RED HAT® ENTERPRISE LINUX®

redhat.

RED HAT ENTERPRISE LINUX FEATURE BRIEF

Application-optimized infrastructure with containers

TECHNOLOGY BRIEF

Four key elements that an operating system should implement that are also addressed by Linux containers:

- Resource management
- Process isolation
- Security
- Tooling

Linux containers combine the flexibility of image-based deployment with lightweight application isolation. Developers choose Linux containers because they simplify application deployment. And many Platform-as-a-Service (PaaS) architectures are built around Linux container technology—including OpenShift by Red Hat.

Red Hat® Enterprise Linux® 7 beta implements Linux containers using core technologies like control groups (cgroups) for resource management, namespaces for process isolation, and Security-Enhanced Linux (SELinux) for security. This allows secure multi-tenancy and reduces the potential for security exploits.

For managing containers in Red Hat Enterprise Linux 7 beta, Docker provides a native toolkit for manipulating core system capabilities such as:

- cgroups
- namespaces
- network interfaces
- firewalls
- kernel features

Red Hat container certification ensures that application containers built using Red Hat Enterprise Linux will operate seamlessly across certified container hosts.

FEATURES AND CAPABILITIES

RESOURCE MANAGEMENT

Resource management for containers is based on cgroups, which allow users to allocate CPU time, system memory, network bandwidth, block IO, or any combination of these resources to a set of user-defined task groups or processes running on a given system. Users can then monitor any cgroups they configure, deny cgroups access to certain resources, and even dynamically reconfigure cgroups on a running system. Cgroups give system administrators fine-grained control over allocating, prioritizing, denying, managing, and monitoring system resources. Hardware resources can be divided among tasks and users, often increasing overall system efficiency.

PROCESS ISOLATION

Process isolation, the core of Linux container architecture, is provided by kernel namespaces. Red Hat Enterprise Linux implements five types of namespaces, each wrapping a particular global system resource in an abstraction, as described in Table 1. This makes each specific resource appear as an isolated instance to the processes within the namespace. This creates the illusion that this group of processes is alone on the system.

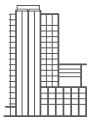

TABLE 1:
NAMESPACES PROVIDING PROCESS ISOLATION

NAMESPACE	DESCRIPTION
PID	Allows processes in different PID namespaces to have the same PID. Containers can be migrated between hosts while keeping the same PIDs for the processes inside the container. This also allows each container to have its own init process that manages various system initialization tasks and the container life cycle.
Network	Provides isolation of network controllers and system resources associated with networking, firewall, and routing tables. Allows each container to have its own virtual network stack associated with a process group.
UTS	Allows each container to have its own hostname and NIS domain name, which is useful for initialization and configuration scripts that tailor actions based on these names.
Mount	Isolates the set of filesystem mount points seen by a group of processes. This facilitates the creation of different read-only filesystems, so that processes in different mount namespaces can have different views of the filesystem hierarchy.
IPC	Isolates certain interprocess communication (IPC) resources, such as System V IPC objects and POSIX message queues.

SECURITY

Security for Linux containers is implemented using SELinux. SELinux applies security labels and policies to Linux containers and their resources, providing an additional layer of security above and beyond the isolation provided by kernel namespaces.

In contrast to the standard virtualization scenario, where different virtual machines share the same physical host, Linux containers can provide secure application isolation and multi-tenancy within a given cloud workload or virtual machine.

For example, consider a hosted customer relationship management (CRM) application. You don't want to stand up a separate virtual machine for each user, but still want to provide full isolation for all the different instances of the application.

ABOUT RED HAT

Red Hat is the world's leading provider of open source solutions, using a community-powered approach to provide reliable and high-performing cloud, virtualization, storage, Linux, and middleware technologies. Red Hat also offers award-winning support, training, and consulting services. Red Hat is an S&P company with more than 70 offices spanning the globe, empowering its customers' businesses.

NORTH AMERICA
1 888 REDHAT1

EUROPE, MIDDLE EAST, AND AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com

facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

redhat.com
#12077347_v2_0414