
Subject: [PATCH 0/7] SUNRPC: initial part of making pipefs work in net ns
Posted by [Stanislav Kinsbursky](#) on Fri, 28 Oct 2011 14:27:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

This patch-set was created in context of clone of git branch:
`git://git.linux-nfs.org/projects/trondmy/nfs-2.6.git`
and rebased on tag "v3.1".

This patch-set implements pipefs superblock creation per network namespace context instead of using single one for all possible contexts. Also, it provides pipefs dentries creation and destruction helpers for kernel routines.

Additional description of the idea about how to make RPC pipefs work per network namespace context can be found in the letter titled "SUNRPC: "RPC pipefs per network namespace" preparations", which has been sent already to `linux-nfs@vger.kernel.org`.

The following series consists of:

Stanislav Kinsbursky (7):

- SUNRPC: create RPC pipefs superblock per network namespace context
- SUNRPC: hold current network namespace while pipefs superblock is active
- SUNRPC: send notification events on pipefs sb creation and destruction
- SUNRPC: pipefs dentry lookup helper introduced
- SUNRPC: put pipefs superblock link on network namespace
- SUNRPC: pipefs per-net operations helper introduced
- SUNRPC: added debug messages to RPC pipefs

```
include/linux/sunrpc/rpc_pipe_fs.h | 15 ++++++
net/sunrpc/netns.h                  | 3 +
net/sunrpc/rpc_pipe.c               | 113 ++++++
net/sunrpc/sunrpc_syms.c            | 1
4 files changed, 130 insertions(+), 2 deletions(-)
```

Subject: [PATCH 3/7] SUNRPC: send notification events on pipefs sb creation and destruction
Posted by [Stanislav Kinsbursky](#) on Fri, 28 Oct 2011 14:27:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

They will be used to notify subscribers about pipefs superblock creation and destruction.
Subscribers will have to create their dentries on passed superblock on mount event and destroy otherwise.

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

```
include/linux/sunrpc/rpc_pipe_fs.h | 8 ++++++++
net/sunrpc/rpc_pipe.c               | 32 +++++++++++++++++++++++++++++++++++++
2 files changed, 40 insertions(+), 0 deletions(-)
```

diff --git a/include/linux/sunrpc/rpc_pipe_fs.h b/include/linux/sunrpc/rpc_pipe_fs.h

index cf14db9..c1cdb2f 100644

--- a/include/linux/sunrpc/rpc_pipe_fs.h

+++ b/include/linux/sunrpc/rpc_pipe_fs.h

```
@@ -44,6 +44,14 @@ RPC_I(struct inode *inode)
    return container_of(inode, struct rpc_inode, vfs_inode);
}
```

```
+extern int rpc_pipefs_notifier_register(struct notifier_block *);
+extern void rpc_pipefs_notifier_unregister(struct notifier_block *);
+
+enum {
+ RPC_PIPEFS_MOUNT,
+ RPC_PIPEFS_UMOUNT,
+};
+
+extern int rpc_queue_upcall(struct inode *, struct rpc_pipe_msg *);
```

struct rpc_clnt;

diff --git a/net/sunrpc/rpc_pipe.c b/net/sunrpc/rpc_pipe.c

index a717564..8abeb9d 100644

--- a/net/sunrpc/rpc_pipe.c

+++ b/net/sunrpc/rpc_pipe.c

```
@@ -28,8 +28,10 @@
#include <linux/sunrpc/rpc_pipe_fs.h>
#include <linux/sunrpc/cache.h>
#include <linux/nsproxy.h>
+#include <linux/notifier.h>
```

```
#include "netns.h"
```

```
+#include "sunrpc.h"
```

```
static struct vfsmount *rpc_mnt __read_mostly;
```

```
static int rpc_mount_count;
```

```
@@ -41,6 +43,20 @@ static struct kmem_cache *rpc_inode_cachep __read_mostly;
```

```
#define RPC_UPCALL_TIMEOUT (30*HZ)
```

```
+static BLOCKING_NOTIFIER_HEAD(rpc_pipefs_notifier_list);
```

```
+
```

```
+int rpc_pipefs_notifier_register(struct notifier_block *nb)
```

```

+{
+ return blocking_notifier_chain_cond_register(&rpc_pipefs_notifier_list, nb);
+}
+EXPORT_SYMBOL_GPL(rpc_pipefs_notifier_register);
+
+void rpc_pipefs_notifier_unregister(struct notifier_block *nb)
+{
+ blocking_notifier_chain_unregister(&rpc_pipefs_notifier_list, nb);
+}
+EXPORT_SYMBOL_GPL(rpc_pipefs_notifier_unregister);
+
static void rpc_purge_list(struct rpc_inode *rpci, struct list_head *head,
void (*destroy_msg)(struct rpc_pipe_msg *), int err)
{
@@ -1010,6 +1026,7 @@ rpc_fill_super(struct super_block *sb, void *data, int silent)
struct inode *inode;
struct dentry *root;
struct net *net = data;
+ int err;

sb->s_blocksize = PAGE_CACHE_SIZE;
sb->s_blocksize_bits = PAGE_CACHE_SHIFT;
@@ -1027,8 +1044,20 @@ rpc_fill_super(struct super_block *sb, void *data, int silent)
}
if (rpc_populate(root, files, RPCAUTH_lockd, RPCAUTH_RootEOF, NULL))
return -ENOMEM;
+ err = blocking_notifier_call_chain(&rpc_pipefs_notifier_list,
+ RPC_PIPEFS_MOUNT,
+ sb);
+ if (err)
+ goto err_depopulate;
sb->s_fs_info = get_net(net);
return 0;
+
+err_depopulate:
+ blocking_notifier_call_chain(&rpc_pipefs_notifier_list,
+ RPC_PIPEFS_UMOUNT,
+ sb);
+ __rpc_depopulate(root, files, RPCAUTH_lockd, RPCAUTH_RootEOF);
+ return err;
}

static struct dentry *
@@ -1043,6 +1072,9 @@ void rpc_kill_sb(struct super_block *sb)
struct net *net = sb->s_fs_info;

put_net(net);
+ blocking_notifier_call_chain(&rpc_pipefs_notifier_list,

```

```
+    RPC_PIPEFS_UMOUNT,  
+    sb);  
    kill_litter_super(sb);  
}
```

Subject: [PATCH 4/7] SUNRPC: pipefs dentry lookup helper introduced

Posted by [Stanislav Kinsbursky](#) on Fri, 28 Oct 2011 14:27:58 GMT

[View Forum Message](#) <> [Reply to Message](#)

In all places, where pipefs dentries are created, only directory inode is actually required to create new dentry. And all this directories has root pipefs dentry as their parent. So we actually don't need this pipefs mount point at all if some pipefs lookup method will be provided.

IOW, all we really need is just superblock and simple lookup method to find root's child dentry with appropriate name. And this patch introduces this method.

Note, that no locking implemented in `rpc_d_lookup_sb()`. So it can be used only in case of assurance, that pipefs superblock still exist. IOW, we can use this method only in pipefs mount-umount notification subscribers callbacks.

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

```
---  
include/linux/sunrpc/rpc_pipe_fs.h | 3 +++  
net/sunrpc/rpc_pipe.c               | 16 ++++++++  
2 files changed, 19 insertions(+), 0 deletions(-)
```

```
diff --git a/include/linux/sunrpc/rpc_pipe_fs.h b/include/linux/sunrpc/rpc_pipe_fs.h
```

```
index c1cdb2f..4a327ad 100644
```

```
--- a/include/linux/sunrpc/rpc_pipe_fs.h
```

```
+++ b/include/linux/sunrpc/rpc_pipe_fs.h
```

```
@@ -52,6 +52,9 @@ enum {
```

```
    RPC_PIPEFS_UMOUNT,
```

```
};
```

```
+extern struct dentry *rpc_d_lookup_sb(const struct super_block *sb,
```

```
+    const unsigned char *dir_name);
```

```
+
```

```
extern int rpc_queue_upcall(struct inode *, struct rpc_pipe_msg *);
```

```
struct rpc_clnt;
```

```
diff --git a/net/sunrpc/rpc_pipe.c b/net/sunrpc/rpc_pipe.c
```

```
index 8abeb9d..4860a56 100644
```

```
--- a/net/sunrpc/rpc_pipe.c
```

```
+++ b/net/sunrpc/rpc_pipe.c
```

```
@@ -1020,6 +1020,22 @@ static const struct rpc_filelist files[] = {
```

```
},
```

```

};

+/*
+ * This call can be used only in RPC pipefs mount notification hooks.
+ */
+struct dentry *rpc_d_lookup_sb(const struct super_block *sb,
+    const unsigned char *dir_name)
+{
+ struct qstr dir = {
+ .name = dir_name,
+ .len = strlen(dir_name),
+ .hash = full_name_hash(dir_name, strlen(dir_name)),
+ };
+
+ return d_lookup(sb->s_root, &dir);
+}
+EXPORT_SYMBOL_GPL(rpc_d_lookup_sb);
+
+static int
+rpc_fill_super(struct super_block *sb, void *data, int silent)
+{

```

Subject: [PATCH 5/7] SUNRPC: put pipefs superblock link on network namespace
 Posted by [Stanislav Kinsbursky](#) on Fri, 28 Oct 2011 14:28:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

We have modules (like, pNFS blocklayout module) which creates pipes on rpc_pipefs. Thus we need per-net operations for them. To make it possible we require appropriate super block. So we have to put sb link on network namespace context. Note, that it's not strongly required to create pipes in per-net operations. IOW, if pipefs wasn't mounted yet, that no sb link reference will present on network namespace and in this case we need just need to pass through pipe creation. Pipe dentry will be created during pipefs mount notification.

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

```

---
net/sunrpc/netns.h | 2 ++
net/sunrpc/rpc_pipe.c | 4 ++++
2 files changed, 6 insertions(+), 0 deletions(-)

```

```

diff --git a/net/sunrpc/netns.h b/net/sunrpc/netns.h
index d013bf2..b384252 100644
--- a/net/sunrpc/netns.h
+++ b/net/sunrpc/netns.h
@@ -9,6 +9,8 @@ struct cache_detail;
 struct sunrpc_net {

```

```

    struct proc_dir_entry *proc_net_rpc;
    struct cache_detail *ip_map_cache;
+
+ struct super_block *pipefs_sb;
};

extern int sunrpc_net_id;
diff --git a/net/sunrpc/rpc_pipe.c b/net/sunrpc/rpc_pipe.c
index 4860a56..5c313d3 100644
--- a/net/sunrpc/rpc_pipe.c
+++ b/net/sunrpc/rpc_pipe.c
@@ -1042,6 +1042,7 @@ rpc_fill_super(struct super_block *sb, void *data, int silent)
    struct inode *inode;
    struct dentry *root;
    struct net *net = data;
+ struct sunrpc_net *sn = net_generic(net, sunrpc_net_id);
    int err;

    sb->s_blocksize = PAGE_CACHE_SIZE;
@@ -1066,6 +1067,7 @@ rpc_fill_super(struct super_block *sb, void *data, int silent)
    if (err)
        goto err_depopulate;
    sb->s_fs_info = get_net(net);
+ sn->pipefs_sb = sb;
    return 0;

err_depopulate:
@@ -1086,6 +1088,9 @@ rpc_mount(struct file_system_type *fs_type,
void rpc_kill_sb(struct super_block *sb)
{
    struct net *net = sb->s_fs_info;
+ struct sunrpc_net *sn = net_generic(net, sunrpc_net_id);

+ sn->pipefs_sb = NULL;
    put_net(net);
    blocking_notifier_call_chain(&rpc_pipefs_notifier_list,
        RPC_PIPEFS_UMOUNT,

```

Subject: [PATCH 6/7] SUNRPC: pipefs per-net operations helper introduced
 Posted by [Stanislav Kinsbursky](#) on Fri, 28 Oct 2011 14:28:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

During per-net pipes creation and destruction we have to make sure, that pipefs sb exists for the whole creation/destruction cycle. This is done by using special mutex which controls pipefs sb reference on network namespace context. Helper consists of two parts: first of them (rpc_get_dentry_net) searches for dentry with specified name and returns with mutex taken on success. When pipe

creation or destructions is completed, caller should release this mutex by `rpc_put_dentry_net` call.

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

```
---
include/linux/sunrpc/rpc_pipe_fs.h | 4 ++++
net/sunrpc/netns.h                  | 1 +
net/sunrpc/rpc_pipe.c               | 36 ++++++++++++++++++++++++++++++++++++++
net/sunrpc/sunrpc_syms.c            | 1 +
4 files changed, 42 insertions(+), 0 deletions(-)

diff --git a/include/linux/sunrpc/rpc_pipe_fs.h b/include/linux/sunrpc/rpc_pipe_fs.h
index 4a327ad..271e1b2 100644
--- a/include/linux/sunrpc/rpc_pipe_fs.h
+++ b/include/linux/sunrpc/rpc_pipe_fs.h
@@ -55,6 +55,10 @@ enum {
extern struct dentry *rpc_d_lookup_sb(const struct super_block *sb,
    const unsigned char *dir_name);

+extern void rpc_pipefs_init_net(struct net *net);
+extern struct super_block *rpc_get_sb_net(const struct net *net);
+extern void rpc_put_sb_net(const struct net *net);
+
extern int rpc_queue_upcall(struct inode *, struct rpc_pipe_msg *);

struct rpc_clnt;
diff --git a/net/sunrpc/netns.h b/net/sunrpc/netns.h
index b384252..11d2f48 100644
--- a/net/sunrpc/netns.h
+++ b/net/sunrpc/netns.h
@@ -11,6 +11,7 @@ struct sunrpc_net {
    struct cache_detail *ip_map_cache;

    struct super_block *pipefs_sb;
+ struct mutex pipefs_sb_lock;
};

extern int sunrpc_net_id;
diff --git a/net/sunrpc/rpc_pipe.c b/net/sunrpc/rpc_pipe.c
index 5c313d3..cbf213e 100644
--- a/net/sunrpc/rpc_pipe.c
+++ b/net/sunrpc/rpc_pipe.c
@@ -1036,6 +1036,40 @@ struct dentry *rpc_d_lookup_sb(const struct super_block *sb,
}
EXPORT_SYMBOL_GPL(rpc_d_lookup_sb);

+void rpc_pipefs_init_net(struct net *net)

```

```

+{
+ struct sunrpc_net *sn = net_generic(net, sunrpc_net_id);
+
+ mutex_init(&sn->pipefs_sb_lock);
+}
+
+/*
+ * This call will be used for per network namespace operations calls.
+ * Note: Function will be returned with pipefs_sb_lock taken if superblock was
+ * found. This lock have to be released by rpc_put_sb_net() when all operations
+ * will be completed.
+ */
+struct super_block *rpc_get_sb_net(const struct net *net)
+{
+ struct sunrpc_net *sn = net_generic(net, sunrpc_net_id);
+
+ mutex_lock(&sn->pipefs_sb_lock);
+ if (sn->pipefs_sb)
+ return sn->pipefs_sb;
+ mutex_unlock(&sn->pipefs_sb_lock);
+ return NULL;
+}
+EXPORT_SYMBOL_GPL(rpc_get_sb_net);
+
+void rpc_put_sb_net(const struct net *net)
+{
+ struct sunrpc_net *sn = net_generic(net, sunrpc_net_id);
+
+ BUG_ON(sn->pipefs_sb == NULL);
+ mutex_unlock(&sn->pipefs_sb_lock);
+}
+EXPORT_SYMBOL_GPL(rpc_put_sb_net);
+
+static int
+rpc_fill_super(struct super_block *sb, void *data, int silent)
+{
@@ -1090,7 +1124,9 @@ void rpc_kill_sb(struct super_block *sb)
+ struct net *net = sb->s_fs_info;
+ struct sunrpc_net *sn = net_generic(net, sunrpc_net_id);
+
+ mutex_lock(&sn->pipefs_sb_lock);
+ sn->pipefs_sb = NULL;
+ mutex_unlock(&sn->pipefs_sb_lock);
+ put_net(net);
+ blocking_notifier_call_chain(&rpc_pipefs_notifier_list,
+ RPC_PIPEFS_UMOUNT,
+ diff --git a/net/sunrpc/sunrpc_syms.c b/net/sunrpc/sunrpc_syms.c
+ index 9d08091..880de8b 100644

```

```

--- a/net/sunrpc/sunrpc_syms.c
+++ b/net/sunrpc/sunrpc_syms.c
@@ -38,6 +38,7 @@ static __net_init int sunrpc_init_net(struct net *net)
    if (err)
        goto err_ipmap;

+ rpc_pipefs_init_net(net);
    return 0;

err_ipmap:

```

Subject: [PATCH 7/7] SUNRPC: added debug messages to RPC pipefs
 Posted by [Stanislav Kinsbursky](#) on Fri, 28 Oct 2011 14:29:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

This patch adds debug messages for notification events.

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

```

---
net/sunrpc/rpc_pipe.c | 8 ++++++++
1 files changed, 8 insertions(+), 0 deletions(-)

diff --git a/net/sunrpc/rpc_pipe.c b/net/sunrpc/rpc_pipe.c
index cbf213e..42e4b6e 100644
--- a/net/sunrpc/rpc_pipe.c
+++ b/net/sunrpc/rpc_pipe.c
@@ -33,6 +33,10 @@
#include "netns.h"
#include "sunrpc.h"

+#define RPCDBG_FACILITY RPCDBG_DEBUG
+
+#define NET_NAME(net) ((net == &init_net) ? "(init_net)" : "")
+
static struct vfsmount *rpc_mnt __read_mostly;
static int rpc_mount_count;

@@ -1095,6 +1099,8 @@ rpc_fill_super(struct super_block *sb, void *data, int silent)
    }
    if (rpc_populate(root, files, RPCAUTH_lockd, RPCAUTH_RootEOF, NULL))
        return -ENOMEM;
+ dprintk("RPC: sending pipefs MOUNT notification for net %p%s\n", net,
+ NET_NAME(net));
    err = blocking_notifier_call_chain(&rpc_pipefs_notifier_list,
        RPC_PIPEFS_MOUNT,
        sb);

```

```
@@ -1128,6 +1134,8 @@ void rpc_kill_sb(struct super_block *sb)
    sn->pipefs_sb = NULL;
    mutex_unlock(&sn->pipefs_sb_lock);
    put_net(net);
+ dprintk("RPC: sending pipefs UMOUNT notification for net %p%s\n", net,
+   NET_NAME(net));
    blocking_notifier_call_chain(&rpc_pipefs_notifier_list,
        RPC_PIPEFS_UMOUNT,
        sb);
```
