may take place far in the future.</p> <p> THIS COMMAND IS <i>VERY</i> DANGEROUS!</p> <p>Use this only in desperate situations, and only manually. Don't call this from scripts. You are forcefully using a sledgehammer, even without <code>--force</code>! The danger is that the <i>true</i> state of other cluster nodes need not be known in case of network problems. Even when it were known, it could be compromised by byzantine failures.</p> <p>It is strongly advised to try this command with <code>--dry-run</code> first.</p> <p>When combined with <code>--force</code>, this command will definitely murder other cluster nodes, possibly after a long while, and even when they are operating in primary mode / having split brains / etc. However, there is no guarantee that other cluster nodes will be <i>really</i> dead – it is (theoretically) possible that they remain only <i>half dead</i>. For example, a half dead node may continue to write data to <i>/mars/</i> and thus lead to overflow somewhen.</p> <p> This command implies a forceful detach, possibly destroying consistency. It is similar in spirit to a STONITH. In particular, when a cluster node was operating in primary mode (<i>/dev/mars/mydata</i> being continuously in use), the forceful detach cannot be carried out until the device is completely unused. In the meantime, the current transaction logfile will be appended to, but the file <i>might</i> be already unlinked (orphan file filling up the disk). After the forceful detach, the underlying disk need not be consistent (although MARS does its best). Since this command deletes any symlinks which normally would indicate the consistency state, no guarantees about consistency can be given after this <i>in general!</i> Always check consistency by hand! When possible / as soon as possible, check the local state on the other nodes in order to <i>really</i> shutdown the resource everywhere (e.g. to <i>really</i> unuse the <i>/dev/mars/mydata</i> device, etc).</p> <p>After this command, you <i>should</i> rebuild the resource under a different name, in order to avoid any clashes caused by unexpected resurrection of “dead” or “half-dead” nodes (beware of shapshot / restores on virtual machines!!). MARS does its best to avoid problems even in case the new resource name should equal the old one, but there can be <i>no guarantee</i> in all possible failure scenarios / usage scenarios.</p> <p> When possible, prefer <code>leave-resource</code> over this!</p>
wait-resource \$res {is-,}{attach, primary, device}{-off,}	no	See section 6.3.3.
Command / Params	Cmp	Description

6.2.2. Operation of the Resource




Common preconditions are the preconditions from section 6.2, plus the respective resource **\$res** must exist, and the local node must be a member of it. With the single exception of `attach` itself, all other operations must be started in `attached` state.




When **\$res** has the special reserved value `all`, the following operations will work on all resources where the current node is a member (analogously to DRBD).

Command / Params	Cmp	Description
Command / Params	Cmp	Description








Command / Params	Cmp	Description
attach \$res	yes	<p>Precondition: the local disk belonging to \$res is not in use by anyone else. Its contents has not been altered in the meantime since the last detach.</p> <p> Mounting <i>read-only</i> is allowed during the detached phase.</p> <p> However, be careful! If you <i>accidentally</i> forget to give the right readonly-mount flags, if you use <code>fsck</code> in repair mode inbetween, or alter the disk content in any other way (beware of LVM snapshots / restores etc), you will almost certainly produce an unnoticed inconsistency (not reported by <code>view-is-consistent</code>)! MARS has <i>no chance</i> to notice suchalike!</p> <p>Postcondition: MARS uses the local disk and is able to work with it (e.g. replay logfiles on it).</p> <p>Note: the local disk is opened in exclusive read-write mode. This should protect against most common misuse, such as opening the disk in parallel to MARS.</p> <p> However, this does not necessarily protect against non-exclusive openers.</p>
detach \$res	yes	<p>Precondition: the local <code>/dev/mars/mydata</code> device (when present) is no longer opened by anybody.</p> <p>Postcondition: the local disk belonging to \$res is no longer in use.</p> <p> In contrast to DRBD, you need not explicitly pause syncing, fetching, or replaying <i>to</i> (as apposed to <i>from</i>) the local disk. These processes are automatically paused. As another contrast to DRBD, the respective processes will usually <i>automatically</i> resume after re-attach, as far as possible in the respective new situation. This will usually work even over <code>rmod</code> or reboot cycles, since the internal symlink tree will automatically persist all todo switches for you (c.f. section 3.5).</p> <p> Notice: only <i>local</i> transfer operations <i>to</i> the local disk are paused by a detach. When another node is remotely running a sync <i>from</i> your local disk, it will likely remain in use for remote reading. The reason is that the server part of MARS is operating purely passively, in order serve all remote requests as best as possible (similar to the original Unix philosophy). In order to really stop all accesses, do a <code>pause-sync</code> on all other resource member where a sync is currently running. You may also try <code>pause-sync-global</code>.</p> <p> WARNING! After this, and ather having paused any remote data access, you might use the underlying disk for your own purposes, such as test-mounting it in <i>readonly</i> mode. Don't modify its contents in any way! Not even by an <code>fsck</code>^a! Otherwise, you will have inconsistencies <i>guaranteed</i>. MARS has no way for knowing of any modifications to your disk when bypassing <code>/dev/mars/*</code>.</p> <p> In case you accidentally modified the underlying disk at the <i>primary</i> side, you may choose to resolve the inconsistencies by <code>marsadm invalide \$res</code> on <i>each</i> secondary.</p> <p>^aSome (but not all) <code>fsck</code> tools for some filesystems have options to start only a test repair / verify mode / dry run, without doing actual modifications to the data. Of course, these modes <i>can</i> be used. But be really sure! Double-check for the right options!</p>
pause-sync \$res	partly	Equivalent to <code>pause-sync-local</code> .
pause-sync-local \$res	partly	<p>Precondition: none additionally.</p> <p>Postcondition: any sync operation targeting the local disk (when not yet completed) is paused after a while (cf section 3.5). When successfully completed, this operation will remember the switch state forever and automatically become relevant if a sync is needed again (e.g. <code>invalidate</code> or <code>resize</code>).</p>
Command / Params	Cmp	Description





6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)

Command / Params	Cmp	Description
<code>pause-sync-global</code> \$res	partly	Like <code>*-local</code> , but operates on all members of the resource.
<code>resume-sync</code> \$res	partly	Equivalent to <code>resume-sync-local</code> .
<code>resume-sync-local</code> \$res	partly	Precondition: additionally, a primary must be designated, and it must not be in emergency mode. Postcondition: any sync operation targeting the local disk (when not yet completed) is resumed after a while. When completed, this operation will remember the switch state forever and become relevant if a sync is needed again (e.g. <code>invalidate</code> or <code>resize</code>).
<code>resume-sync-global</code> \$res	partly	Like <code>*-local</code> , but operates on all members of the resource.
<code>pause-fetch</code> \$res	partly	Equivalent to <code>pause-fetch-local</code> .
<code>pause-fetch-local</code> \$res	partly	Precondition: none additionally. The resource <i>should</i> be in secondary role. Otherwise the switch has <i>no immediate</i> effect, but will come (possibly unexpectedly) into effect whenever secondary role is entered later for whatever reason. Postcondition: any transfer of (parts of) transaction logfiles which are present at another primary host to the local <code>/mars/</code> storage are paused at their current stage.  This switch works independently from <code>{pause,resume}-replay</code> .
<code>pause-fetch-global</code> \$res	partly	Like <code>*-local</code> , but operates on all members of the resource.
<code>resume-fetch</code> \$res	partly	Equivalent to <code>resume-fetch-local</code> .
<code>resume-fetch-local</code> \$res	partly	Precondition: none additionally. The resource <i>should</i> be in secondary role. Otherwise the switch has <i>no immediate</i> effect, but will come (possibly unexpectedly) into effect whenever secondary role is entered later for whatever reason. Postcondition: any (parts of) transaction logfiles which are present at another primary host should be transferred to the local <code>/mars/</code> storage as far as not yet locally present.  This works independently from <code>{pause,resume}-replay</code> .
<code>resume-fetch-global</code> \$res	partly	Like <code>*-local</code> , but operates on all members of the resource.
<code>pause-replay</code> \$res	partly	Equivalent to <code>pause-replay-local</code> .
<code>pause-replay-local</code> \$res	partly	Precondition: none additionally. The resource <i>should</i> be in secondary role. Otherwise the switch has <i>no immediate</i> effect, but will come (possibly unexpectedly) into effect whenever secondary role is entered later for whatever reason. Postcondition: any local replay operations of transaction logfiles to the local disk are paused at their current stage.  This works independently from <code>{pause,resume}-fetch</code> resp. <code>{dis,}connect</code> .
<code>pause-replay-global</code> \$res	partly	Like <code>*-local</code> , but operates on all members of the resource.
<code>resume-replay</code> \$res	partly	Equivalent to <code>pause-replay-local</code> .
<code>resume-replay-local</code> \$res	partly	Precondition: must be in secondary role. Postcondition: any (parts of) locally existing transaction logfiles (whether replicated from other hosts or produced locally) are started for replay to the local disk, as far as they have not yet been applied.
<code>resume-replay-global</code> \$res	partly	Like <code>*-local</code> , but operates on all members of the resource.
Command / Params	Cmp	Description

Command / Params	Cmp	Description
connect \$res	partly	Equivalent to connect-local and to resume-fetch-local.  Note: although this sounds similar to DRBD's <code>drbdadm connect</code> , there are subtle differences. DRBD has exactly one connection per resource, which is associated with <i>pairs</i> of nodes. In contrast, MARS may create multiple connections per resource at runtime, and these are associated with the <i>target</i> host (not with <i>pairs</i> of hosts). As a consequence, the fetch may <i>potentially</i> occur from any other source host which happens to be reachable (although the current implementation prefers the current designated primary, but this may change in future). In addition, <code>marsadm disconnect</code> does not stop <i>all</i> communication. It only stops fetching logfiles. The symlink update running in background is <i>not</i> stopped, in order to always propagate as much metadata as possible in the cluster. In case of a later incident, chances are higher for a better knowledge of the <i>real</i> state of the cluster.
connect-local \$res	partly	Equivalent to resume-fetch-local.
connect-global \$res	partly	Equivalent to resume-fetch-global.
disconnect \$res	partly	Equivalent to disconnect-local and to pause-fetch-local.  See above note at <code>connect</code> .
disconnect-local \$res	partly	Equivalent to pause-fetch-local.
disconnect-global \$res	partly	Equivalent to pause-fetch-global.
up \$res	yes	Equivalent to <code>attach</code> followed by <code>resume-fetch</code> followed by <code>resume-replay</code> followed by <code>resume-sync</code> .
down \$res	yes	Equivalent to <code>pause-sync</code> followed by <code>pause-fetch</code> followed by <code>pause-replay</code> followed by <code>detach</code> .  Hint: consider to prefer plain <code>detach</code> over this, because <code>detach</code> will remember the last state of all switches, while <code>down</code> will <i>not</i> .
Command / Params	Cmp	Description

6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)

Command / Params	Cmp	Description
primary \$res	almost	<p>Precondition: sync must have finished at any resource member. All relevant transaction logfiles must be either already locally present, or be fetchable (see <code>resume-fetch</code> and <code>resume-replay</code>). When some logfile data is locally missing, there must be enough space on <i>/mars/</i> to fetch it. Any replay must not have been interrupted by a replay error (see macro <code>%replay-code{}</code> or diskstate <code>DefectiveLog</code>). The current designated primary must be reachable over network. When there is no designated primary (i.e. <code>marsadm secondary</code> had been executed before, which is explicitly <i>not recommended</i>), <i>all</i> other members of the resource must be reachable (since we have no memory who was the old primary before), and then they must also match the same preconditions. When another host is currently primary (whether designated or not), it must match the preconditions of <code>marsadm secondary</code> (that means, its local <i>/dev/mars/mydata</i> device must not be in use any more). A split brain must not already exist.</p> <p>Postcondition: <i>/dev/mars/\$dev_name</i> appears locally and is usable; the current host is in primary role.</p> <p>Switches the designated primary. There are two variants:</p> <p>1) Handover when <i>not</i> giving <code>--force</code>: when another host is currently primary, it is first asked to leave its primary role, and it is waited until it actually has become secondary. After that, the local host is asked to become primary. Before actually becoming primary, all relevant logfiles are transferred over the network and replayed, in order to avoid accidental creation of split brain as best as possible^a. Only after that, <i>/dev/mars/\$dev_name</i> will appear. When network transfers of the symlink tree are very slow (or currently impossible), this command may take a very long time.</p> <p>In case a split brain is already detected at the initial situation, the local host will refuse to switch the designated primary without <code>--force</code>.</p> <p> In case of $k > 2$ replicas: if you want to handover between host A and B while a sync is currently running at host C, you have the following options:</p> <ol style="list-style-type: none"> 1. wait until the sync has finished (see macro <code>sync-rest</code>, or <code>marsadm view</code> in general). 2. do a <code>leave-resource</code> on host C, and later <code>join-resource</code> after the handover completed successfully. <p>2) Forced switching: by giving <code>-force</code> while <code>pause-fetch</code> is active (but not <code>pause-replay</code>), most preconditions are ignored, and MARS does its best to actually become primary even if some logfiles are missing or incomplete or even defective.</p> <p> <code>primary --force</code> is a potentially harmful variant, because it will provoke a split brain in many cases, and therefore in turn will lead to data loss because one of your split brain versions must be discarded later in order to resolve the split brain (see section 3.4.3).</p> <p> Never call <code>primary --force</code> when <code>primary</code> without <code>--force</code> is sufficient! If <code>primary</code> without <code>--force</code> complains that the device is in use at the former primary side, take it seriously! Don't override with <code>--force</code>, but rather unmount^b the device at the other side!</p> <p> Only use <code>primary --force</code> when something is <i>already broken</i>, such as a network outage, or a node crash, etc. During ordinary operations (network OK, nodes OK), you should never need <code>primary --force</code>!</p> <p> If you unmount <i>/dev/mars/mydata</i> on the old primary A, and then wait until <code>marsadm view</code> (or another suitable macro) on the target host B shows that everything is <code>UpToDate</code>, you can prevent a split brain by yourself even when giving <code>primary --force</code> afterwards. However, checking / assuring this is <i>your</i> responsibility!</p> <p> <code>primary --force</code> switches the <i>designated</i> primary. In some extremely rare cases, when <i>multiple</i> faults have accumulated in a <i>weird</i> situation, it <i>might</i> be impossible becoming the / an actual primary. Typically you may be <i>already</i> in a split brain situation. This has not been observed for a long operations time on recent versions of MARS, but in general becoming primary via <code>--force</code> cannot be guaranteed always, although MARS does its best. In split brain situations, or if you ever encounter such a problem, you <i>must</i> resolve the split brain immediately after giving this command (see section 3.4.3).</p> <p> Hint in case of $k > 2$ replicas: <code>marsadm invalidate</code> cannot always resolve a split brain at other secondaries (which are neither the old</p>



Command / Params	Cmp	Description
secondary \$res	almost	<p>Precondition: the local <code>/dev/mars/\$dev_name</code> is no longer in use (e.g. unmounted). Postcondition: There exists no designated primary any more. During split brain and when the network is OK (again), all actual primaries (including the local host) will leave primary ASAP (i.e. when their <code>/dev/mars/mydata</code> is no longer in use). Any secondary will start following (old) logfiles (even from backlogs) by replaying transaction logs if it is <i>uniquely</i> possible (which is often violated during split brain). On any secondary, <code>/dev/mars/\$dev_name</code> will have disappeared.</p> <p> Notice: in difference to DRBD, you don't need this command during normal operation, including handover. Any resource member which is <i>not</i> designated as primary will <i>automatically</i> go into secondary role. For example, if you have $k = 4$ replicas, only <i>one of them</i> can be designated as a primary. When the network is OK, all other 3 nodes will know this fact, and they will <i>automatically</i> go into secondary mode, following the transaction logs from the (new) primary.</p> <p> Hint: avoid this command. It turns off <i>any</i> primary, globally^a. You cannot start a sync after that (e.g. <code>invalidate</code> or <code>join-resource</code> or <code>resume-sync</code>), because it is <i>not unique</i> wherefrom the data shall be fetched. In split brain situations (when the network is OK again), this may have further drawbacks. It is much better / easier to directly switch the designated primary from one node to another via the <code>primary</code> command. See also section 3.4.2.2.</p> <p> There is only one valid use case where you <i>really</i> need this command: before finally destroying a resource via the <i>last</i> <code>leave-resource</code> (or the dangerous <code>delete-resource</code>), you will need this before you can do that.</p> <p>^aA serious misconception among some people is when they believe that they can switch "a certain node to secondary". It is not possible to switch individual nodes to secondary, without affecting other nodes! The concept of "designated primary" is global throughout a resource!</p>
wait-umount \$res	no	See section 6.3.3.
log-purge-all \$res	no	<p>Precondition: none additionally. Postcondition: all locally known logfiles and version links are removed, whenever they are not / no longer reachable by any split brain version. Rationale: remove hindering split-brain / <code>leave-resource</code> leftovers. Use this only when split brain does not go away by means of <code>leave-resource</code> (which <i>could</i> happen in very weird scenarios such as MARS running on virtual machines doing a restore of their snapshots, or otherwise unexpected resurrection of dead or half-dead nodes).</p> <p> THIS IS POTENTIALLY DANGEROUS! This command <i>might</i> destroy some valuable logfiles / other information in case the local information is outdated or otherwise incorrect. MARS does its best for checking anything, but there is no guarantee. Hint: use <code>--dry-run</code> beforehand for checking!</p>
resize \$res [\$size]	almost	<p>Precondition: The local host must be primary. All disks in the cluster participating in <code>\$res</code> must be physically larger than the logical resource size (e.g. by use of <code>lvm</code>; can be checked by macros <code>%disk-size{}</code> and <code>%resource-size{}</code>). When the optional <code>\$size</code> argument is present, it must be smaller than the minimum of all physical sizes, but larger than the current logical size of the resource. Postcondition: the logical size of <code>/dev/mars/\$dev_name</code> will reflect the new size after a while.</p>
Command / Params	Cmp	Description

6.2.3. Logfile Operations

6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)

Command / Params	Cmp	Description
<code>log-rotate</code> <code>\$res</code>	no	Precondition: the local node <code>\$host</code> must be primary at <code>\$res</code> . Postcondition: after a while, a new transaction logfile <code>/mars/resource-\$res/log-\$new_nr-\$host</code> will be used instead of <code>/mars/resource-\$res/log-\$old_nr-\$host</code> where <code>\$new_nr = \$old_nr + 1</code> . Without <code>--force</code> , this will only carry out actions at the primary side since it makes no sense on secondaries. With <code>--force</code> , secondaries are <i>trying</i> to <i>remotely</i> trigger a log-rotate, but without any guarantee (likely even a split-brain may result instead, so use this only if you are <i>really</i> desperate).
<code>log-delete</code> <code>\$res</code>	no	Precondition: the local node must be a member of <code>\$res</code> . Postcondition: when there exists an old transaction logfile <code>/mars/resource-\$res/log-\$old_nr-\$some_host</code> where <code>\$old_nr</code> is the minimum existing number and that logfile is no longer referenced by any of the symlinks <code>/mars/resource-\$res/replay-*</code> , that logfile is marked for deletion in the whole cluster. When no such logfile exists, nothing will happen.
<code>log-delete-all</code> <code>\$res</code>	no	Like <code>log-delete</code> , but mark <i>all</i> currently unreferenced logfiles for deletion.
Command / Params	Cmp	Description

6.2.4. Consistency Operations

Command / Params	Cmp	Description
<code>invalidate</code> <code>\$res</code>	no	Precondition: the local node must be in secondary role at <code>\$res</code> . A <i>designated</i> primary must exist. When having $k > 2$ replicas, no split brain must exist (otherwise, or when <code>invalidate</code> does not work in case of $k = 2$, use the <code>leave-resource ; join-resource</code> method described in section 3.4.3). Postcondition: the local disk is marked as inconsistent, and a fast fullsync from the designated primary will start after a while. Notice that <code>marsadm {pause,resume}-sync</code> will influence whether the sync really starts. When the fullsync has finished successfully, the local node will be consistent again.
<code>fake-sync</code> <code>\$res</code>	no	Precondition: the local node must be in secondary role at <code>\$res</code> . Postcondition: when a fullsync is running, it will stop after a while, and the local node will be <i>marked</i> as consistent as if it were consistent again.  ONLY USE THIS IF YOU REALLY KNOW WHAT YOU ARE DOING! See the WARNING in section 3.3 Use this only <i>before</i> creating a fresh filesystem inside <code>/dev/mars/\$res</code> .
<code>set-replay</code>	no	 ONLY FOR ADVANCED HACKERS WHO KNOW WHAT THEY ARE DOING! This command is deliberately not documented. You need the competence level RTFS (“read the fucking sources”).
Command / Params	Cmp	Description

6.3. Further Operations

6.3.1. Inspection Commands

Command / Params	Cmp	Description
<code>view-macroname</code> <code>\$res</code>	no	Display the output of a macro evaluation. See section 3.6 for a thorough description.
<code>view</code> <code>\$res</code>	no	Equivalent to <code>view-default</code> .
<code>role</code> <code>\$res</code>	no	Deprecated. Use <code>view-role</code> instead.
<code>state</code> <code>\$res</code>	no	Deprecated. Use <code>view-state</code> instead.
<code>cstate</code> <code>\$res</code>	no	Deprecated. Use <code>view-cstate</code> instead.
Command / Params	Cmp	Description

Command / Params	Cmp	Description
<code>dstate</code> \$res	no	Deprecated. Use <code>view-dstate</code> instead.
<code>status</code> \$res	no	Deprecated. Use <code>view-status</code> instead.
<code>show-state</code> \$res	no	Deprecated. Don't use it. Use <code>view-state</code> instead, or other macros.
<code>show-info</code> \$res	no	Deprecated. Don't use it. Use <code>view-info</code> instead, or other macros.
<code>show</code> \$res	no	Deprecated. Don't use it. Use or implement some macros instead.
<code>show-errors</code> \$res	no	Deprecated. Use <code>view-the-err-msg</code> or <code>view-resource-err</code> similar macros.
<code>cat</code> \$file	no	Write the file content to stdout, but replace all occurrences of numeric timestamps converted to a human-readable format. This is most useful for inspection of status and log files, e.g. <code>marsadm cat /mars/5.total.log</code>
Command / Params	Cmp	Description

6.3.2. Setting Parameters

6.3.2.1. Per-Resource Parameters

Command / Params	Cmp	Description
<code>set-emergency-limit</code> \$res <i>n</i>	no	The argument <i>n</i> must be percentage between 0 and 100 %. When the remaining store space in <code>/mars/</code> undershoots the given percentage, the resource will go <i>earlier</i> into emergency mode than by the global computation described in section 4.4. 0 means unlimited.
<code>get-emergency-limit</code> \$res	no	Inquiry of the preceding value.
Command / Params	Cmp	Description


6.3.2.2. Global Parameters

Command / Params	Cmp	Description
<code>set-sync-limit-value</code> <i>n</i>	no	Limit the concurrency of sync operations to some maximum number. 0 means unlimited.
<code>get-sync-limit-value</code>	no	Inquiry of the preceding value.
<code>set-sync-pref-list</code> res1,res2,resn	no	Set the order of preferences for syncing. The argument must be comma-separated list of resource names.
<code>get-sync-pref-list</code>	no	Inquiry of the preceding value.
<code>set-connect-pref-list</code> host1,host2,hostn	no	Set the order of preferences for connections when there are more than 2 hosts participating in a cluster. The argument must be comma-separated list of node names.
<code>get-connect-pref-list</code>	no	Inquiry of the preceding value.
Command / Params	Cmp	Description

6.3.3. Waiting

Command / Params	Cmp	Description
Command / Params	Cmp	Description

6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)

Command / Params	Cmp	Description
<code>wait-cluster</code>	no	Precondition: the <code>/mars/</code> filesystem must be mounted and it must contain a valid MARS symlink tree produced by the other <code>marsadm</code> commands. The kernel module must be loaded. Postcondition: none. Wait until <i>all</i> nodes in the cluster have sent a message, or until timeout. The default timeout is 30 s (exceptionally) and  Be may be changed by <code>--timeout=\$seconds</code>
<code>wait-resource</code> <code>\$res</code> <code>{is-,}{attach,</code> <code>primary,</code> <code>device}{-off,}</code>	no	Precondition: the local node must be a member of the resource <code>\$res</code> . Postcondition: none. Wait until the local node reaches a specified condition on <code>\$res</code> , or until timeout. The default timeout of 60 s may be changed by <code>--timeout=\$seconds</code> . The last argument denotes the condition. The condition is inverted if suffixed by <code>-off</code> . When preceded by <code>is-</code> (which is the most useful case), it is checked whether the condition is actually reached. When the <code>is-</code> prefix is left off, the check is whether another <code>marsadm</code> command has been already given which <i>tries</i> to achieve the intended result (typically, you may use this after the <code>is-</code> variant has failed).
<code>wait-connect</code> <code>\$res</code>	almost	This is an alias for <code>wait-cluster</code> waiting until only those nodes are reachable which belong to <code>\$res</code> (instead of waiting for the <i>full</i> cluster).
<code>wait-umount</code> <code>\$res</code>	no	Precondition: none additionally. Postcondition: the local <code>/dev/mars/\$dev_name</code> is no longer in use (e.g. unmounted).
Command / Params	Cmp	Description

6.3.4. Low-Level Expert Commands

These commands are for experts and advanced sysadmins only. The interface is not stable, i.e. the meaning may change at any time. Use at your own risk!

Command / Params	Cmp	Description
<code>set-link</code>	no	RTFS.
<code>get-link</code>	no	RTFS.
<code>delete-file</code>	no	RTFS.
Command / Params	Cmp	Description

The following commands are for manual setup / repair of cluster membership. Only to be used by experts who know what they are doing! In general, cluster-wide operations on IP addresses may need to be repeated at all hosts in the cluster iff the communication is not (yet) possible and/or not (yet) actually working (e.g. firewalling problems etc).

Command / Params	Cmp	Description
<code>lowlevel-ls-host-ips</code>	no	List all configured cluster members together with their currently configured IP addresses, as known <i>locally</i> .
<code>lowlevel-set-host-ip</code> <code>\$hostname</code> <code>\$ip</code>	no	Change the assignment of IP addresses <i>locally</i> . May be used when hosts are moved to different network locations, or when different network interfaces are to be used for replication (e.g. dedicated replication IPs). Notice that the names of hosts must not change at all, only their IP addresses may be changed. Check active connections with <code>netstat</code> & friends. Updates may need some time to proceed (socket timeouts etc).
<code>lowlevel-delete-host</code> <code>\$hostname</code>	no	Remove a host from the cluster membership <i>locally</i> , together with its IP address assignment. This does not remove any further information. In particular, resource memberships are untouched.
Command / Params	Cmp	Description

6.3.5. Senseless Commands (from DRBD)

Command / Params	Cmp	Description
<code>syncer</code>	no	
<code>new-current-uuid</code>	no	
<code>create-md</code>	no	
<code>dump-md</code>	no	
<code>dump</code>	no	
<code>get-gi</code>	no	
<code>show-gi</code>	no	
<code>outdate</code>	no	
<code>adjust</code>	yes	Implemented as NOP (not necessary with MARS).
<code>hidden-commands</code>	no	
Command / Params	Cmp	Description

6.3.6. Forbidden Commands (from DRBD)

These commands are not implemented because they would be dangerous in MARS context:

Command / Params	Cmp	Description
<code>invalidate-remote</code>	no	This would be too dangerous in case you have multiple secondaries. A similar effect can be achieved with the <code>--host=</code> option.
<code>verify</code>	no	This would cause unintended side effects due to races between log-file transfer / application and block-wise comparison of the underlying disks. However, <code>marsadm join-resource</code> or <code>invalidate</code> will do the same as DRBD <code>verify</code> followed by DRBD <code>resync</code> , i.e. this will automatically correct any found errors;. Note that the fast-fullsync algorithm of MARS will minimize network traffic.
Command / Params	Cmp	Description

6.4. The `/proc/sys/mars/` and other Expert Tweaks

In general, you shouldn't need to deal with any tweaks in `/proc/sys/mars/` because everything should already default to reasonable predefined values. This interface allows access to some internal kernel variables of the `mars.ko` kernel module at runtime. Thus it is *not* a stable interface. It is not only specific for MARS, but may also change between releases without notice.

This section describes only those tweaks intended for sysadmins, not those for developers / very deep internals.

6.4.1. Syslogging

All internal messages produced by the kernel module belong to one of the following classes:

- 0 debug messages
- 1 info messages
- 2 warnings
- 3 error messages
- 4 fatal error messages
- 5 any message (summary of 0 to 4)

6.4.1.1. Logging to Files

These classes are used to produce status files `$class.*.status` in the `/mars/` and/or in the `/mars/resource-mydata/` directory / directories.

6. The Sysadmin Interface (*marsadm* and */proc/sys/mars/*)

When you create a file `$class.*.log` in parallel to any `$class.*.status`, the `*.log` file will be appended forever with the same messages as in `*.status`. The difference is that `*.status` is regenerated anew from an empty starting point, while `*.log` can (potentially) increase indefinitely unless you remove it, or rename it to something else.



Beware, any permanently present `*.log` file can easily fill up your `/mars/` partition until the problems described in section 4.4 will appear. Use `*.log` only for a **limited time**, and **only for debugging!**

6.4.1.2. Logging to Syslog

The classes also play a role in the following `/proc/sys/mars/` tweaks:

`syslog_min_class` (rw) The *minimum* class number for *permanent* syslogging. By default, this is set to -1 in order to switch off permanent logging completely. Permanent logging can easily flood your syslog with such huge amounts of messages (in particular when `class=0`), that your system as a whole may become unusable (because vital kernel threads may be blocked too long or too often by the userspace syslog daemon). Instead, please use the flood-protected syslogging described below!

`syslog_max_class` (rw) The *maximum* class number for *permanent* syslogging. Please use the flood-protected version instead.

`syslog_flood_class` (rw) The minimum class of flood-protected syslogging. The maximum class is always 4.

`syslog_flood_limit` (rw) The maximum number of messages after which the flood protection will start. This is a hard limit for the the number of messages written to the syslog.

`syslog_flood_recovery_s` (rw) The number of seconds after which the internal flood counter is reset (after flood protection state has been reached). When no new messages appear after this time, the flood protection will start over at count 0.



The rationale behind flood protected syslogging: sysadmins are usually only interested in the point in time where some problems / incidents / etc have *started*. They are usually not interested in capturing *each* and *every* single error message (in particular when they are flooding the system logs).



If you *really* need complete error information, use the `*.log` files described above, compress them and save them to somewhere else *regularly* by a cron job. This bears much less overhead than filtering via the syslog daemon, or even remote syslogging in real time which will almost surely screw up your system in case of network problems co-inciding with flood messages, such as caused in turn by those problems. Don't rely on real-time concepts, just do it the old-fashioned batch job way.

6.4.1.3. Tuning Verbosity of Logging

`show_debug_messages` Boolean switch, 0 or 1. Mostly useful only for developers. This can easily flood your logs if our are not careful.

`show_log_messages` Boolean switch, 0 or 1.

`show_connections` Boolean switch, 0 or 1. Show detailed internal statistics on sockets.

`show_statistics_local` / `show_statistics_global` Only useful for kernel developers. Shows some internal information on internal brick instances, memory usage, etc.

6.4.2. Tuning the Sync

`sync_flip_interval_sec` (rw) The sync process must not run in parallel to logfile replay, in order to easily guarantee consistency of your disk. If logfile replay would be paused for the full duration of very large or long-lasting syncs (which could take some days over very slow networks), your `/mars/` filesystem could overflow because no replay would be possible in the meantime. Therefore, MARS regularly flips between actually syncing and actually replaying, if both is enabled. You can set the time interval for flipping here.

`sync_limit` (rw) When > 0 , this limits the maximum number of sync processes actually running parallel. This is useful if you have a large number of resources, and you don't want to overload the network with sync processes.

`sync_nr` (ro) Passive indicator for the number of sync processes currently running.

`sync_want` (ro) Passive indicator for the number of sync processes which *demand* running.

7. Tips and Tricks


7.1. Avoiding Inappropriate Clustermanager Types for Medium and Long-Distance Replication

This section addresses some wide-spread misconceptions. Its main target audience is developers, but sysadmins will profit from **detailed explanations of problems and pitfalls**. When the problems described in this section are solved somewhen in future, this section will be shortened and some relevant parts moved to the appendix.

Doing **High Availability (HA)** wrong at *concept level* may easily get you into trouble, and may cost you several millions of € or \$ in larger installations, or even knock you out of business when disasters are badly dealt with at higher levels such as clustermanagers.

7.1.1. General Cluster Models

The most commonly known cluster model is called **shared-disk**, and typically controlled by clustermanagers like **PaceMaker**:

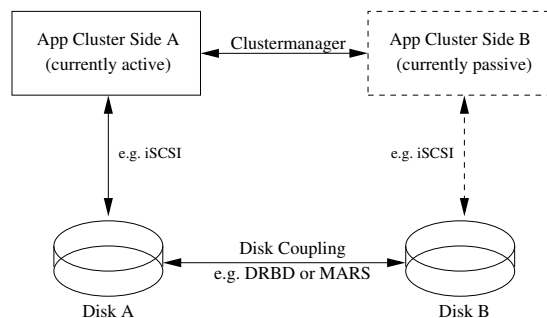


The most important property of shared-disk is that there exists only a single disk instance. Nowadays, this disk often has some *internal* redundancy such as RAID. At *system* architecture layer / network level, there exists no redundant disk at all. Only the application cluster is built redundant.



It should be immediately clear that shared-disk clusters are only suitable for short-distance operations in the same datacenter. Although running one of the data access lines over short distances between very near-by datacenters (e.g. 1 km) would be theoretically possible, there would be no sufficient protection against failure of a whole datacenter.

Both DRBD and MARS belong to a different architectural model called **shared-nothing**:



The characteristic feature of a shared-nothing model is (additional) **redundancy at network level**.



Shared-nothing “clusters¹” could theoretically be built for *any* distances, from short to medium to long distances. However, concrete technologies of disk coupling such as synchronous operation may pose practical limits on the distances (see chapter 2).

In general, clustermanagers must fit to the model. Some clustermanager can be configured to fit to multiple models. If so, this must be done properly, or you may get into serious trouble.

Some people don’t know, or they don’t believe, that different architectural models like shared-disk or shared-nothing will *require* an *appropriate* type of clustermanager and/or a different configuration. Failing to do so, by selection of an inappropriate clustermanager type and/or an inappropriate configuration may be hazardous.



Selection of the right model alone is not sufficient. Some, if not many, clustermanagers have not been designed for long distances. As explained in section 7.1.5, long distances have further **hard requirements**. Disregarding them may be also hazardous!

7.1.2. Handover / Failover Reasons and Scenarios

From a sysadmin perspective, there exist a number of different **reasons** why the application workload must be switched from the currently active side A to the currently passive side B:

1. Some **defect** has occurred at cluster side A or at some corresponding part of the network.
2. Some **maintenance** has to be done at side A which would cause a longer downtime (e.g. security kernel update or replacement of core network equipment or maintenance of UPS or of the BBU cache etc - hardware isn’t 24/7/365 in practice, although some vendors *claim* it - it is either not really true, or it becomes *extremely* expensive).

Both reasons are valid and must be automatically handled in larger installations. In order to deal with all of these reasons, the following basic mechanisms can be used in either model:

1. **Failover** (triggered either manually or automatically)
2. **Handover** (triggered manually²)

It is important to not confuse handover with failover at concept level. Not only the reasons / preconditions are very different, but also the *requirements*. Example: precondition for handover is that *both* cluster sides are healthy, while precondition for failover is that *some relevant(!)* failure has been *detected* somewhere (whether this is *really* true is another matter). Typically, failover must be able to run in masses, while planned handover often has lower scaling requirements.

Not all existing clustermanagers are dealing with all of these cases (or their variants) equally well, and some are not even dealing with some of these cases / variants *at all*.

Some clustermanagers cannot easily express the concept of “automatic triggering” versus “manual triggering” of an action. There exists simply no cluster-global switch which selects either “manual mode” or “automatic mode” (except when you start to hack the code and/or write new plugins; then you might notice that there is almost no architectural layering / sufficient separation between mechanism and strategy). Being forced to permanently use an automatic mode for several hundreds or even thousands of clusters is not only boring, but bears a considerable risk when automatics do a wrong decision at hundreds of instances in parallel.

7.1.3. Granularity and Layering Hierarchy for Long Distances

Many existing clustermanager solutions are dealing with a single cluster instance, as the term “clustermanager” suggests. However, when running several hundreds or thousands of cluster instances, you likely will not want to manage each of them individually. In addition, failover


¹Notice that the term “cluster computing” usually refers to short-distance only. Long-distance coupling should be called “grid computing” in preference. As known from the scientific literature, grid computing requires different concepts and methods in general. Only for the sake of simplicity, we use “cluster” and “grid” interchangeably.

²Automatic triggering could be feasible for prophylactic treatments.

7. Tips and Tricks

should *not only* be *triggered* (not to be confused with *executed*) individually at cluster level, but likely *also* at a higher granularity such as a room, or a whole datacenter. Otherwise, some chaos is likely to happen.

Here is what you probably will **need**, possibly in difference to what you may find on the market (whether OpenSource or not). For simplicity, the following diagram shows only two levels of granularity, but can be easily extended to multiple layers of granularity, or to some concept of various *subsets of clusters*:



Notice that many existing clustermanager solutions are not addressing the datacenter granularity at all. Typically, they use concepts like **quorums** for determining failures *at cluster level* solely, and then immediately executing failover of the cluster, sometimes without clean architectural distinction between trigger and execution (similar to the “separation of concerns” between **mechanism** and **strategy** in Operating Systems). Sometimes there is even no internal software layering / modularization according to this separation of concerns at all.



When there is no distinction between different levels of granularity, you are hopelessly bound to a non-extensible and thus non-adaptable system when you need to operate masses of clusters.



A lacking distinction between automatic mode and manual mode, and/or lack of corresponding **architectural software layers** is not only a blatant ignorance of well-established best practices of **software engineering**, but will bind you even more firmly to an inflexible system.



Terminology: for practical reasons, we use the general term “clustermanager” also for speaking about layers dealing with higher granularity, such as datacenter layers, and also for long-distance replication scenarios, although some terminology from grid computing would be more appropriate in a scientific background.

Please consider the following: when it comes to long-distance HA, the above layering architecture is also motivated by vastly different numbers of instances for each layer. Ideally, the topmost automatics layer should be able to overview several datacenters in parallel, in order to cope with (almost) global network problems such as network partitions. Additionally, it should also detect single cluster failures, or intermediate problems like “rack failure” or “room failure”, as well as various types of (partial / intermediate) (replication) network failures. Incompatible decisions at each of the different granularities would be a no-go in practice. Somewhere and somehow, you need one single³ top-most *logical* problem detection / ranking instance, which should be *internally distributed* of course, typically using some **distributed consensus protocol**; but in difference to many published distributed consensus algorithms it should be able to work with multiple granularities at the same time.

7.1.4. Methods and their Appropriateness

7.1.4.1. Failover Methods

Failover methods are only needed in case of an incident. They should not be used for regular handover.

³If you have *logical pairs of datacenters* which are firmly bound together, you could also have several topmost automatics instances, e.g. for each *pair* of datacenters. However, that would be very **inflexible**, because then you cannot easily mix locations or migrate your servers between datacenters. Using $k > 2$ replicas with MARS would also become a nightmare. In your own interest, please don't create any concepts where masses of hardware are firmly bound to fixed constants at some software layers.

STONITH-like Methods STONITH = Shoot The Other Node In The Head

These methods are widely known, although they have several serious drawbacks. Some people even believe that *any* clustermanager must *always* have some STONITH-like functionality. This is wrong. There *exist* alternatives, as shown in the next paragraph.

The most obvious drawback is that STONITH will always create a **damage**, by definition.

Example: a typical contemporary STONITH implementation uses IPMI for automatically powering off your servers, or at least pushes the (virtual) reset button. This will *always* create a certain type of damage: the affected systems will definitely not be available, at least for some time until they have (manually) rebooted.

This is a conceptual contradiction: the reason for starting failover is that you want to restore availability as soon as possible, but in order to do so you will first *destroy* the availability of a particular *component*. This may be counter-productive.

Example: when your hot standby node B does not work as expected, or if it works even *worse* than A before, you will lose some time until you *can* become operational again at the old side A.

Here is an example method for handling a failure scenario. The old active side A is assumed to be no longer healthy anymore. The method uses a sequential state transition chain with a STONITH-like step:

Phase1 Check whether the hot standby B is currently usable. If this is violated (which may happen during certain types of disasters), abort the failover for any affected resources.

Phase2 Try to shutdown the damaged side A (in the *hope* that there is no *serious* damage).

Phase3 In case phase2 did not work during a grace period / after a timeout, assume that A is badly damaged and therefore STONITH it.

Phase4 Start the application at the hot standby B.

Notice: any cleanup actions, such as **repair** of defective hard- or software etc, are outside the scope of failover processes. Typically, they are executed much later when restoring redundancy.

Also notice: this method is a *heavily* distributed one, in the sense that sequential actions are alternated multiple times on different hosts. This is known to be cumbersome in distributed systems, in particular in presence of network problems.

Phase4 in more detail for DRBD, augmented with some pseudo code for application control:

1. at side B: `drbdadm disconnect all`
2. at side B: `drbdadm primary --force all`
3. at side B: `applicationmanager start all`

The same phase4 using MARS:

1. at side B: `marsadm pause-fetch all`
2. at side B: `marsadm primary --force all`
3. at side B: `applicationmanager start all`

This sequential 4-phase method is far from optimal, for the following reasons:

- The method tries to handle both failover and handover scenarios with one single sequential recipe. In case of a true failover scenario where it is *already known for sure* that side A is badly damaged, this method will unnecessarily waste time for phase 2. This could be fixed by introduction of a conceptual distinction between handover and failover, but it would not fix the following problems.
- Before phase4 is started (which will re-establish the service from a user's perspective), a lot of time is wasted by *both* phases 2 *and* 3. Even if phase 2 would be skipped, phase 3 would unnecessarily cost some time. In the next paragraph, an alternative method is explained which eliminates any unnecessary waiting time at all.
- The above method is adapted to the shared-disk model. It does not take advantage of the shared-nothing model, where further possibilities for better solutions exist.

7. Tips and Tricks

- In case of long-distance network partitions and/or sysadmin / system management sub-network outages, you may not even be able to (remotely) start STONITH at all. Thus the above method misses an important failure scenario.

Some people seem to have a *binary* view at the healthiness of a system: in their view, a system is either operational, or it is damaged. This kind of view is ignoring the fact that some systems may be half-alive, showing only *minor* problems, or occurring only from time to time.

It is obvious that damaging a healthy system is a bad idea by itself. Even *generally* damaging a half-alive system in order to “fix” problems is not generally a good idea, because it may increase the damage when you don’t know the *real* reason⁴.

Even worse: in a distributed system⁵ you sometimes *cannot(!)* know whether a system is healthy, or to what degree it is healthy. Typical STONITH methods as used in some contemporary cluster managers are **assuming a worst case**, even if that worst case is currently not for real.

Therefore, avoid the following **fundamental flaws** in failover concepts and healthiness models, which apply to implementors / configurators of cluster managers:

- Don’t mix up knowledge with conclusions about a (sub)system, and also don’t mix this up with the real state of that (sub)system. In reality, you don’t have any knowledge about a complex distributed system. You only may have *some* knowledge about *some* parts of the system, but you cannot “see” a complex distributed system as a whole. What you think is your knowledge, isn’t knowledge in reality: in many cases, it is *conclusion*, not knowledge. Don’t mix this up!
- Some systems are more complex than your model of it. Don’t neglect important parts (such as networks, routers, switches, cables, plugs) which may lead you to wrong conclusions!
- Don’t restrict your mind to boolean models of healthiness. Doing so can easily create unnecessary damage by construction, and even at concept level. You should know from software engineering that defects in concepts or models are much more serious than simple bugs in implementations. Choosing the wrong model cannot be fixed as easily as a typical bug or a typo.
- Try to deduce the state of a system as **reliably** as possible. If you don’t know something for sure, don’t generally assume that it has gone wrong. Don’t confuse missing knowledge with the conclusion that something is bad. Boolean algebra restricts your mind to either “good” or “bad”. Use at least **tri-state algebra** which has a means for expressing “**unknown**”. Even better: attach a probability to anything you (believe to) know. Errare humanum est: nothing is absolutely sure.
- Oversimplification: don’t report an “unknown” or even a “broken” state for a complex system whenever a smaller subsystem exists for which you have some knowledge (or you can conclude something about it with reasonable evidence). Otherwise, your users / sysadmins may draw wrong conclusions, and assume that the whole system is broken, while in reality only some minor part has some minor problem. Users could then likely make wrong decisions, which may then easily lead to bigger damages.
- Murphy’s law: **never assume that something can’t go wrong!** Doing so is a blatant misconception at topmost level: the *purpose* of a cluster manager is creating High Availability (HA) out of more or less “unreliable” components. It is the damn duty of both a cluster manager and its configurator to try to compensate *any* failures, *regardless of their probability*⁶, as best as possible.

⁴Example, occurring in masses: an incorrectly installed bootloader, or a wrong BIOS boot priority order which unexpectedly lead to hangs or infinite reboot cycles once the DHCP or BOOTP servers are not longer available / reachable.

⁵Notice: the STONITH concept is more or less associated with short-distance scenarios where **crossover cables** or similare equipment are used. The assumption is that crossover cables can’t go defective, or at least it would be an extremely unlikely scenario. For long-distance replication, this assumption is simply not true.

⁶Never claim that something has only low probability (and therefore it were not relevant). In the HA area, you simply **cannot know** that, because you typically have *sporadic* incidents. In extreme cases, the *purpose* of your HA solution is protection against 1 failure per 10 years. You simply don’t have the time to wait for creating an incident statistics about that!

7.1. Avoiding Inappropriate Clustermanager Types for Medium and Long-Distance Replication

- Never confuse **probability** with **expectancy value!** If you don't know the mathematical term "expectancy value", or if you don't know what this means *in practice*, don't take responsibility for millions of € or \$.
- When operating masses of hard- and software: never assume that a particular failure can occur only at a low number of instances. There are **unknown(!) systematic errors** which may pop up at the wrong time and in huge masses when you don't expect them.
- Multiple layers of fallback: *any* action can fail. Be prepared to have a plan B, and even a plan C, and even better a plan D, wherever possible.
- Never increase any damage anywhere, unnecessarily! Always try to *mimize* any damage! It can be mathematically proven that in deterministic probabilistic systems having finite state, increases of a damage level *at the wrong place* will *introduce* an *additional risk* of getting into an **endless loop**. This is also true for nondeterministic systems, as known from formal language theory⁷.
- Use the **best effort principle**. You should be aware of the following fact: in general, it is impossible to create an *absolutely reliable system* out of unreliable components. You can *lower* the risk of failures to any $\epsilon > 0$ by investing a lot of resources and of money, but whatever you do: $\epsilon = 0$ is impossible. Therefore, be careful with boolean algebra. Prefer approximation methods / optimizing methods instead. Always do *your* best, instead of trying to reach a *global* optimum which likely does not exist at all (because the ϵ can only *converge* to an optimum, but will never actually reach it). The best effort principle means the following: if you discover a method for improving your operating state by reduction of a (potential) damage in a reasonable time and with reasonable effort, then **simply do it**. Don't argue that a particular step is no 100% solution for all of your problems. *Any improvement* is valuable. **Don't miss any valuable step** having reasonable costs with respect to your budget. Missing valuable measures which have low costs are certainly a violation of the best effort principle, because you are not doing *your* best. Keep that in mind.

If you have *understood* this (e.g. deeply think at least one day about it), you will no longer advocate STONITH methods *in general*, when there are alternatives. STONITH methods are only valuable when you *know in advance* that the final outcome (after reboot) will most likely be better, and that waiting for reboot will most likely *pay off*. In general, this condition is *not true* if you have a healthy hot standby system. This should be easy to see. But there exist well-known clustermanager solutions / configurations blatantly ignoring⁸ this. Only when the former standby system does not work as expected (this means that *all* of your redundant systems are not healthy enough for your application), *only then*⁹ STONITH is inevitable as a *last resort* option.

In short: blindly using STONITH without true need during failover is a violation of the best effort principle. You are simply not doing your best.

- When your budget is limited, carefully select those improvements which make your system **as reliable as possible**, given your fixed budget.

⁷Finite automata are known to be transformable to deterministic ones, usually by an exponential increase in the number of states.


⁸For some *special(!)* cases of the shared-disk model, there exist some justifications for doing STONITH *before* starting the application at the hot standby. Under certain circumstances, it can happen that system A running amok could destroy the data on your single shared disk (example: a filesystem doubly mounted *in parallel*, which will certainly destroy your data, except you are using `ocfs2` or suchalike). This argument is only valid for *passive* disks which are *directly* attached to *both* systems A and B, such that there is no *external* means for fencing the disk. In case of iSCSI running over ordinary network equipment such as routers or switches, the argument "fencing the disk is otherwise not possible" does not apply. You can interrupt iSCSI connection at the network gear, or you can often do it at cluster A or at the iSCSI target. Even commercial storage appliances speaking iSCSI can be remotely controlled for forcefully aborting iSCSI sessions. In modern times, the STONITH method has no longer such a justification. The justification stems from ancient times when a disk was a purely passive mechanical device, and its disk controller was part of the server system.

⁹Notice that STONITH may be needed for (manual or partially automatic) *repair* in some cases, e.g. when you know that a system has a kernel crash. Don't mix up the repair phase with failover or handover phases. Typically, they are executed at different times. The repair phase is outside the scope of this section.

7. Tips and Tricks

- Create statistics on the duration of your actions. Based on this, try to get a *balanced* optimum between time and costs.
- Whatever actions you can **start in parallel** for saving time, do it. Otherwise you are disregarding the best effort principle, and your solution will be sub-optimal. You will require deep knowledge of parallel systems, as well as experience with dealing with problems like (distributed) races. Notice that *any* distributed system is *inherently parallel*. Don't believe that sequential methods can deliver an optimum solution in such a difficult area.
- If you don't have the **necessary skills** for (a) recognizing already existing parallelism, (b) dealing with parallelism at concept level, (c) programming and/or configuring parallelism race-free and deadlock-free (or if you even don't know what a race condition is and where it may occur in practice), then don't take responsibility for millions of € or \$.
- Avoid hard timeouts wherever possible. Use **adaptive timeouts** instead. Reason: depending on hardware or workload, the same action A may take a very short time on cluster 1, but take a very long time on cluster 2. If you need to guard action A from hanging (which is almost always the case because of Murphy's law), don't configure any fixed timeout for it. When having several hundreds of clusters, you would need to use the *worst case value*, which is the longest time occurring somewhere at the very slow clusters / slow parts of the network. This wastes a lot of time in case one of the fast clusters is hanging. Adaptive timeouts work differently: they use a kind of "progress bar" to monitor the *progress* of an action. They will abort only if there is *no progress* for a certain amount of time. Hint: among others, `marsadm view-*-rest` commands or macros are your friend.


ITON = Ignore The Other Node This means **fencing from application traffic**, and can be used as an alternative to STONITH when done properly.




Fencing from application traffic is best suited for the shared-nothing model, but can also be adapted to the shared-disk model with some quirks.

The idea is simple: always route your application network traffic to the current (logically) active side, whether it is currently A or B. Just don't route any application requests to the current (logically) passive side at all.

For failover (and *only* for that), you *should not care about* any split brain occurring at the low-level generic block device:



Although having a split brain at the generic low-level block device, you now define the "logically active" and "logically passive" side by yourself by *logically ignoring* the "wrong" side as defined by yourself:



This is possible because the generic block devices provided by DRBD or MARS are completely **agnostic** of the "meaning" of either version A or B. Higher levels such as cluster managers (or humans like sysadmins) can assign them a meaning like "relevant" or "not relevant", or "logically active" or "logically passive".

7.1. Avoiding Inappropriate Clustermanager Types for Medium and Long-Distance Replication

As a result of fencing from application traffic, the “logically passive” side will *logically* cease any actions such as updating user data, even if it is “physically active” during split-brain (when two primaries exist in DRBD or MARS sense¹⁰).

If you already have some load balancing, or BGP, or another *mechanism* for dynamic routing, you already have an important part for the ITON method. Additionally, ensure by an appropriate *strategy* that your balancer status / BGP announcement etc does always coincide with the “logically active” side (recall that even during split-brain *you* must define “logically active” **uniquely**¹¹ by yourself).

Example:

Phase1 Check whether the hot standby B is currently usable. If this is violated (which may happen during certain types of disasters), abort the failover for any affected resources.

Phase2 Do the following *in parallel*¹²:

- Start all affected applications at the hot standby B. This can be done with the same DRBD or MARS procedure as described on page 83.
- Fence A by fixedly routing all affected application traffic to B.

That’s all which has to be done for a shared-nothing model. Of course, this will likely produce a split-brain (even when using DRBD in place of MARS), but that will not matter from a user’s perspective, because the users will no longer “see” the “logically passive” side A through their network. Only during the relatively small time period where application traffic was going to the old side A while not replicated to B due to the incident, a very small number of updates *could* have gone lost. In fields like webhosting, this is taken into account. Users will usually not complain when some (smaller amount of) data is lost due to split-brain. They will complain when the service is unavailable.

This method is the fastest for restoring availability, because it doesn’t try to execute any (remote) action at side A. Only from a sysadmin’s perspective, there remain some cleanup tasks to be done during the following repair phase, such as split-brain resolution, which are outside the scope of this treatment.

By running the application fencing step *sequentially* (including wait for its partial successfulness such that the old side A can no longer be reached by any users) in front of the failover step, you may minimize the amount of lost data, but at the cost of total duration. Your service will take longer to be available again, while the amount of lost data is typically somewhat smaller.



A few people might clamour when some data is lost. In long-distance replication scenarios with high update traffic, there is *simply no way at all* for guaranteeing that no data can be lost ever. According to the laws of Einstein and the laws of Distributed Systems like the famous CAP theorem, this isn’t the fault of DRBD+proxy or MARS, but simply the *consequence* of having long distances. If you want to protect against data loss as best as possible, then don’t use $k = 2$ replicas. Use $k \geq 4$, and spread them over different distances, such as mixed small + medium + long distances. Future versions of MARS will support adaptive pseudo-synchronous modes, which will allow individual adaptation to network latencies / distances.

¹⁰Hint: some clustermanagers and/or some people seem to define the term “split-brain” differently from DRBD or MARS. In the context of generic block devices, split brain means that the *history* of both versions has been split to a Y-like **fork** (for whatever reason), such that re-joining them *incrementally* by ordinary write operations is no longer guaranteed to be possible. As a slightly simplified definition, you might alternatively use the definition “two incompatible primaries are existing in parallel”, which means almost the same in practice. Details of formal semantics are not the scope of this treatment.

¹¹A possible strategy is to use a Lamport clock for route changes: the change with the most recent Lamport timestamp will always win over previous changes.

¹²For database applications where no transactions should get lost, you should slightly modify the order of operations: first fence the old side A, then start the application at standby side B. However, be warned that even this cannot guarantee that no transaction is lost. When the network between A and B is interrupted *before* the incident happens, DRBD will automatically disconnect, and MARS will show a lagbehind. In order to fully eliminate this possibility, you can either use DRBD and configure it to hang forever during network outages (such that users will be unable to commit any transactions at all), or you can use the shared-disk model instead. But in the latter case, you are introducing a SPOF at the single shared disk. The former case is logically almost equivalent to shared-disk, but avoiding some parts of the physical SPOF. In a truly distributed system, the famous CAP theorem is limiting your possibilities. Therefore, no general solution exists fulfilling all requirements at the same time.

7. Tips and Tricks

The ITON method can be adapted to shared-disk by additionally fencing the common disk from the (presumably) failed cluster node A.

7.1.4.2. Handover Methods

Planned handover is conceptually simpler, because both sides must be (almost) healthy as a *precondition*. There are simply no pre-existing failures to deal with.

Here is an example using DRBD, some application commands denoted as pseudo code:

1. at side A: `applicationmanager stop all`
2. at side A: `drbdadm secondary all`
3. at side B: `drbdadm primary all`
4. at side B: `applicationmanager start all`

MARS already has a conceptual distinction between handover and failover. With MARS, it becomes even simpler, because a generic handover procedure is already built in:

1. at side A: `applicationmanager stop all`
2. at side B: `marsadm primary all`
3. at side B: `applicationmanager start all`

7.1.4.3. Hybrid Methods

In general, a planned handover may fail at any stage. Notice that such a failure is also a failure, but (partially) caused by the planned handover. You have the following alternatives for automatically dealing with such cases:

1. In case of a failure, switch back to the old side A.
2. Instead, forcefully switch to the new side A, similar to the methods described in section [7.1.4.1](#).

Similar options exist for a failed failover (at least in theory), but chances are lower for actually recovering if you have only $k = 2$ replicas in total.

Whatever you decide to do in what case in whatever priority order, whether you decide it in advance or during the course of a failing action: it simply means that according to the best effort principle, you should **never leave your system in a broken state** when there exists a chance to recover availability with any method.

Therefore, you should *implement* neither handover nor failover in their pure forms. Always implement hybrid forms following the best effort principle.

7.1.5. Special Requirements for Long Distances

Most contemporary cluster managers have been constructed for short distance shared-nothing clusters, or even for *local* shared-nothing clusters (c.f. DRBD over crossover cables), or even for shared-disk clusters (*originally*, when their *concepts* were developed). Blindly using them for long-distance replication without modification / adaptation bears some additional risks.

- Notice that long-distance replication always *requires* a **shared-nothing** model.
- As a consequence, **split brain** can appear *regularly* during failover. There is no way for preventing it! This is an *inherent property* of distributed systems, not limited to MARS (e.g. also occurring with DRBD if you try to use it over long distances). Therefore, you *must* deal with occurrences of split-brain as a *requirement*.
- The probability of **network partitions** is much higher: although you should have been required by Murphy's law to deal with network partitions already in short-distance scenarios, it now becomes *mandatory*.
- Be prepared that in case of certain types of (more or less global) internet partitions, you may not be able to trigger STONITH actions *at all*. Therefore, **fencing of application traffic** is *mandatory*.

7.2. Creating Backups via Pseudo Snapshots

When all your secondaries are all homogenously located in a standby datacenter, they will be almost idle all the time. This is a waste of computing resources.

Since MARS is no substitute for a full-fledged backup system, and since backups may put high system load onto your active side, you may want to utilize your passive hardware resources in a better way.

MARS supports this thanks to its ability to switch the `pause-replay` *independently* from `pause-fetch`.

The basic idea is simple: just use `pause-replay` at your secondary site, but leave the replication of transaction logfiles intact by deliberately *not* saying `pause-fetch`. This way, your secondary replica (block device) will stay frozen for a limited time, without loosing your redundancy: since the transaction logs will continue to replicate in the meantime, you can start `resume-replay` at any time, in particular when a primary-side incident should happen unexpectedly. The former secondary will just catch up by replaying the outstanding parts of the transaction logs in order to become recent.

However, some *details* have to be obeyed. In particular, the current version of MARS needs an additional `detach` operation, in order to release exclusive access to the underlying disk `/dev/lv/$res`. Future versions of MARS are planned to support this more directly, without need for an intermediate `detach` operation.



Beware: `mount -o ro /dev/vg/$res` can lead to **unnoticed write operations** if you are not careful! Some journalling filesystems like `xfs` or `ext4` may replay their journals onto the disk, leading to *binary* differences and thus **destroying your consistency** later when you re-enable `resume-replay`!



Therefore, you may use small LVM snapshots (only in such cases). Typically, `xfs` journal replay will require only a few megabytes. Therefore you typically don't need much temporary space for this. Here is a more detailed description of steps:

1. `marsadm pause-replay $res`
2. `marsadm detach $res`
3. `lvcreate --size 100m --snapshot --name ro-$res /dev/vg/$res`
4. `mount -o ro /dev/vg/ro-$res /mnt/tmp`
5. Now draw your backup from `/mnt/tmp/`
6. `umount /mnt/tmp`
7. `lvremove -f /dev/vg/ro-$res`
8. `marsadm up $res`

Hint: during the backup, the transaction logs will accumulate on `/mars/`. In order to avoid overflow of `/mars/` (c.f. section 4.4), don't unnecessarily prolong the backup duration.

8. MARS for Developers

This chapter is organized strictly top-down.

If you are a sysadmin and want to inform yourself about internals (useful for debugging), the relevant information is at the beginning, and you don't need to dive into all technical details at the end.

If you are a kernel developer and want to contribute code to the emerging MARS community, please read it (almost) all. Due to the top-down organization, sometimes you will need to follow some forward references in order to understand details. Therefore I recommend reading this chapter twice in two different reading modes: in the first reading pass, you just get a raw network of principles and structures in your brain (you don't want to grasp details, therefore don't strive for a full understanding). In the second pass, you will exploit your knowledge from the first pass for a deeper understanding of the details.

Alternatively, you may first read the sections about general architecture, and then start a bottom-up scan by first reading the last section about generic objects and aspects, and working in reverse *section* order (but read *subsections* in-order) until you finally reach the kernel interfaces / symlink trees.

8.1. Motivation / Politics

MARS is not yet upstream in the Linux kernel. This section tries to clear up some potential doubts. Some people have asked why MARS uses its own internal framework instead of *directly*¹ being based on some already existing Linux kernel infrastructures like the device mapper. Here is a list of technical reasons:

1. The existing device mapper infrastructure is based on `struct bio`. In contrast, the new XIO personality of the generic brick infrastructure is based on the concept of AIO (Asynchronous IO), which is a **true superset** of block IO.
2. In particular, `struct bio` is firmly referencing to `struct page` (via intermediate `struct bio_vec`), using types like `sector_t` in the field `bi_sector`. Basic transfer units are blocks, or sectors, or pages, or the like. In contrast, `struct aio_object` used by the XIO personality can address **arbitrary granularity** memory with byte resolution even at odd² positions in (virtual) files / devices, similar to classical Unix file IO, but *asynchronously*. Practical experience shows that even non-functional properties like performance of many datacenter workloads are profiting from that³. The AIO/XIO abstraction contains no fixed link to kernel abstractions and should be **easily portable** to other environments. In summary, the new personality provides a uniform abstraction which abstracts away from multiple different kernel interfaces; it is designed to be useful even in userspace.
3. Kernel infrastructures for the concept of *direct IO* are different from those for *buffered IO*. The XIO personality used by MARS subsumes both concepts as use case *variants*.

¹Notice that *indirect* use of pre-existing Linux infrastructure is not only possible, but actually implemented, by using it *internally* in brick *implementations* (black-box principle). However, such bricks are not portable to other environments like userspace.

²Some brick *implementations* (as opposed to the capabilities of the *interface*) may be (and, in fact, *are*) restricted to `PAGE_SIZE` operations or the like. This is no general problem, because IOP can automatically insert some translator bricks extending the capabilities to universal granularity (of course at some performance costs).

³The current transaction logger uses variable-sized headers at “odd” addresses. Although this increases `mempcpy()` load due to “misalignment”, the *overall performance* was provably better than in variants where sector / page alignment was strictly obeyed, but space was wasted for alignments. Such functionality is only possible if the XIO infrastructure *allows for* (but doesn't force) “mis-aligned” IO operations. In future, many different transaction logfile formats showing different runtime behaviour (e.g. optimized for high-throughput SSD loads) may co-exist in parallel. Note that properly aligned XIO operations bear no noticeable overhead compared to classical block IO, at least in typical datacenter RAID scenarios.

Buffering is an optional internal property of XIO bricks (almost non-functional property with support for consistency guarantees).

4. The AIO/XIO personality is generically designed for remote operations over networks, at arbitrary places in the IO stack, with (almost⁴) no semantic differences to local operations (built-in **network transparency**). There are universal provisions for mixed operation of different versions (**rolling software updates** in clusters / grids).
5. The generic brick infrastructure (as well as its personalities like XIO or any other future personality) supports **dynamic re-wiring / re-configuration** *during* operation (even while parallel IO requests are flying, some of them taking different paths in the IO stack in parallel). This is absolutely needed for MARS logfile rotation. In the long term, this would be useful for many advanced new features and products, not limited to multipathing.
6. The generic brick infrastructure (and in turn all personalities) provide **additional comfort** to the programmer while enabling **increased functionality**: by use of a generalization of **aspect orientation**⁵, the programmer need no longer worry about dynamic memory allocations for *local state* in a brick instance. MARS is **automating local state** even when dynamically instantiating new bricks (possibly having the same brick type) at runtime. Specifically, XIO is automating **request stacking** at the completion path this way, even while dynamically reconfiguring the IO stack⁶. A similar automation⁷ does not exist in the rest of the Linux kernel.
7. The generic brick infrastructure, together with personalities like XIO, enables **new long-term functional and non-functional opportunities** by use of concepts from instance-oriented programming (IOP⁸). The application area is **not limited to device drivers**. For example, a new personality for *stackable filesystems* could be developed in future.

In summary, anyone who would insist that MARS should be *directly*⁹ based on pre-existing kernel structures / frameworks instead of contributing a new framework would cause a *massive regression of functionality*.

- On one hand, all code contributed by the MARS project is **non-intrusive** into the rest of the Linux kernel. From the viewpoint of other parts of the kernel, the whole addition *behaves like* a driver (although its infrastructure is much more than a driver).
- On the other hand, if people are interested, the contributed infrastructure *may* be used to *add* to the power of the Linux kernel. It is designed to be **open for contributions**.

⁴By default, automatic network connection re-establishment and infinite network retries are already implemented in the `xio_client` and `xio_server` bricks to provide fully transparent semantics. However, this may be undesirable in case of fatal crashes. Therefore, abort operations are also configurable, as well as network timeouts which are then mapped to classical IO errors.

⁵Similar to AOP, insertion of IOP bricks for checking / debugging etc is one of the key advantages of the generic brick infrastructure. In contrast to AOP where debugging is usually {en,dis}abled statically at compile time, IOP allows for *dynamic* (re-)configuration of debugging bricks, automatic repair, and many more features promoted by *organic computing*.

⁶The generic aspect orientation approach leads to better **separation of concerns**: local state needed by brick implementations is not visible from outside by default. In other words, local state is also **private state**. Accidental hampering of internal operations is impeded.

Example from the kernel: in `include/linux/blkdev.h` the definition of `struct request` contains the following comment: `/* the following two fields are internal, NEVER access directly */`. It appears that `struct request` contains not only fields relevant for the caller, but also **internal fields** needed only in *some specific* callees. For example, `rb_node` is documented to be used only in IO schedulers.

XIO goes one step further: there need not exist exactly one IO scheduler instance in the IO stack for a single device. Future `xio_scheduler_{deadline,cfq,...}` brick types could be each instantiated many times, and in arbitrary places, even for the same (logical) device. The equivalent of `rb_node` would then be automatically instantiated multiple times for the same IO request, by automatically instantiating the right local aspect instances.

⁷DM can achieve stacking and dynamic routing by a workaround called *request cloning*, potentially leading to mass creation of temporary / intermediate object instances.


⁸See http://athomux.net/papers/paper_inst2.pdf

⁹Notice that kernel-specific structures like `struct bio` are of course used by MARS, but only *inside* the blackbox implementation of bricks like `mars_bio` or `mars_if` which act as **adaptors** to/from that structure. It is possible to write further adaptors, e.g. for direct interfacing to the device mapper infrastructure.

8. MARS for Developers

- A *possible* (but not the only possible) way to do this is giving the generic brick framework / the XIO personality as well as future personalities / the MARS application the status of a *subsystem* inside the kernel (in the long term), similar to the SCSI subsystem or the network subsystem. Noone is forced to use it, but anybody may use it if he/she likes.
- Politically, the author is a FOSS advocate willing to collaborate and to support anyone interested in contributions. The author's personal interest is long-term and is open for both in-tree and out-of-tree extensions of both the framework and MARS by any other party obeying the GPL and not hazarding FOSS by patents (instead supporting organizations like the Open Invention Network). The author is open to closer relationships with the Linux Foundation and other parts of the Linux ecosystem.

8.2. Architecture Overview

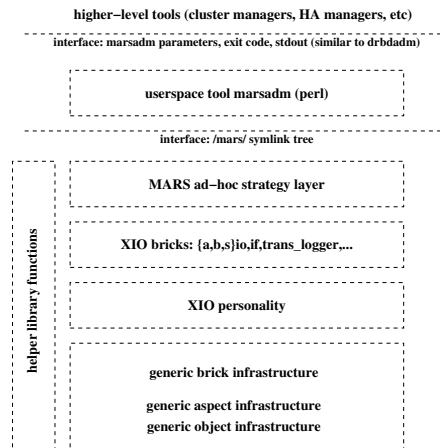


8.3. Some Architectural Details

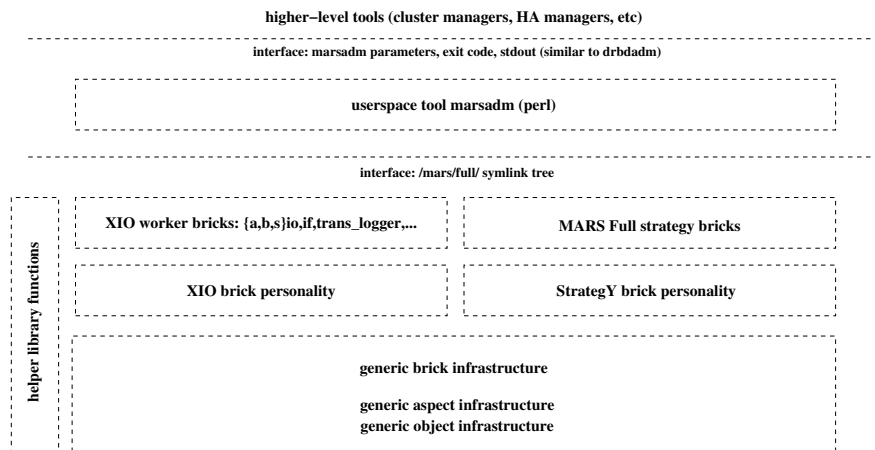
The following pictures show some “zones of responsibility”, not necessarily a strict hierarchy (although Dijkstra’s famous layering rules from THE are tried to be respected as much as possible). The construction principle follows the concept of **Instance Oriented Programming** (IOP) described in http://athomux.net/papers/paper_inst2.pdf. Please note that MARS is only instance-based¹⁰, while MARS Full is planned to be fully instance-oriented.

8.3.1. MARS Architecture

¹⁰Similar to OOP, where “object-based” means a weaker form of “object-oriented”, the term “instance-based” means that the *strategy* brick layer need not be fully modularized according to the IOP principles, but the *worker* brick layer already is.



8.3.2. MARS Full Architecture (planned)



8.4. Documentation of the Symlink Trees

The `/mars/` symlink tree is serving the following purposes, all at the same time:

1. For **communication** between cluster nodes, see sections 4.2 and 4.3. This communication is even the *only* communication between cluster nodes (apart from the *contents* of transaction logfiles and sync data).
2. **Internal interface** between the kernel module and the userspace tool `marsadm`.
3. **Internal persistent repository** which keeps state information between reboots (also in case of node crashes). It is even the *only* place where state information is kept. There is no other place like `/etc/drbd.conf`.



Because of its internal character, its representation and semantics may change at any time without notice (e.g. via an *internal* upgrade procedure between major releases). It is *not* an external interface to the outer world. Don't build anything on it.

However, knowledge of the symlink tree is useful for advanced sysadmins, for **human inspection** and for **debugging**. And, of course, for developers.

As an “official” interface from outside, only the `marsadm` command should be used.

8. *MARS for Developers*

8.4.1. Documentation of the MARS Symlink Tree

8.5. XIO Worker Bricks

8.6. StrategY Worker Bricks

NYI

8.7. The XIO Brick Personality

8.8. The Generic Brick Infrastructure Layer

8.9. The Generic Object and Aspect Infrastructure

A. Technical Data MARS

MARS has some built-in limitations which should be overcome¹ by the future MARS Full. Please don't exceed the following limits:

- maximum 10 nodes per cluster
- maximum 10 resources per cluster
- maximum 100 logfiles per resource

¹Some internal algorithms are quadratic. The reason is that MARS evolved from a lab prototype which wasn't originally intended for enterprise grade usage, but should have been succeeded by the fully instance-oriented MARS Full much earlier.

B. Handout for Midnight Problem Solving

Here are generic instructions for the generic `marsadm` and commandline level. Other levels (e.g. different types of cluster managers, PaceMaker, control scripts / rc scripts / upstart scripts, etc should be described elsewhere.

B.1. Inspecting the State of MARS

For manual inspection, please prefer the new `marsadm view all` over the old `marsadm view-1and1 all`. It shows more appropriate / detailed information.

Hint: this might change in future when somebody will program better macros for the `view-1and1` variant, or create even better other macros.

```
# watch marsadm view all
```

Checking the low-level network connections at runtime:

```
# watch "netstat -tcp | grep 777"
```

Meaning of the port numbers (as currently configured into the kernel module, may change in future):

- 7777 = metadata / symlink propagation
- 7778 = transfer of transaction logfiles
- 7779 = transfer of sync traffic

7777 must be always active on a healthy cluster. 7778 and 7779 will appear only on demand, when some data is transferred.

Hint: when one of the columns Send-Q or Recv-Q are constantly at high values, you might have a network bottleneck.

B.2. Replication is Stuck

Indications for a stuck:

- One of the flags shown by `marsadm view all` or `marsadm view-flags all` contain a symbol "-" (dash). This means that some switch is currently switched off (deliberately). Please check whether there is a valid reason why somebody else switched it off. If the switch-off is just by accident, use the following command to fix the stuck:

```
# marsadm up all
```

(or replace `all` by a particular resource name if you want to start only a specific one).

Note: `up` is equivalent to the sequence `attach; resume-fetch; resume-replay; resume-sync`.

Instead of switching each individual knob, use `up` as a shortcut for switching on anything which is currently off.

- `netstat --tcp | grep 7777` does not show anything. Please check the following:

- Is the kernel module loaded? Check `lsmod | grep mars`. When necessary, run `modprobe mars`.
 - Is the network interface down? Check `ifconfig`, and/or `ethtool` and friends, and fix it when necessary.
 - Is a `ping <partner-host>` possible? If not, fix the network / routing / firewall / etc. When fixed, the MARS connections should automatically appear after about 1 minute.
 - When `ping` is possible, but a MARS connection to port 7777 does not appear after a few minutes, try to connect to remote port 7777 by hand via `telnet`. But don't type anything, just abort the connection immediately when it works! Typing anything will almost certainly throw a harsh error message at the other server, which could unnecessarily alarm other people.
- Check whether `marsadm view all` shows some progress bars somewhere. Example:

```

istore-test-bap1:~# marsadm view all
----- resource lv-0
lv-0 OutDated[F] PausedReplay dCAS-R Secondary istore-test-bs1
  replaying: [ >.....] 1.21% (12/1020)MiB logs: [2..3]
    > fetch: 1008.198 MiB rate: 0 B/sec remaining: --:--:-- hrs
    > replay: 0 B rate: 0 B/sec remaining: 00:00:00 hrs

```

At least one of the `rate:` values should be greater than 0. When none of the `rate:` values indicate any progress for a longer time, try `marsadm up all` again. If it doesn't help, check and repair the network. If even this does not help, check the hardware for any IO hangups, or kernel hangups. First, check the RAID controllers. Often (but not certainly), a stuck kernel can be recognized when many processes are *permanently* in state "D", for a long time: `ps ax | grep "D" | grep -v grep` or similar. Please check whether there is just an overload, or *really* a true kernel problem. Discrimination is not easy, and requires experience (as with any other system; not limited to MARS). A truly stuck kernel can only be resurrected by rebooting. The same holds for any hardware problems.

- Check whether `marsadm view all` reports any lines like `WARNING: SPLIT BRAIN at '' detected`. In such a case, check that there is *really* a split brain, before obeying the instructions in section B.4. Notice that network outages or missing `marsadm log-delete-all all` may continue to report an old split brain which has gone in the meantime.
- Check whether `/mars/` is too full. For a rough impression, `df /mars/` may be used. For getting authoritative values as internally used by the MARS emergency-mode computations, use `marsadm view-rest-space` (the unit is GiB). In practice, the differences are only marginal, at least on bigger `/mars/` partitions. When there is only few rest space (or none at all), please obey the instructions in section B.3.

B.3. Resolution of Emergency Mode

Emergency mode occurs when `/mars/` runs out of space, such that no new logfile data can be written anymore.

In emergency mode, the primary will write any write requests *directly* to the underlying disk, as if MARS were not present at all. Thus, your application will continue to run. Only the *replication* as such is stopped.

Notice: emergency mode means that your secondary nodes are usually in a *consistent*, but *outdated* state (exception: when a sync was running in parallel to the emergency mode, then the sync will be automatically started over again). You can check consistency via `marsadm view-flags all`. Only when a local disk shows a lower-case letter "d" instead of an uppercase "D", it is known to be inconsistent (e.g. during a sync). When there is a dash instead, it usually means that the disk is detached or misconfigured or the kernel module is not started. Please fix these problems first before believing that your local disk is unusable. Even if it is really inconsistent (which is very unlikely, typically occurring only as a consequence of hardware

B. Handout for Midnight Problem Solving

failures, or of the above-mentioned exception), you have a big chance to recover most of the data via `fsck` and `friends`.

A currently existing Emergency mode can be detected by

```
primary:~# marsadm view-is-emergency all
secondary:~# marsadm view-is-emergency all
```

Notice: this delivers the current state, telling nothing about the past.

Currently, emergency mode will also show something like `WARNING: SPLIT BRAIN at "detected`. This ambiguity will be resolved in a future MARS release. It is however not crucial: the resolution methods for both cases are very similar. If in doubt, start emergency resolution first, and only proceed to split brain resolution if it did not help.

Preconditions:

- Only current version of MARS: the space at the primary side should have been already released, and the emergency mode should have been already left. Otherwise, you might need the split-brain resolution method from section [B.4](#).
- The network **must** be working. Check that the following gives an entry for each secondary:

```
primary:~# netstat --tcp | grep 7777
```

When necessary, fix the network first (see instructions above).

Emergency mode should now be resolved via the following instructions:

```
primary:~# marsadm view-is-emergency all
primary:~# du -s /mars/resource-* | sort -n
```

Remember the affected resources. Best practice is to do the following, starting with the *biggest* resource as shown by the `du | sort` output in reverse order, but *starting* the following only with the *affected* resources in the first place:

```
secondary1:~# marsadm invalidate <res1>
secondary1:~# marsadm log-delete-all all
... dito with all resources showing emergency mode
... dito on all other secondaries
primary:~# marsadm log-delete-all all
```

Hint: during the resolution process, some other resources might have gone into emergency mode concurrently. In addition, it is possible that some secondaries are stuck at particular resources while the corresponding primary has *not yet* entered emergency mode. Please repeat the steps in such a case, and look for emergency modes at secondaries additionally. When necessary, extend your list of *affected* resources.

Hint: be patient. Deleting large bulks of logfile data may take a long time, at least on highly loaded systems. You should give the cleanup processes at least 5 minutes before concluding that an `invalidate` followed by `log-delete-all` had no effect! Don't forget to give the `log-delete-all` at all cluster nodes, even when seemingly unaffected.

In very complex scenarios, when the primary roles of different resources are spread over different hosts (aka mixed operation), you may need to repeat the whole cycle iteratively for a few cycles until the jam is resolved.

If it does not go away, you have another chance by the following split-brain resolution process, which will also cleanup emergency mode as a side effect.

B.4. Resolution of Split Brain and of Emergency Mode

Hint: in many cases (but not guaranteed), the previous recipe for resolution of emergency mode will also cleanup split brain. Good chances are in case of $k = 2$ total replicas. Please collect your own experiences which method works better for you!

Precondition: the network must be working. Check that the following gives an entry for each secondary:

```
primary:~# netstat --tcp | grep 7777
```

When necessary, fix the network first (see instructions above).

Inspect the split brain situation:

```
primary:~# marsadm view all
primary:~# du -s /mars/resource-* | sort -n
```

Remember those resources where a message like `WARNING: SPLIT BRAIN at '' detected` appears. Do the following only for *affected* resources, starting with the biggest one (before proceeding to the next one).

Do the following with only *one* resource at a time (before proceeding to the next one), and repeat the actions on that resource at every secondary (if there are multiple secondaries):

```
secondary1:~# marsadm leave-resource $res1
secondary1:~# marsadm log-delete-all all
```

Check whether the split brain has vanished everywhere. Startover with other resources at their secondaries when necessary.

Finally, when no split brain is reported at any (former) secondary, do the following on the primary:

```
primary:~# marsadm log-delete-all all
primary:~# sleep 30
primary:~# marsadm view all
```

Now, the split brain should be gone even at the primary. If not, repeat this step.

In case even this should fail on some `$res` (which is very unlikely), read the PDF manual before using `marsadm log-purge-all $res`.

Finally, when the split brain is gone everywhere, rebuild the redundancy at every secondary via

```
secondary1:~# marsadm join-resource $res1 /dev/<lv-x>/$res1
```

If even this method does not help, setup the whole cluster afresh by `rmmod mars` everywhere, and creating a fresh `/mars/` filesystem everywhere, followed by the same procedure as installing MARS for the first time (which is outside the scope of this handout).

B.5. Handover of Primary Role

When there exists a method for primary handover in higher layers such as cluster managers, please prefer that method (e.g. `cm3` or other tools).

If such alike doesn't work, or if you need to handover some resource `$res1` by hand, do the following:

- Stop the load / application corresponding to `$res1` on the old primary side.
- `umount /dev/mars/$res1`, or otherwise close any openers such as iSCSI.
- At the new primary: `marsadm primary $res1`
- Restart the application at the new site (in reverse order to above). In case you want to switch *all* resources which are not yet at the new side, you may use `marsadm primary all`.

B.6. Emergency Switching of Primary Role

Emergency switching is necessary when your primary is no longer reachable over the network for a *longer* time, or when the hardware is defective.

Emergency switching will very often lead to a split brain, which requires lots of manual actions to resolve (see above). Therefore, try to avoid emergency switching when possible!

Hint: MARS can automatically recover after a primary crash / reboot, as well as after secondary crashes, just by executing `modprobe mars` after `/mars/` had been mounted. Please consider to wait until your system comes up again, instead of risking a split brain.

B. Handout for Midnight Problem Solving

The decision between emergency switching and continuing operation at the same primary side is an operational one. MARS can support your decision by the following information at the potentially new primary side (which was in secondary mode before):

```
istore-test-bap1:~# marsadm view all
----- resource lv-0
lv-0 InConsistent Syncing dcAsFr Secondary istore-test-bs1
syncing: [====>.....] 27.84% (567/2048)MiB rate: 72583.00 KiB/sec remaining: 00:00:20 hrs
> sync: 567.293/2048 MiB rate: 72583 KiB/sec remaining: 00:00:20 hrs
replaying: [>::::::::::] 0.00% (0/12902)KiB logs: [1..1]
> fetch: 0 B rate: 38 KiB/s remaining: 00:00:00
> replay: 12902.047 KiB rate: 0 B/s remaining: ---:---:---
```

When your target is syncing (like in this example), you cannot switch to it (same as with DRBD). When you had an emergency mode before, you should first resolve that (whenever possible). When a split brain is reported, try to resolve it first (same as with DRBD). Only in case you *know* that the primary is really damaged, or it is really impossible to run the application there for some reason, emergency switching is desirable.

Hint: in case the secondary is inconsistent for some reason, e.g. because of an incremental fast full-sync, you have a last chance to recover most data after forceful switching by using a filesystem check or suchalike. This might be even faster than restoring data from the backup. But use it only if you are *really* desperate!

The amount of data which is *known* to be missing at your secondary is shown after the `> fetch:` in human-readable form. However, in cases of networking problems this information may be outdated. You *always* need to consider further facts which cannot be known by MARS.

When there exists a method for emergency switching of the primary in higher layers such as cluster managers, please prefer that method in front of the following one.

If suchalike doesn't work, or when a handover attempt has failed several times, or if you *really need* forceful switching of some resource `$res1` by hand, you can do the following:

- When possible, stop the load / application corresponding to `$res1` on the old primary side.
- When possible, `umount /dev/mars/$res1`, or otherwise close any openers such as iSCSI.
- When possible (if you have some time), wait until as much data has been propagated to the new primary as possible (watch the `fetch:` indicator).
- At the new primary: `marsadm disconnect $res1; marsadm primary --force $res1`
- Restart the application at the new site (in reverse order to above).
- After the application is known to run reliably, check for split brains and cleanup them when necessary.

C. Alternative Methods for Split Brain Resolution

Instead of `marsadm invalidate`, the following steps may be used. In preference, start with the old “wrong” primaries first:

1. `marsadm leave-resource mydata`
2. After having done this on one cluster node, check whether the split brain is already gone (e.g. by saying `marsadm view mydata`). There are chances that you don’t need this on all of your nodes. Only in very rare¹ cases, it might happen that the preceding `leave-resource` operations were not able to clean up all logfiles produced in parallel by the split brain situation.
3. Read the documentation about `log-purge-all` (see page 73) and use it.
4. If you want to restore redundancy, you can follow-up a `join-resource` phase to the old resource name (using the correct device name, double-check it!) This will restore your redundancy by overwriting your bad split brain version with the correct one.



It is important to resolve the split brain *before* you can start the `join-resource` reconstruction phase! In order to keep as many “good” versions as possible (e.g. for emergency cases), don’t re-join them all in parallel, but rather start with the oldest / most outdated / worst / inconsistent version first. It is recommended to start the next one only when the previous one has successfully finished.

¹When your network had partitioned in a very awkward way for a long time, and when your partitioned primaries did several `log-rotate` operations indendently from each other, there is a small chance that `leave-resource` does not clean up *all* remains of such an awkward situation. Only in such a case, try `log-purge-all`.

D. Alternative De- and Reconstruction of a Damaged Resource

In case `leave-resource --host=` does not work, you may use the following fallback. On the surviving new designated primary, give the following commands:

1. `marsadm disconnect-all mydata`
2. `marsadm down mydata`
3. Check by hand whether your local disk is consistent, e.g. by test-mounting it readonly, `fsck`, etc.
4. `marsadm delete-resource mydata`
5. Check whether the other vital cluster nodes don't report the dead resource any more, e.g. `marsadm view all` at *each* of them. In case the resource has not disappeared anywhere (which may happen during network problems), do the `down` ; `delete-resource` steps also there (optionally again with `--force`).
6. Be sure that the resource has disappeared *everywhere*. When necessary, repeat the `delete-resource` with `--force`.
7. `marsadm create-resource newmydata ...` at the *correct* node using the *correct* disk device containing the *correct* version, and further steps to setup your resource from scratch, preferably under a different name to minimize any risk.

In any case, **manually check** whether a split brain is reported for any resource on any of your *surviving* cluster nodes. If you find one there (and only then), please (re-)execute the split brain resolution steps on the affected node(s).

E. Cleanup in case of Complicated Cascading Failures

MARS does its best to recover even from multiple failures (e.g. **rolling disasters**). Chances are high that the instructions from sections 3.4.3 3.4.4 or appendix C D will work even in case of multiple failures, such as a network failure plus local node failure at only 1 node (even if that node is the former primary node).

However, in general (e.g. when more than 1 node is damaged and/or when the filesystem `/mars/` is badly damaged) there is no general guarantee that recovery will *always* succeed under *any* (weird) circumstances. That said, your chances for recovery are *very* high when some disk remains usable at least at one of your surviving secondaries.



It should be very hard to finally trash a secondary, because the transaction logfiles are containing `md5` checksums for all data records. Any attempt to replay corrupted logfiles is refused by MARS. In addition, the sequence numbers of `log-rotated` logfiles are checked for contiguity. Finally, the *sequence path* of logfile applications (consisting of logfile names plus their respective length) is additionally secured by a `git`-like incremental checksum over the whole path history (so-called “version links”). This should detect split brains even if logfiles are appended / modified *after* a (forceful) switchover has already taken place.



That said, your risk of final data loss is very high if you remove the **BBU** from your hardware RAID controller before all hot data has been flushed to the physical disks. Therefore, never try to “repair” a seemingly dead node before your replication is up again somewhere else! Only unplug the network cables when advised, but never try to repair the hardware instantly!

In case of desperate situations where none of the previous instructions have succeeded, your last chance is rebuilding all your resources from intact disks as follows:

1. Do `rmmod mars` on all your cluster nodes and/or reboot them. Note: if you are less desperate, chances are high that the following will also work when the kernel module remains active and everywhere a `marsadm down` is given instead, but for an *ultimate* instruction you should eliminate *potential* kernel problems by `rmmod / reboot`, at least if you can afford the downtime on concurrently operating resources.
2. For safety, physically remove the storage network cables on *all* your cluster nodes. Note: the same disclaimer holds. MARS really does its best, even when `delete-resource` is given while the network is fully active and multiple split-brain primaries are actively using their local device in parallel (approved by some testcases from the automatic test suite, but note that it is impossible to catch all possible failure scenarios). Don’t challenge your fate if you are desperate! Don’t *rely* on this! Nothing is absolutely fail-safe!
3. **Manually** check which surviving disk is usable, and which is the “best” one for your purpose.
4. Do `modprobe mars only` on that node. If that fails, `rmmod` and/or reboot again, and start over with a completely fresh `/mars/` partition (`mkfs.ext4 /mars/` or similar) *everywhere* on *all* cluster nodes, and continue with step 7.
5. If your old `/mars/` works, and you did not already (forcefully) switch your designated primary to the final destination, do it now (see description in section 3.4.2.2). Wait until any old logfile data has been replayed.
6. Say `marsadm delete-resource mydata --force`. This will cleanup all internal symlink tree information for the resource, but will leave your disk data intact.

E. Cleanup in case of Complicated Cascading Failures

7. Locally build up the new resource(s) as usual, out of the underlying disks.
8. Check whether the new resource(s) work in standalone mode.
9. When necessary, repeat these steps with other resources.

Now you can choose how to rebuild your cluster. If you rebuilt `/mars/` anywhere, you *must* rebuild it on *all* new cluster nodes and start over with a fresh `join-cluster` on each of them, from scratch. It is not possible to mix the old cluster with the new one.

10. Finally, do all the necessary `join-resources` on the respective cluster nodes, according to your new redundancy scenario after the failures (e.g. after activating spare nodes, etc). If you have $k > 2$ replicas, start `join-resource` on the worst / most damaged version first, and start the next preferably only after the previous sync has completed successfully. This way, you will be permanently retaining some (old and outdated, but hopefully potentially usable) replicas while a sync is running. Don't start too many syncs in parallel.



Never use `delete-resource` twice on the same resource name, after you have already a working standalone primary¹. You might accidentally destroy your again-working copy! You *can* issue `delete-resource` multiple times on different nodes, e.g. when the network has problems, but doing so *after* re-establishment of the initial primary bears some risk. Therefore, the safest way is first deleting the resources everywhere, and then starting over afresh.

Before re-connecting any network cable on any non-primary (new secondaries), ensure that all `/dev/mars/mydata` devices are no longer in use (e.g. from an old primary role before the incident happened), and that each local disk is detached. Only after that, you should be able to safely re-connect the network. The `delete-resource` given at the new primary should propagate now to each of your secondaries, and your local disk should be usable for a `re-join-resource`.



When you did not rebuild your cluster from scratch with fresh `/mars/` filesystems, and one of the old cluster nodes is supposed to be removed permanently, use `leave-resource` (optionally with `--host=` and/or `--force`) and finally `leave-cluster`.

¹Of course, when you don't have created the *same* resource anew, you may repeat `delete-resource` on other cluster nodes in order to get rid of local files / symlinks which had not been propagated to other nodes before.

F. Experts only: Special Trick Switching and Rebuild

The following is a further alternative for **experts** who really know what they are doing. The method is very simple and therefore well-suited for coping with mass failures, e.g. **power blackout of whole datacenters**.

In case a primary datacenter fails as a whole for whatever reason and you have a backup datacenter, do the following steps in the backup datacenter:

1. Fencing step: by means of firewalling, **ensure** that the (virtually) damaged datacenter nodes **cannot** be reached over the network. For example, you may place REJECT rules into all of your local iptables firewalls at the backup datacenter. Alternatively / additionally, you may block the routes at the appropriate central router(s) in your network.
2. Run the sequence `marsadm disconnect all; marsadm primary --force all` on all nodes in the backup datacenter.
3. Restart your services in the backup datacenter (as far as necessary). Depending on your network setup, further steps like switching BGP routes etc may be necessary.
4. Check that *all* your services are *really* up and running, before you try to repair anything! Failing to do so may result in data loss when you execute the following restore method for *experts*.

Now your backup datacenter should continue servicing your clients. The final reconstruction of the originally primary datacenter works as follows:

1. At the damaged primary datacenter, ensure that nowhere the MARS kernel module is running. In case of a power blackout, you shouldn't have executed an automatic `modprobe mars` anywhere during reboot, so you should be already done when all your nodes are up again. In case some nodes had no reboot, execute `rmmod mars` everywhere. If `rmmod` refuses to run, you may need to unmount the `/dev/mars/mydata` device first. When nothing else helps, you may just mass reboot your hanging nodes.
2. At the failed side, do `rm -rf /mars/resource-$mydata/` for all those resources which had been primary before the blackout. Do this *only* for those cases, otherwise you will need unnecessary `leave-resources` or `invalidates` later (e.g. when half of your nodes were already running at the surviving side). In order to avoid unnecessary traffic, please do this only as far as really necessary. Don't remove any other directories. In particular, `/mars/ips/` *must* remain intact. In case you accidentally deleted them, or you had to re-create `/mars/` from scratch, try `rsync` with the correct options.



Caution! before doing this, check that the corresponding directory exists at the backup datacenter, and that it is *really* healthy!

3. Un-Fencing: restore your network firewall / routes and check that they work (`ping` etc).
4. Do `modprobe mars` everywhere. All missing directories and their missing symlinks should be automatically fetched from the backup datacenter.
5. Run `marsadm join-resource $res`, but only at those places where the directory was removed previously, while using the same disk devices as before. This will minimize actual traffic thanks to the fast full sync algorithm.

F. Experts only: Special Trick Switching and Rebuild



It is **crucial** that the fencing step **must** be executed *before* any **primary --force!** This way, no split brain will be *visible* at the backup datacenter side, because there is simply no chance for transferring different versions over the network. It is also crucial to remove any (potentially diverging) resource directories *before* the **modprobe!** This way, the backup datacenter never runs into split brain. This saves you a lot of detail work for split brain resolution when you have to restore bulks of nodes in a short time.



In case the repair of a full datacenter should take so extremely long that some **/mars/** partitions are about to run out of space at the surviving side, you may use the **leave-resource --host=failed-node** trick described earlier, followed by **log-delete-all**. Best if you have prepared a fully automatic script long before the incident, which executes such alike only as far as necessary in each individual case.



Even better: train such scenarios in advance, and prepare scripts for mass automation. Look into section **5.3**.

G. GNU Free Documentation License

GNU Free Documentation License
Version 1.3, 3 November 2008

Copyright (C) 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.
<<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not

G. GNU Free Documentation License

allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

The "publisher" means any person or entity that distributes copies of the Document to the public.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If

G. GNU Free Documentation License

- there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
 - K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
 - L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
 - M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
 - N. Do not retitling any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
 - O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements",

and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to

G. GNU Free Documentation License

60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

"Massive Multiauthor Collaboration Site" (or "MMC Site") means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A "Massive Multiauthor Collaboration" (or "MMC") contained in the site means any set of copyrightable works thus published on the MMC site.

"CC-BY-SA" means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

"Incorporate" means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is "eligible for relicensing" if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.