

# SUSE<sup>®</sup> Linux Enterprise Server 11

What's new ?

**Dennis Deitermann**

Technology Specialist Data Center

<dd@novell.com>

**Novell<sup>®</sup>**

# SUSE® Linux Enterprise Server 11

## Key Features

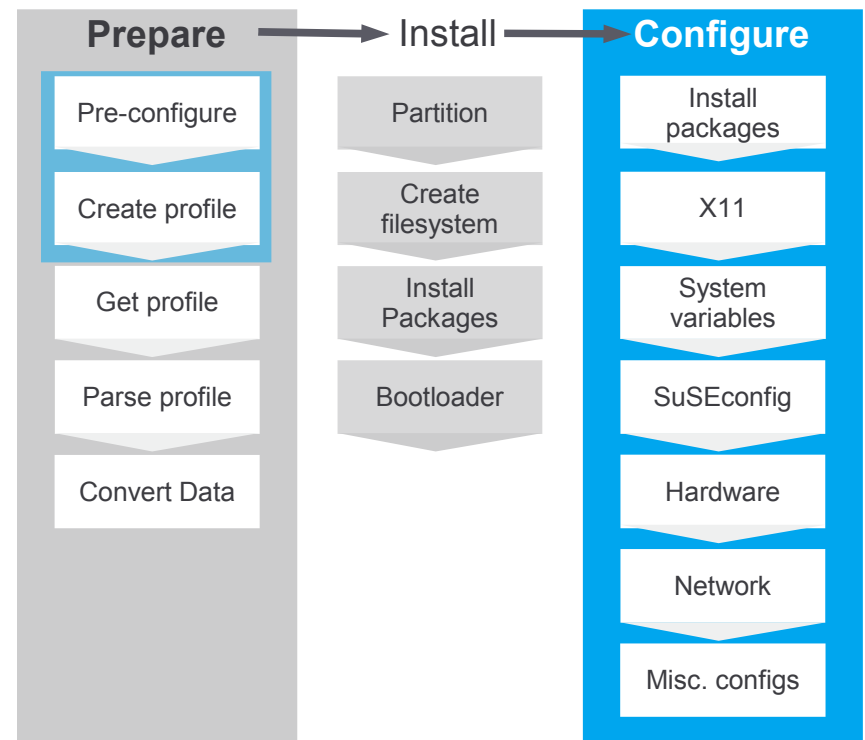
- **Integrated Systems Management**
  - Single server (YaST, Novell Customer Center)
  - Multiple server (AutoYaST, Subscription Management Tool, ZENworks® Linux Management, PlateSpin® Orchestrate)
  - Fast update stack (ZYpper)
  - Novell® Support Link, Support Advisor
- **Reliability, Availability, Serviceability**
  - Diskless server: Swap over NFS
  - Multipath I/O, Network channel bonding
  - Control groups, CPUset
- **Cross Platform Virtualization**
  - Integrated Xen hypervisor
  - Optimized as a guest OS for ESX, Hyper-V, and Xen
- **Interoperable with other Platforms**
  - Cross platform directory integration
  - Heterogeneous systems management (e.g. CIM, WS-Man)
  - POSIX compliance for interoperability with UNIX
- **Green IT**
  - Tickless idle
  - Granular power profiles
- **Comprehensive Security**
  - AppArmor application security
  - Trusted Platform Module (TPM) support
- **Modular OS and Functional Extensions**
  - SUSE Linux Enterprise High Availability Extension
  - SUSE Linux Enterprise Mono. Extension
  - SUSE Linux Enterprise Real Time Extension
- **Scalability and Flexibility**
  - Supports up to 4096 cores (Itanium) and 512 cores (x86\_64)
  - Supports up to 23 TB RAM
  - Runs on x86, x86\_64, Itanium, Power and Mainframe processor architectures
  - Supports diverse workloads (e.g. third party ISV, Java, .NET)
  - Optimized performance on physical and virtual systems

# Key Feature Details

The slide features a solid blue background. At the bottom, there are several horizontal, glowing white lines that create a sense of motion or depth, resembling a stylized horizon or a modern design element.

# Integrated Systems Management YaST and AutoYaST

- YaST installation and configuration management framework
  - Unified and consistent single interface to all systems management tasks via modules
  - Configure every aspect of the server
  - Enhanced YaST Partitioner
  - CIM Standard adopted
- AutoYaST auto-installation
  - A tool for installing SUSE® Linux Enterprise on systems with as much (or little) automation as you want
  - Works in networked and non-networked environments
  - Very flexible, very scalable, easy to use



# Integrated Systems Management Package Management Subsystem

- ZYpp: Core of the new software management stack
- ZYpper: Command line package manager
  - Can manipulate package repositories (like the Build Service)
  - Can search for packages, install, remove, or update packages and more
- libzypp: Software management library
  - Provides all the functionality for a package manager
    - > API for package repository management
    - > API for solving packages, products, patterns and patches (installation, removal, update and distribution upgrade operations) dependencies
    - > API for committing the transaction to the system over a rpm target
    - > API for browsing available and installed software
    - > A suite of maintained solving testcases



# Integrated Systems Management Common Information Model (CIM)

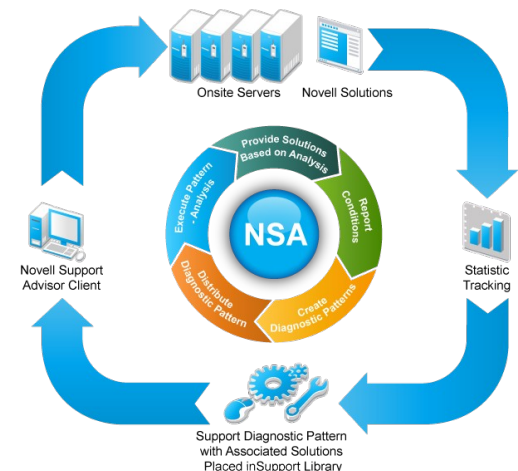
- Support for open Common Information Model management industry standards, basis for development of standards-based cross-platform management tools
- Key Components:
  - Small Footprint CIM Broker (SFCB)
  - Storage Management Provider (SMI-S providers for Volume Management and Snapshot, libstorage)
  - Virtualization Management Provider (libvirt)
  - Power Management Provider

## Technology Preview:

### Software Management Provider Package Kit

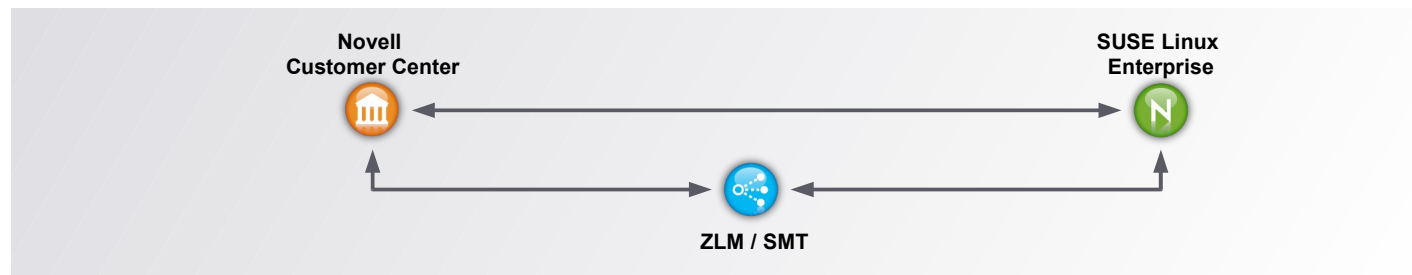
# Integrated Systems Management Supportability Infrastructure

- Novell® Support Link
  - Built-in into SLE, easily file a support incident with Novell Technical Services<sup>SM</sup>
  - YaST module enables quick reporting of an incident
  - Review any and all data before submitting, and discard any bit considered sensitive
  - Enables diagnostics performed by the Novell Support Advisor
- Enablement for Novell Support Advisor
  - System Discovery based upon IP address, IP range, or DNS name
  - SUSE® Linux Enterprise Server Health Diagnosis, with recommended solutions
  - Authentication support for Novell Customer Center
  - Automated Support Diagnostic Pattern deployment
  - Centralized Novell Technical Services Link

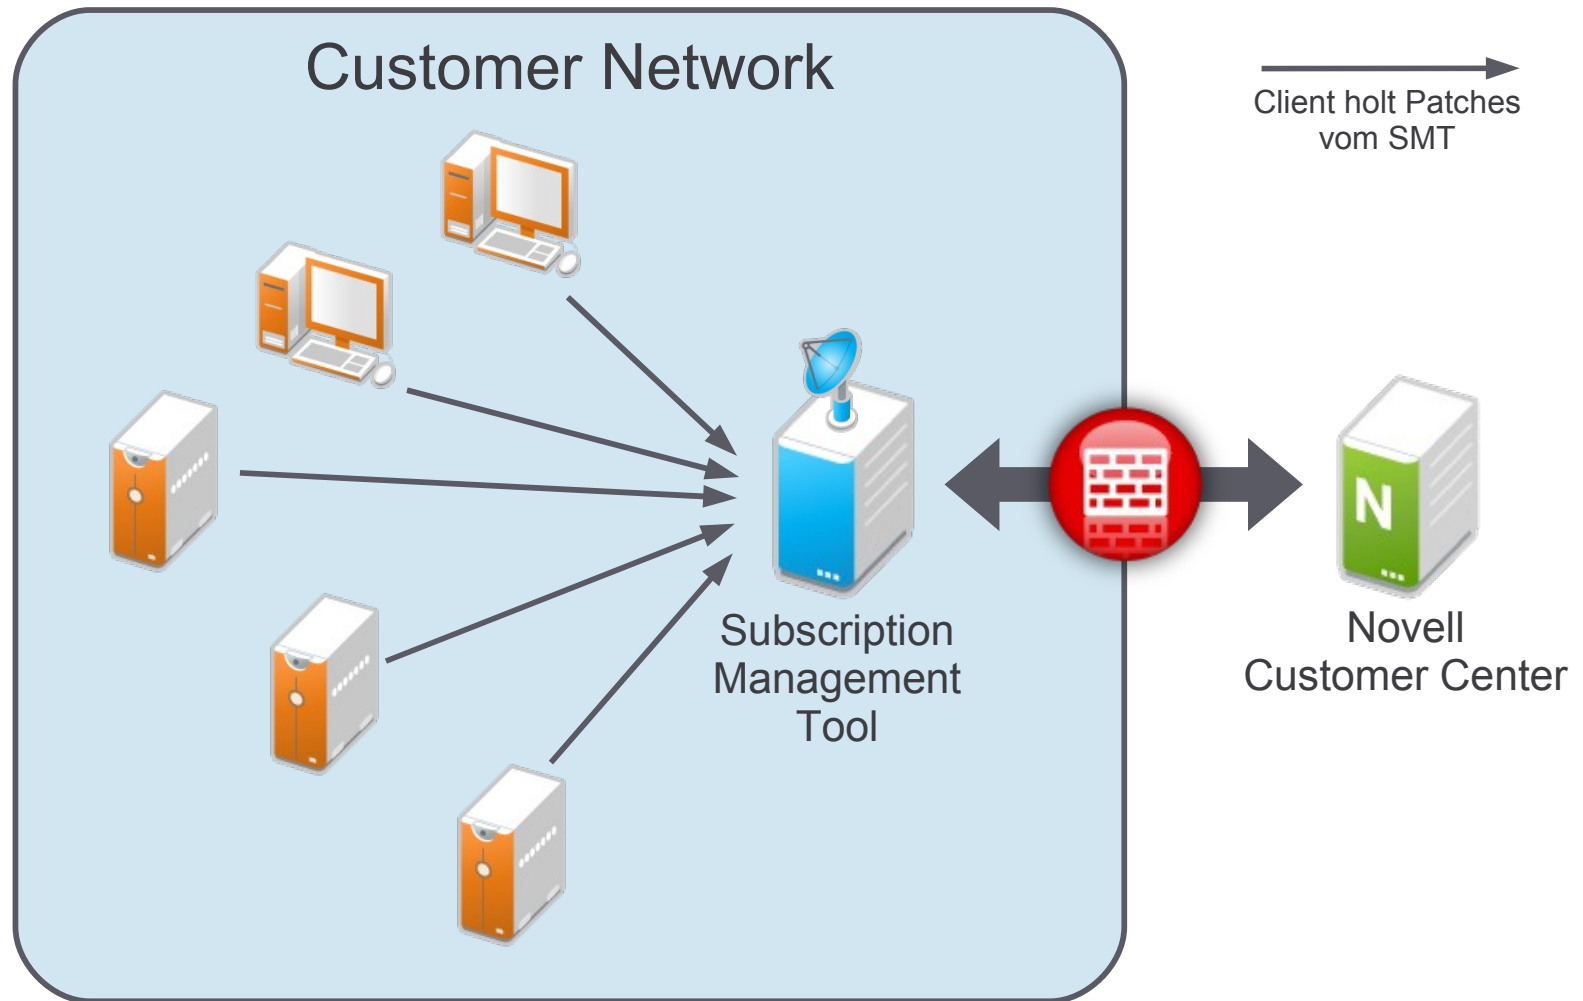


# Integrated Systems Management Subscription and Patch Management

- Novell® Customer Center
  - remains at the heart of the Novell relationship with the customers to manage SUSE® Linux Enterprise subscriptions
  - in cooperation with the SMT provides ability to automatically assign registration codes
- Subscription Management Tool for SUSE Linux Enterprise
  - is a local package proxy tightly integrated with the Novell® support infrastructure
  - is included at no cost in SLES entitlement
  - supports SLE 11 as a client at SLE 11 GA



# Integrated Systems Management Subscription Management Tool

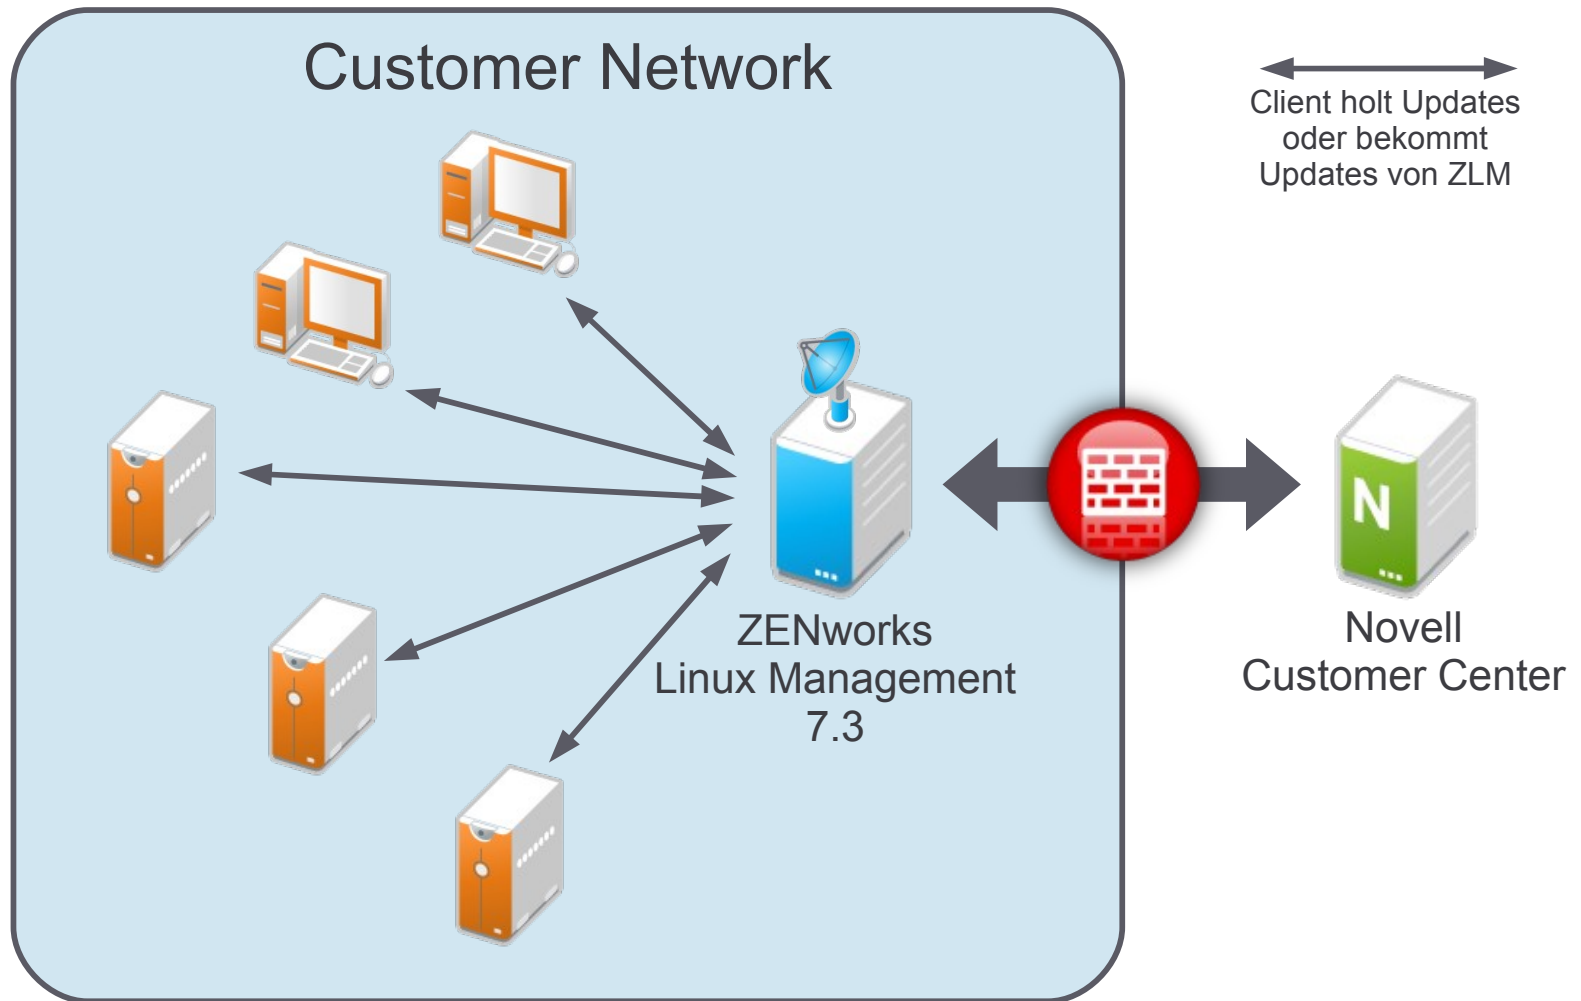


# Integrated Systems Management Enablement for ZLM

- ZENworks® Linux Management is the Novell® enterprise offering for Linux Management
  - Software management
  - Software, hardware inventory
  - Policy-driven management
  - Imaging engine
- ZENworks Linux Management 7.3 will support SLE 11 devices at SLE 11 GA
  - Easily network bootstrapped from any SLE install
  - Easily packaged in add-on format for roll-out



# ZENworks Linux Management 7.3



# Reliability, Availability, Serviceability File Systems

- **ext3** – default journaling file system in SUSE® Linux Enterprise 11
- **XFS** - high-performance journaling file system originally developed by SGI for use in its IRIX systems
- **ReiserFS** - alternative journaling file system

## Technology Preview (not supported):

- ext4
- eCryptfs Filesystem
- Read-only Root Filesystem

# Reliability, Availability, Serviceability Storage Management: LVM2

- Logical Volume Manager 2 for the Linux kernel
- Resize volume groups and logical volumes online
- Create read-write snapshots of logical volumes
- Stripe whole or parts of logical volumes across multiple PVs
- Mirror whole or parts of logical volumes
- Move online logical volumes between PVs
- Split / merge volume groups in situ

# Reliability, Availability, Serviceability Swap over NFS

- Protects systems against application restarts and downtime
- Usage scenarios:
  - If your stations do not have enough memory and do not have local drives, you may want to use swap over NFS
  - Cluster/blades, where having local disks is a cost issue (maintenance of failures, heat, etc)
  - Virtualization, where dumping the storage on a networked storage unit makes for trivial migration

# Cross-platform Virtualization

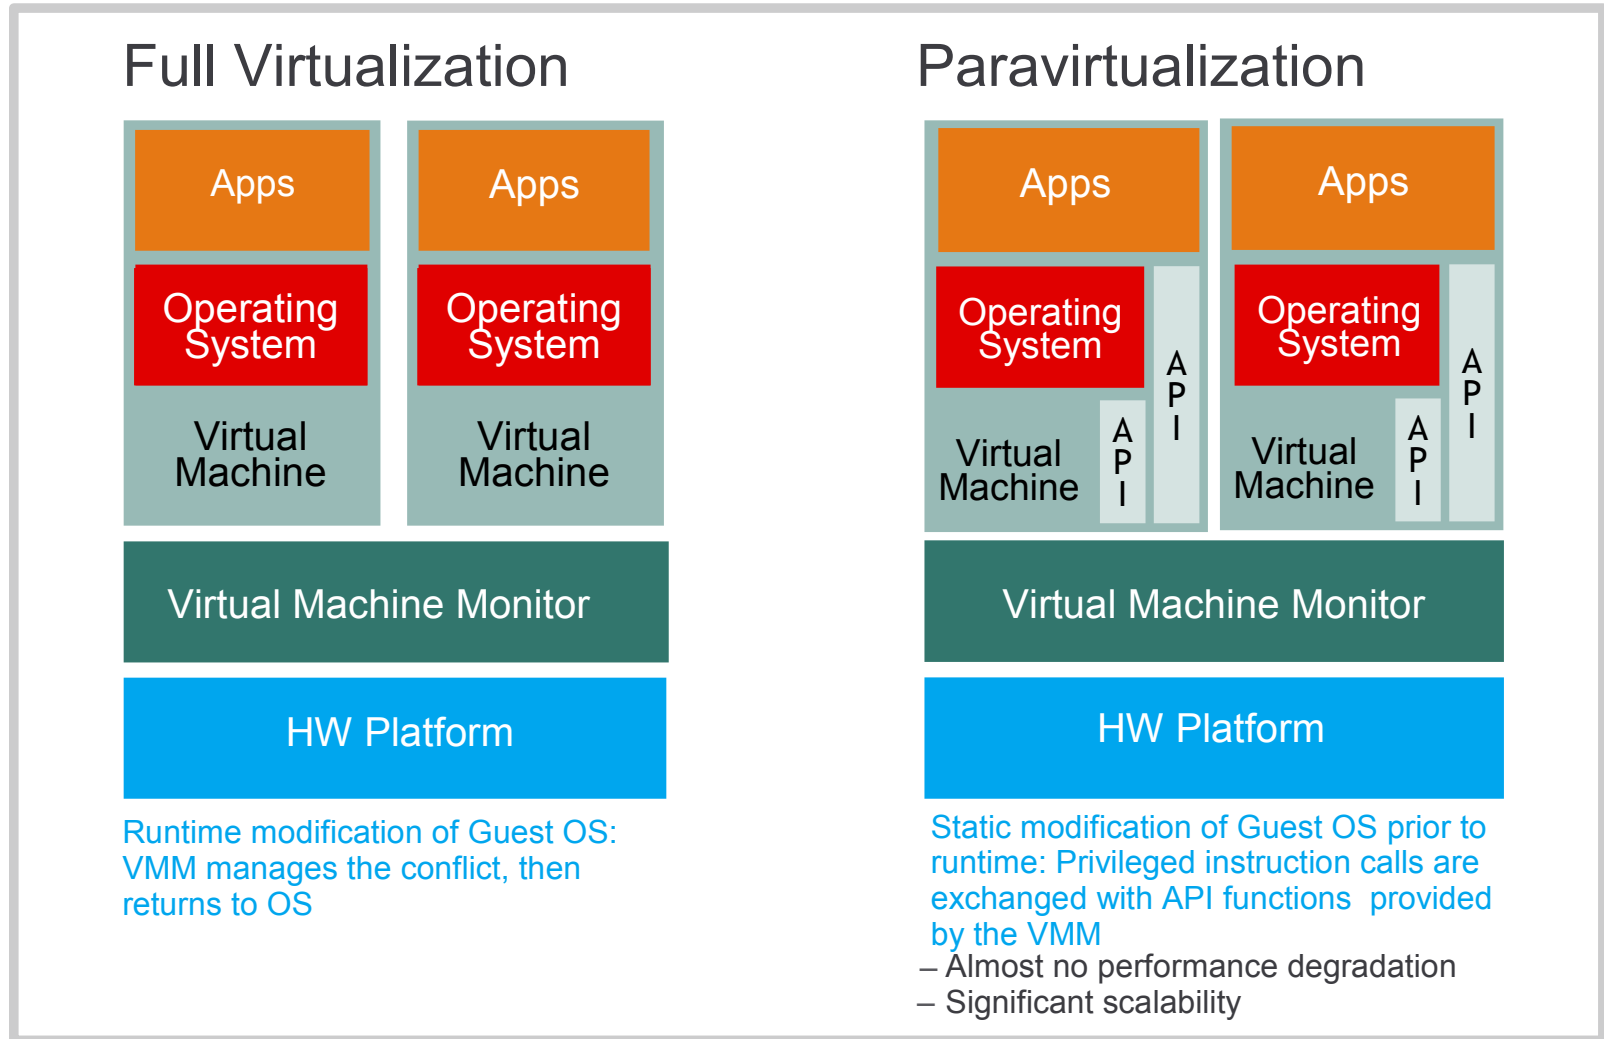
## Xen 3.3 Hypervisor-based Virtualization

- Hardware virtualization
  - Intel VT
  - AMD-V
- Integrated management tools
- Multiple guest operating systems
- No additional subscription fees for guest images
- Additional management solutions available

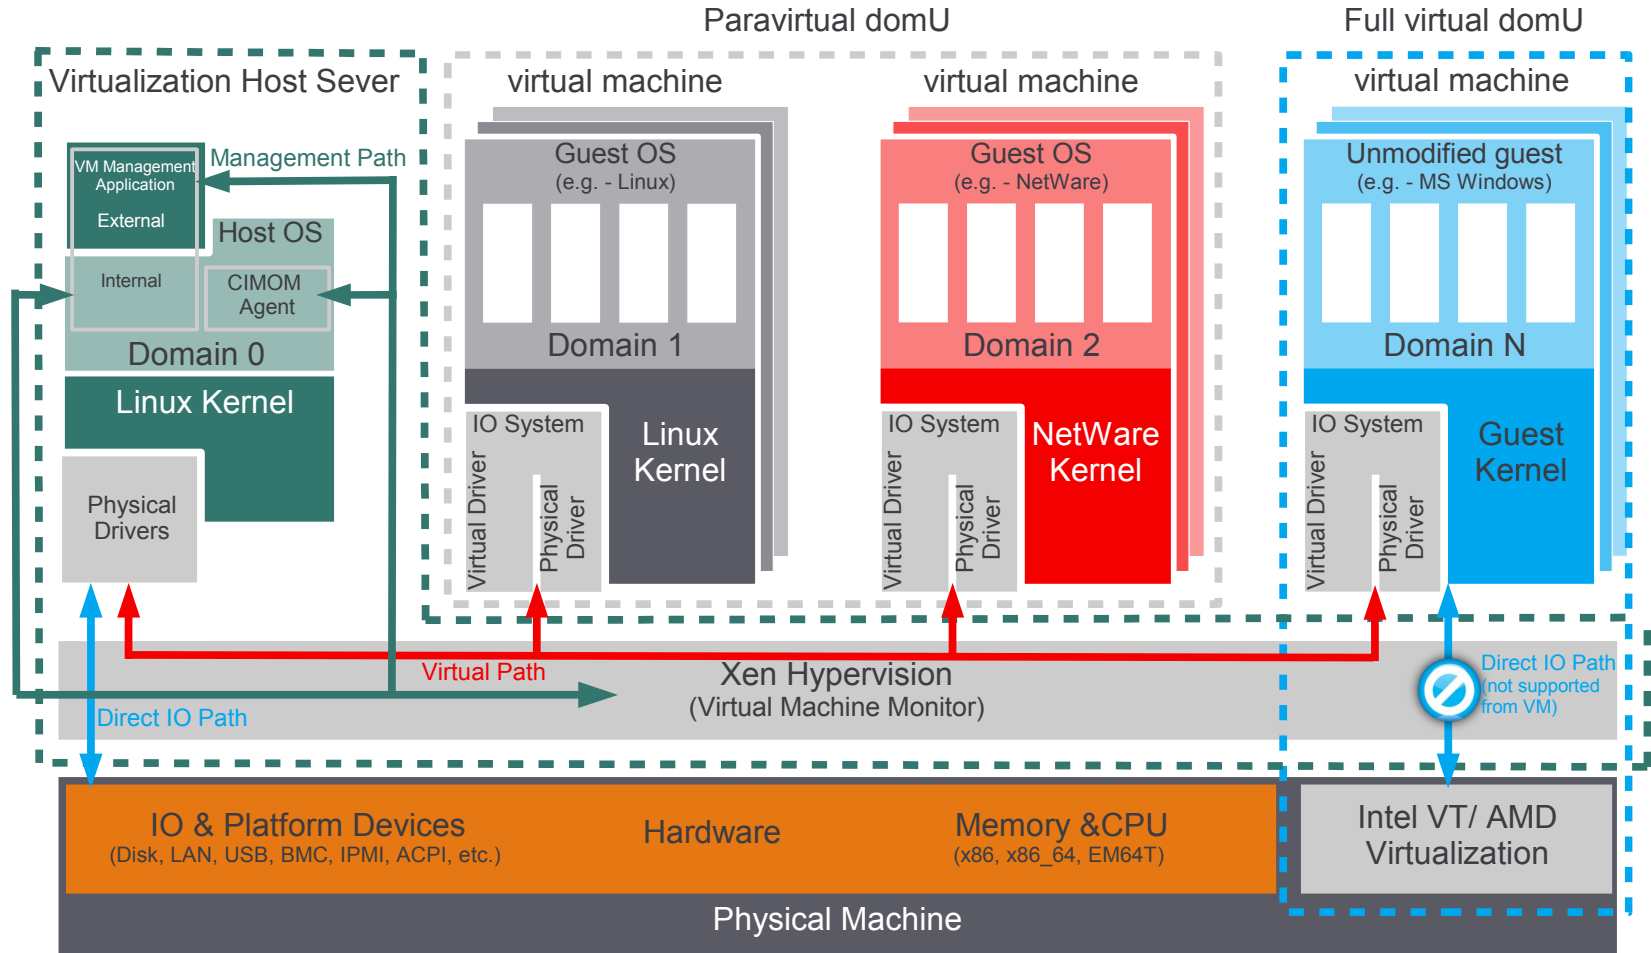


# Cross-platform Virtualization

## Full and Paravirtualization Overview



# Cross-platform Virtualization Virtual Machine Architecture



# Cross-platform Virtualization Virtual Machine Driver Pack

- Run your virtual services on unmodified Windows and Linux operating systems with near native performance
  - Run them on a wide variety of operating systems
  - Extend the useful life of your custom and legacy applications
  - Run them with the assurance of enterprise support from Novell.
- Device drivers improve communication between the Xen hypervisor and guest operating system, accelerating network and storage I/O
  - Drivers for Windows Server 2008, 2003, 2000 and XP on Xen
  - Drivers for Red Hat Enterprise Linux 4 and 5 on Xen

# XEN Features

- Save & Restore von laufenden Instanzen
- Dynamisches zuweisen / wegnehmen von Speicher und Virtuellen CPUs während der Laufzeit
- Migration von laufenden Gästen auf einen anderen Server
- Hotplugging von Block Devices & Netzwerkkarten
- Exklusives zuweisen von Hardware Komponenten an Virtuelle Instanzen (nicht im laufenden Betrieb)
- Fein zuteilbare Rechenleistung an einzelne Instanzen

# Cross-platform Virtualization Operating Systems Supported

Paravirtual Guest			Full Virtual Guest			
officially supported on SLES 11:	32 bit	64 bit	officially supported on SLES 11:	Paravirtual Drivers	32 bit	64 bit
SLES 10 SP1			SLES 9 and 10	yes		
SLES 10 SP2			SLES 11	yes		
SLES 11			OES2-Linux	yes		
OES2-NetWare. (SP7)			Windows 2000 Server	yes		
OES2-Linux			Windows XP	yes		
			Windows Server 2003	yes		
			Windows Server 2008 (enlightened)	yes		
			RHEL 4	yes		
			RHEL 5	yes		

\*Windows Paravirtual Drivers are closed source due to Microsoft DDK license restrictions

# Cross-platform Virtualization Virtual Machine Guest

- SUSE® Linux Enterprise Server is the “Perfect Guest” operating system for virtual machines
- Optimized for all three major hypervisors
  - VMware ESX
  - Microsoft Hyper-V
  - Xen
- Most attractive VM guest pricing: Unlimited virtual machine guests for the price of a single subscription
- Close partnerships with VMware, Microsoft and the Xen community



# Green IT

## Tickless Idle

- Prior versions of Linux kernel
  - Used a periodic timer for each CPU for e.g. process accounting, scheduler load balancing, and maintaining per-CPU timer events
  - Timer tick happened periodically whether processor was idle or busy
  - Idle processors had to wake up every 1, 4, or 10 milliseconds.
  - Lead to unnecessary power consumption in servers
- Kernel with "tickless idle"
  - Linux kernel has eliminated periodic timer tick when the CPU is idle
  - Allows the CPU to remain in power saving states for a longer period of time
  - Reducing the overall system power consumption.
- Benefits of tickless idle are lost if CPU is frequently awakened by unnecessary timer events

# Green IT

## More Features

- Processor C-states
  - Cut the clock signal and power from idle units inside the CPU
- Virtualization
  - Server Consolidation
  - Automatic Management of Resources
- Management of Power Profiles via
  - YaST modules
  - CIM



# Comprehensive Security

# AppArmor Security Technology

- Creates firewall around any Linux program (custom, open source, third party)
- Prevents the exploitation of unknown or undiscovered application vulnerabilities
- Easy to use GUI tools with static analysis and learning-based profile development
- Default policies included
- Create custom policy in hours, not days



# Comprehensive Security Capabilities More Features

- Integrated VPN
  - Integrated Firewall
  - Trusted Platform Module
  - Anti-virus protection
  - Certificate management
  - Secure network connections (ssh, OpenSSL, SASL)
  - Linux audit framework
  - YaST: Easiest security configuration and management
- 



# Reliability, Availability, Serviceability CPU Management and System Activity

- Control Groups
- CPUset System (cpuset)
  - Full CPU administration capabilities for Linux kernel
  - Based on Control Groups
- Cpuset
  - Command-line tool for using cpuset facilities
- System Activity: sysstat package
  - sar: collects and reports system activity information
  - iostat: reports CPU and I/O statistics
- Interrupt request balance
  - Switches the target CPUs for interrupt requests, preventing CPU0 from being used for all interrupt requests

# Reliability, Availability, Serviceability Crash Dump, Debug, Trace, Monitor

- Crash dumping analysis and debugging
  - “crash”, kdump, ltrace, GDB
- Faster reboots
  - e.g. using kexec (~Linux kernel warm-restart)
- System Tap
  - One tool to analyze systemic problems all the way from applications to operating system
- Nagios
  - Open Source host, service and network monitoring program

## **Technology Preview (not supported): *PerfMon2***

- Monitors performance on unmodified or for an entire system

# High Performance Computing Advanced Input/Output Processing

- Asynchronous I/O (AIO)
  - A form of input/output processing that permits other processing to continue before the transmission has finished
- Modular I/O Scheduler
  - Algorithm most suitable for workload can be chosen dynamically
- Multi-core/hyper-threading processor support
  - Execute threads in parallel within each individual processor
- Intel I/O Acceleration
  - Offloads the CPU towards the network card, thus allowing the system to continue processing data while I/O is taking place



# High Performance Computing Flexible and Pluggable I/O Scheduler

- CFQ
  - Complete Fair Queuing, suitable for a wide variety of applications, default I/O scheduler
- Deadline I/O scheduler
  - A valuable choice on a server with intelligent hardware, or if latency is the primary goal
- Anticipatory I/O scheduler
  - Suitable for file servers but does not work as well for database workloads
- No-op
  - Better suited for systems with storage solutions that have their own cache and their own virtualization Modular I/O Schedule


# High Performance Computing More Features

- Enhanced and tuned 2.6 series Linux Kernel for optimal performance
- CPU Management and System Activity
  - Cpuset System, Cpuset command line tool, Sysstat package, IRQbalance
- Non-Uniform Memory Access or Non-Uniform Memory Architecture (NUMA)
  - A computer memory design used in multiprocessors, where the memory access time depends on the memory location relative to a processor
- OpenFabrics Enterprise Distribution (OFED)
  - Remote Direct Memory Access (RDMA) switched fabric technologies, high-speed data transport technologies for server and storage connectivity
- SystemTap

# Advanced Networking

## FCoE and DCB

- Fibre Channel over Ethernet (FCoE)
  - Mapping of Fibre Channel frames over full duplex IEEE 802.3 networks
  - Leverage 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol
- Data Center Bridging (DCB)
  - Enhanced Transmission Selection (Priority Grouping)
  - Priority-based Flow Control (PFC)
  - Congestion Notification



# SUSE Linux Enterprise High Availability Extension

# SUSE® Linux Enterprise High Availability Extension

- An integrated suite of robust open source clustering technologies
  - Affordable
  - Integrated
  - Virtualization-aware
- Used with SUSE Linux Enterprise Server, it helps you
  - Maintain business continuity
  - Protect data integrity
  - Reduce unplanned downtime for your mission critical Linux workloads

# SUSE® Linux Enterprise High Availability Extension Pricing



## • Pricing

-x86 and x86\_64

- > USD 699 per year per server
- > Support level inherited by base SUSE Linux Enterprise Server
- Power, Itanium, System z
  - > Bundled with base SUSE Linux Enterprise Server at no additional charge
  - > Support level inherited by base SUSE Linux Enterprise Server

# SUSE® Linux Enterprise High Availability Extension Promotion



## • Existing Customers

### -Free of charge subscription

- > For all valid SUSE Linux Enterprise Server subscriptions
- > Effective date: June 1<sup>st</sup> 2009
- > Valid for subsequent subscription periods if base SUSE Linux Enterprise Server is renewed on time

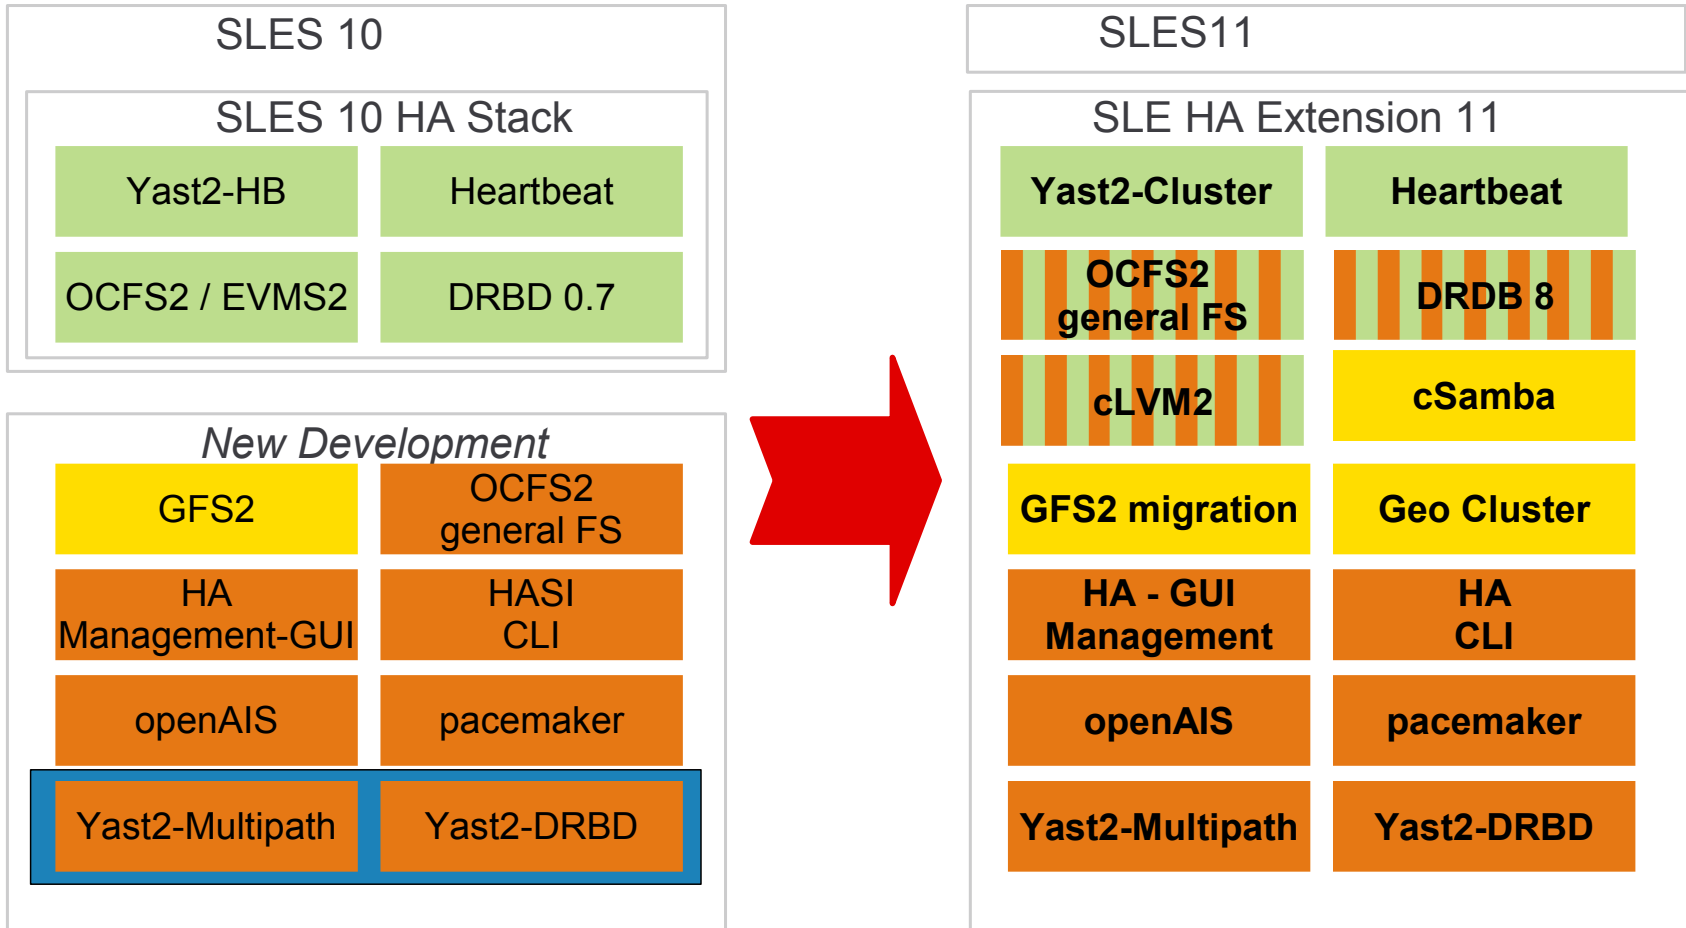
# Packaging: Add-on Product

- Oracle Cluster File System 2
  - Kernel modules, user land packages
- DRBD8
  - Kernel modules, user land packages, YaST module
- CLVM2
  - Kernel modules, cLVM daemon, user land packages, not affecting the non-clustered use of LVM2 in SLES
- Linux-HA stack
  - incl. heartbeat, openAIS, pacemaker, HA-GUIs and directly related packages, YaST modules
- Linux Virtual Server

# SUSE® Linux Enterprise High Availability Extension Clustering Scenarios

- Active / Active
- Active / Passive
  - N + 1
  - N + M
  - N to 1
  - N to M
- Hybrid physical and virtual

# SUSE Linux Enterprise High Availability Extension



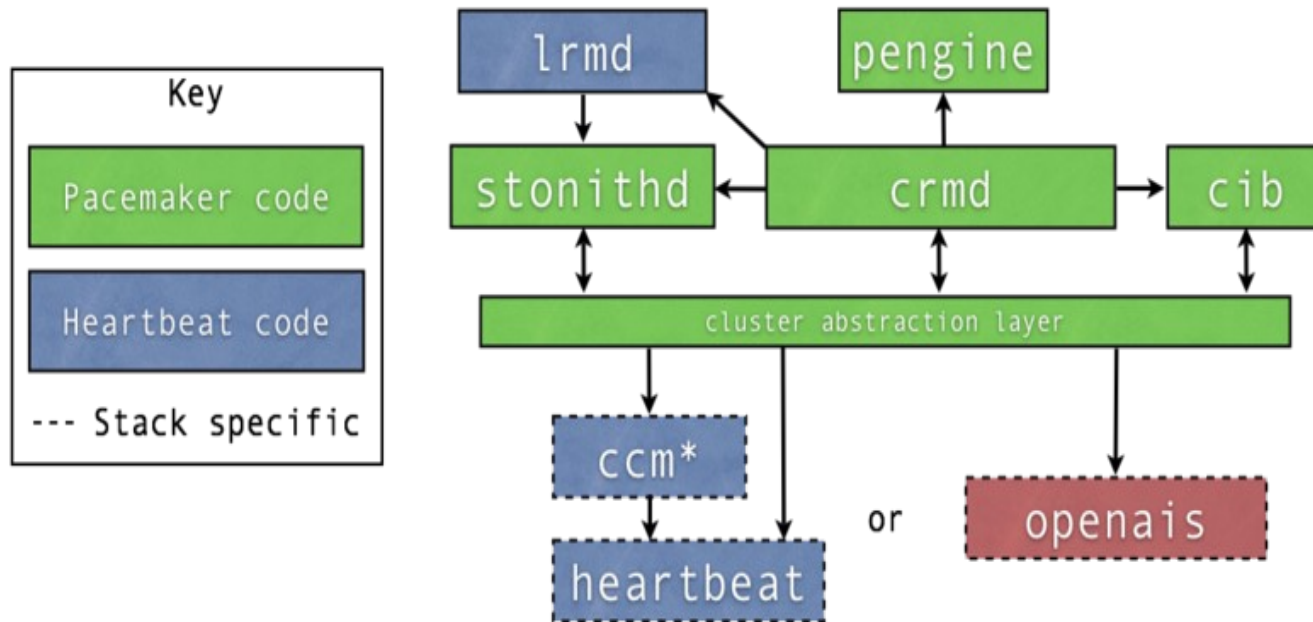
Part of SLES10

New for SLES11

Major new code for SLES11

Planned for SLES11 SPx

# Pacemaker Cluster Stack



- **Stonithd** The Heartbeat fencing subsystem.
- **lrm** Local Resource Management Daemon, non-cluster aware, presents a common interface to the supported resource types. Interacts directly with resource agents (scripts).
- **Pengine** Policy Engine. Computes the next state of the cluster based on the current state and the configuration. Produces a transition graph contained a list of actions and dependencies.
- **cib** Cluster Information Base. Contains definitions of all cluster options, nodes, resources, their relationships to one another and current status. Synchronizes updates to all cluster nodes.
- **crmd** Short for Cluster Resource Management Daemon. Largely a message broker for the PEngine and LRM, it also elects a leader to co-ordinate the activities (including starting/stopping resources) of the cluster.
- **ccm** Consensus Cluster Membership. The Heartbeat membership layer.
- **openAIS** The OpenAIS messaging and membership layer.
- **Heartbeat** The Heartbeat messaging layer, an alternative to OpenAIS.

**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.

