Content Router™

Complete, automated data management and distribution

Intelligently and securely automate the distribution of data from Caringo® CAStor®

IT organizations are having difficulty complying with data protection and disaster recovery regulations due to the increasing size of data sets and shrinking budgets. In addition the move to an all-digital workflow is driving the need to keep data online and accessible for years or even decades.

To effectively solve these issues we developed Content Router (CR) – an object routing engine integrated into our object storage platform that intelligently and securely automates the distribution of data between your Caringo® CAStor® clusters and can facilitate specialized operations within a cluster.

Granular control driven by customizable metadata

CR leverages user-defined metadata per object to enables sophisticated content distribution, content processing and data protection policies based on your business.

Content Router Benefits

Automated data protection – No manual migration, no provisioning and no storage management required.

Optimized access – Intelligent and secure routing of content closer to the end user according to your business needs.

Flexibility – CR is a highly versatile replication product that delivers:

- Remote replication over any distance to were the data is needed for increased protection or access.
- Disaster recovery with rules based queries for restoration down to the object level.
- Real-time access to data within a local cluster or at a remote cluster based on load balancing and availability of an active/active clustered configuration.

Unlimited scale-out – CR enables storage clusters to grow from a single site to distributed clusters across many sites with easy web based administration.

Enables high availability – Distributed clusters containing replicated data can be used for fail over requirements.

Easy integration with value added applications – CR is built on a Publisher/Subscriber model. The Subscriber API enables integration of custom services such as compression, encryption, virus scanning, indexing and search.

Content Router™

Complete, automated data management and distribution

Recommended Infrastructure

The following networking services are recommended for the management of a CAStor cluster that includes Content Router.

Syslog server for receiving critical alerts

NTP time server to provide clock synchronization

Gigabit Ethernet is the recommended connection speed between Content Router servers and CAStor cluster nodes. A Content Router node should use the same speed connection as the fastest CAStor node it communicates with to prevent bottlenecks.

Operating System

Content Router runs on Red Hat Enterprise Linux 5.5 or 6.0.

Hardware Configuration

CAStor Content Router will run on a variety of x86 hardware. The following are recommended characteristics of a Content Router server.

CPU: 64-bit, Intel Xeon, AMD Athlon64, or equivalent

Hard Drive: Standard SATA, SAS, or SCSI with RAID configuration for redundancy

Other Drive: CD ROM to install the operating system

Memory: Publisher Service – 2GB recommended. Replicator Service – 512MB recommended

Network Interface: Dual Gigabit Ethernet

Storage nodes: CAStor™ Storage Cluster Storage cluster of 3 or more server nodes.

About Caringo

To learn more about Caringo and our products visit our website or e-mail info@caringo.com.

Caringo, Inc.

6801 North Capital of Texas Highway
Building 2, Suite 200
Austin, Texas 78731
512-782-4490
www.caringo.com

CR Features

Fully Integrated with CAStor® – an add-on to CAStor storage clusters enabling a policy based object routing engine for data distribution

Metadata driven object routing engine – using metadata rules and policies for intelligent content distribution and processing

Network management – bandwidth optimization and tolerance for high latency and intermittent connectivity during replication over wide area networks (WANs)

Subscriber API – snap in third-party content processing applications such as indexing and search engine, and virus scanning

Virtual Machine support – can seamlessly run in a virtual machine for optimization and flexibility in the cluster

Active/Active support – Enables replication of content to another location for failover scenarios

Red Hat Enterprise Linux support – runs on any x86 platforms running Red Hat 5.5 and 6.0

Multiple replication and distribution topology support – Provides high availability options of safeguarding content by mirroring data between clusters in one or many distributed locations. A wide range of flexible deployment models can be configured to meet DR and business continuity SLAs – (1:1), (1:M), (M:1), (M:M)

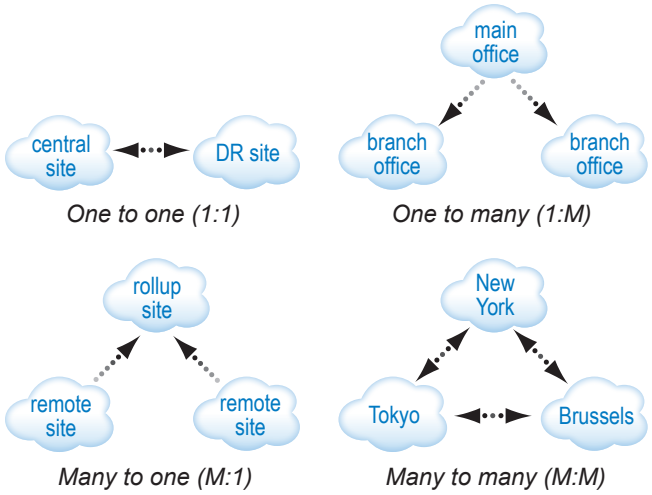

1:1 – DR solution for central site to remote site

1:M – Business continuity solution with bunker topology addressing geographic disasters

M:1 – DR solution with spoke and hub topology addressing data protection for multiple sites to a central site

M:M – Business continuity solution with replication to a central site and a separate long distance DR site

caringo