



SV1 Cluster User Data Base Feature

Richard Lagerstrom

Resource Administration



**41st Cray User Group
Conference
Minneapolis, Minnesota**

SV1 Cluster UDB Feature



Introduction

- UDB (User Data Base) contains user authorization, limits, quotas , etc.
- Each SV1 node has a UDB – same as stand-alone Unicos
- Existing tools are the core components of this feature

SV1 Cluster UDB Feature



Objective: To help maintain a consistent UDB on all nodes of an SV1 cluster

- Maintain UDB records from any node
- Designed in resiliency features
- Distribute user initiated password and shell changes
- Support local needs

SV1 Cluster UDB Feature

sgi

Support a configurable degree of consistency among the nodes

- Require fields such as home directory and password to be invariant
- Allow fields such as limits to vary from node to node
- Specify fields such as share usage node local

SV1 Cluster UDB Feature



Resiliency

- UDB maintenance may be done on any node – there is no “master node”
- nfs or other shared file support is required by this feature
- Pending changes will be applied when a node rejoins the cluster

SV1 Cluster UDB Feature



Implementation

- Minor changes to `udbgen` and `udbsee` to support node specific information
- No kernel or library changes were made
- Four new commands, a server and a module in which to place local support hooks have been added

SV1 Cluster UDB Feature



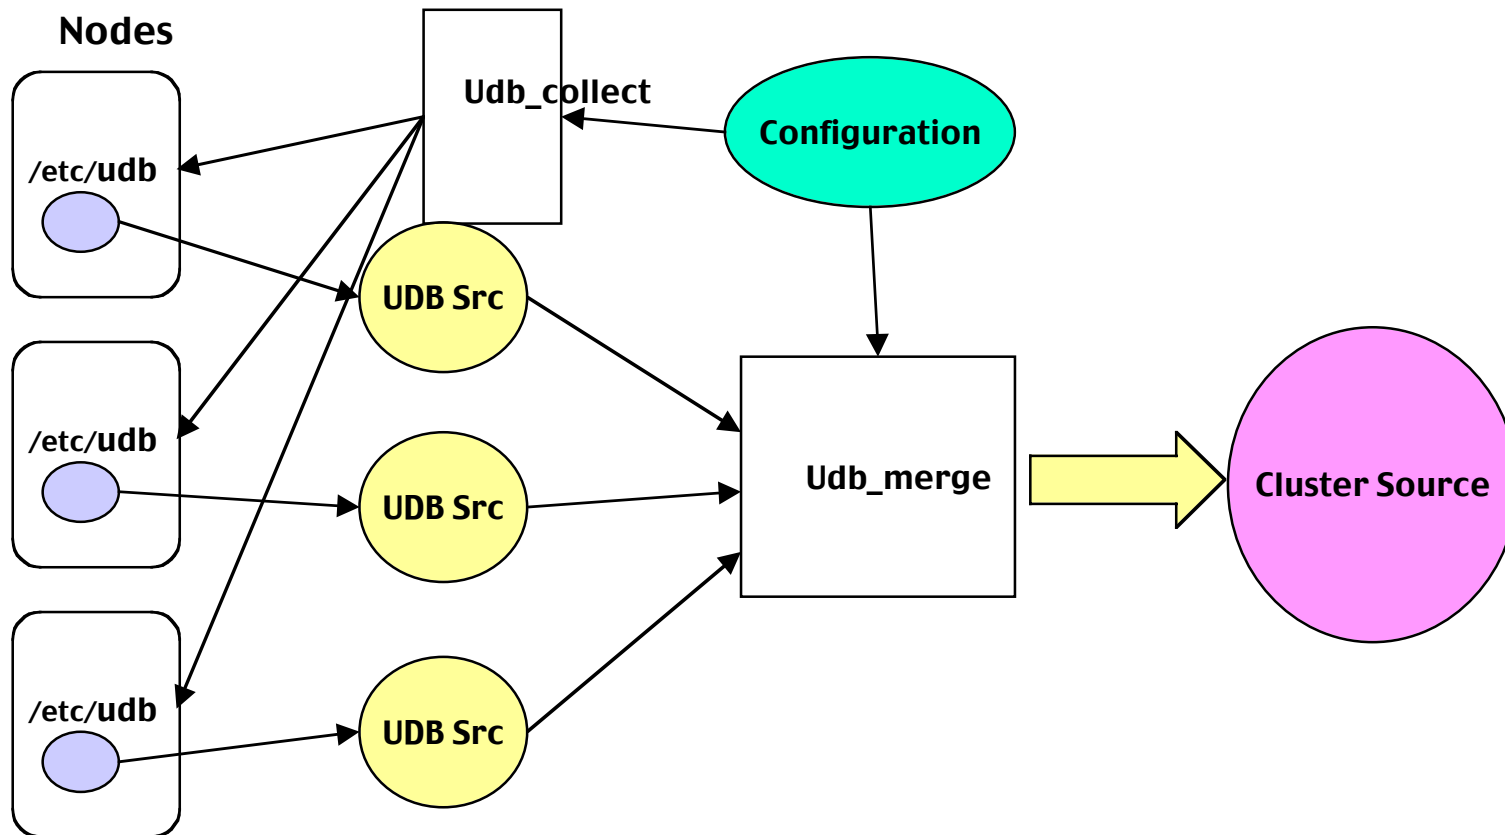
Security Issues

- File system security is used
- Only a limited set of commands are accessible to the server
- A common configuration file is highly recommended
- Cluster UDB commands are root-only

SV1 Cluster UDB Feature

sgi

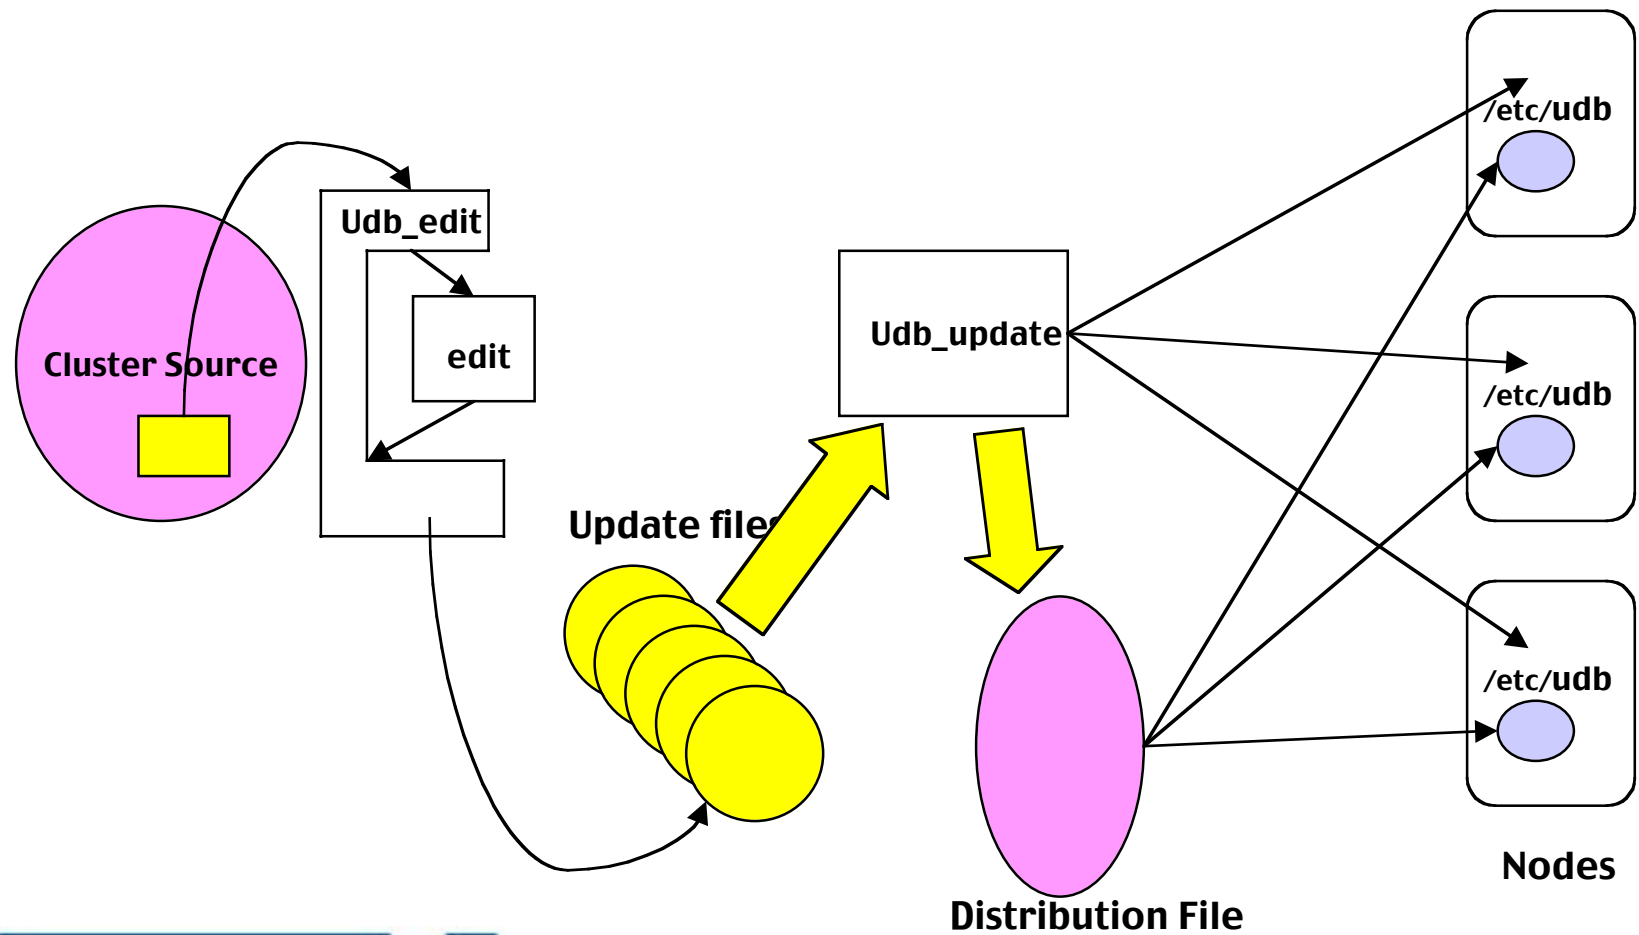
Getting Started



SV1 Cluster UDB Feature

sgi

Editing and Updating UDB Records



SV1 Cluster UDB Feature



Questions?

*Richard Lagerstrom
rnl@sgi.com*