

High-Availability on Linux

The Future

Lars Marowsky-Brée <Imb@suse.de>

Principal Engineer & Architect HA/Storage

Novell[®]

Agenda

- Introduction
- Problem statement
- History
- Plan for the future
- Present

Problem statement

- Independent Software Vendors,
- Independent Hardware Vendors,
- distributors,
- enterprise customers,
- all hate choice.

Choice means:

- Increased effort in
 - Design,
 - Development,
 - Testing,
 - Support,
 - Certification,
 - Training,
 - Documentation,
 - Evaluation,
 - Management.
- Good for exploring different directions.

History

- Linux-HA:
 - Heartbeat since 1999, multi-node since 2004
 - 2.x Cluster Resource Manager became Pacemaker
 - Primarily backed by SUSE/Novell/IBM
- Red Hat Cluster Suite:
 - Dates back to Kimberlite
 - rgmanager
- File systems & storage:
 - GFS{1,2} versus OCFS2
 - C-EVMS2 versus cLVM2

Cluster Developer Summit
Prague, 2008-09-29 – 2008-10-02

Dogs, cats, and mice ...



Topics discussed:

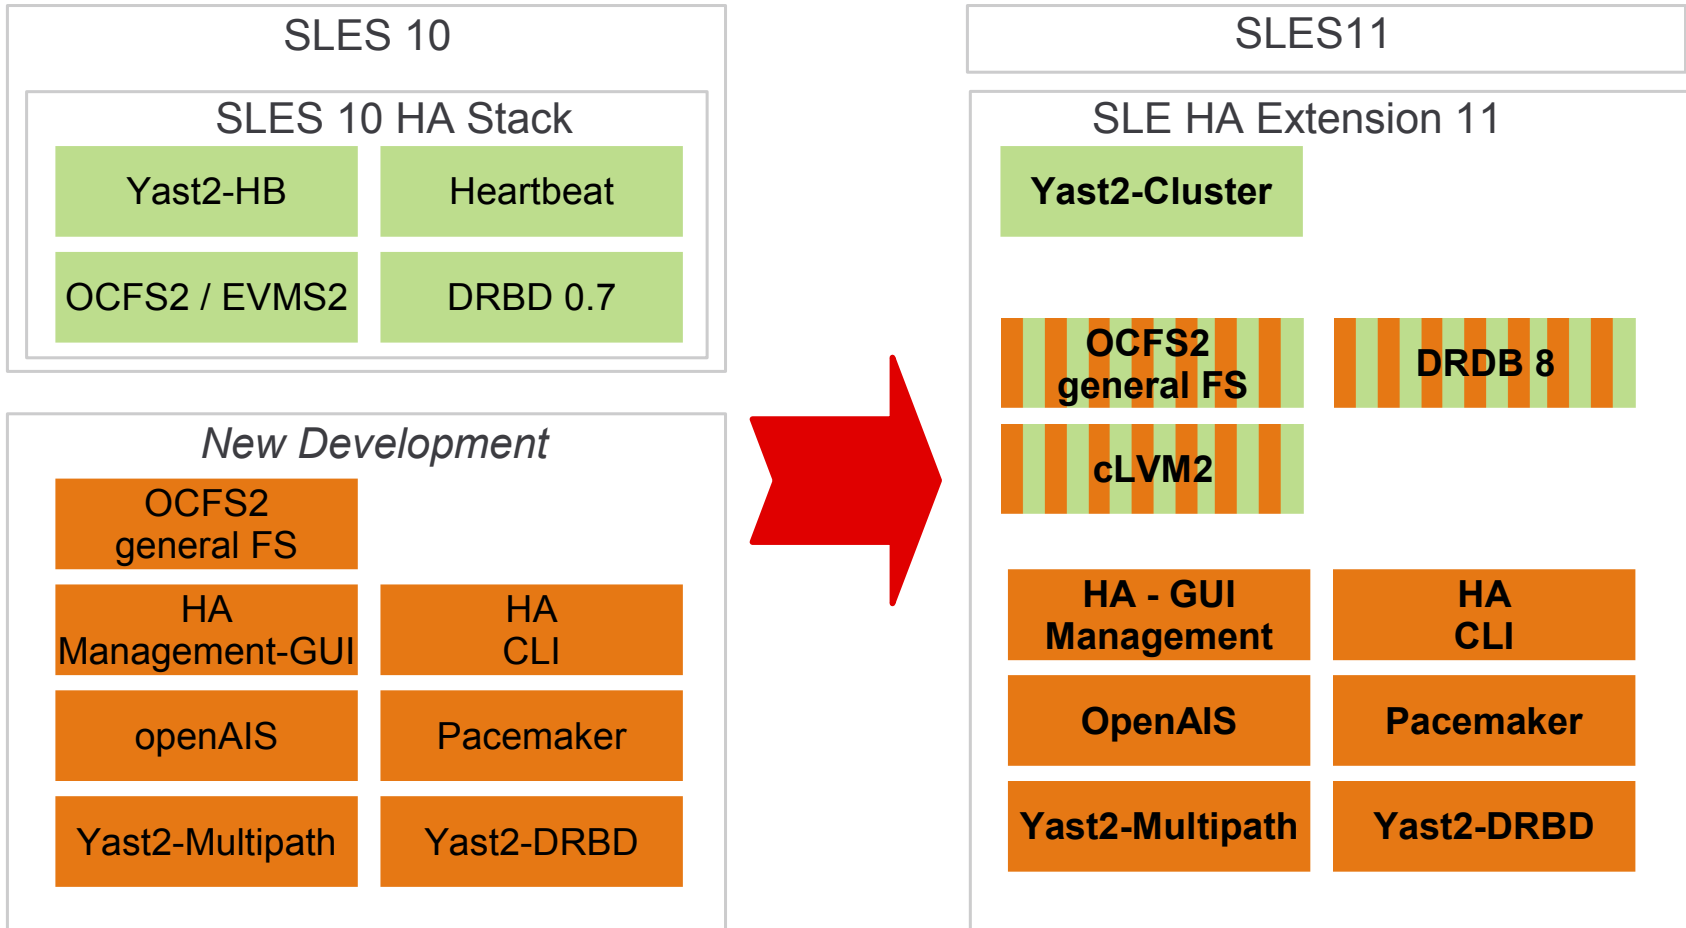
- Merging of stacks
 - Tree splitting
- Community release versus “Enterprise” distributions
- Quorum
- Fencing
- Resource Agents + common library
- Standard logging
- Standard configuration(?)
- Shared root clusters
- Clustered Samba

The future stack:

- ~~heartbeat~~ openAIS (corosync)
- Local Resource Manager
- ~~rgmanager~~ Pacemaker
- OCF Resource Agents
- Fencing component
- ~~C-EVMS2~~ cLVM2
- OCFS2
- GFS2

Present state

SUSE® Linux Enterprise High Availability Extension From 10 to 11



Part of SLES10

Major new code for SLES11

New in SLES11

SUSE® Linux Enterprise High Availability Extension Key Features (overview)

- High Availability clustering solution
 - Service availability 24h a day
- Oracle Cluster File System 2 (OCFS2)
 - Sharing and Scaling data-access by multiple nodes
- Distributed Remote Block Device 8 (DRBD)
 - Provide low-cost “SAN” through disk replication via TCP
- IP Load-balancing
 - Scale network services using multiple nodes
- User-friendly tools
- Documentation and whitepapers

Red Hat

- Pacemaker included with Fedora
- Tech-preview in Red Hat Enterprise Linux 6

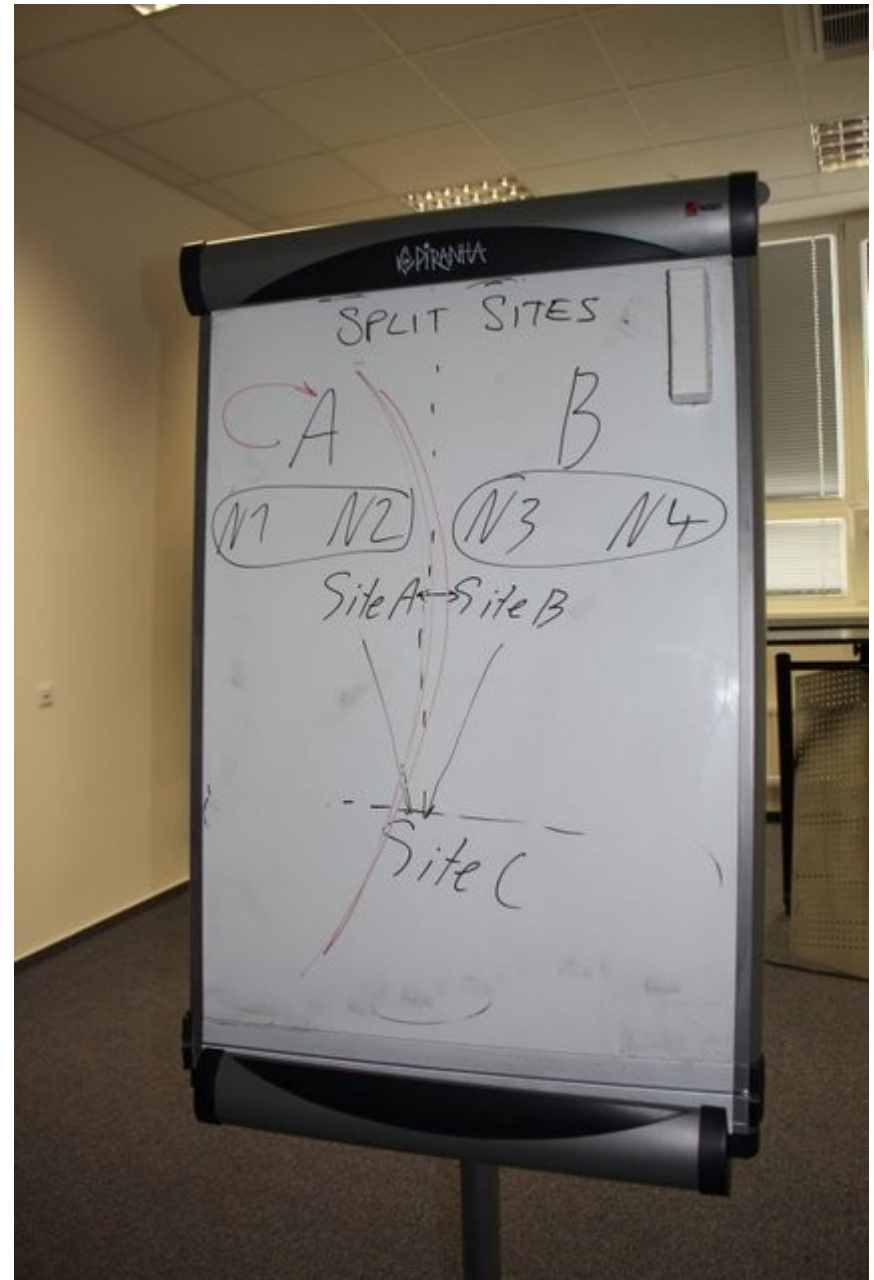
Future directions

Future issues

- Unifying fencing/STONITH
 - Split-site clustering
 - Virtualization integration
 - Improve OCF Resource Agent interfaces
 - Consolidate supporting libraries
-
- Next Cluster Summit probably collocated with LinuxCon/Linux Plumbers Conference 2010

Split-site clusters

- Wide-area:
 - High latency
 - Low bandwidth
 - Redundancy expensive
- Metro-area:
 - Low latency
 - High bandwidth
 - Redundancy expensive
- Quorum?
- Fencing?
- Storage?



Questions? Answers!

Novell®

Unpublished Work of Novell, Inc. All Rights Reserved.

This work is an unpublished work and contains confidential, proprietary, and trade secret information of Novell, Inc. Access to this work is restricted to Novell employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of Novell, Inc. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

General Disclaimer

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. Novell, Inc. makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for Novell products remains at the sole discretion of Novell. Further, Novell, Inc. reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All Novell marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

