R-1
Deployment Platform

RepliWeb

RepliWeb R-1 & Real-Time Continuous Update Add-on

1. Overview

RepliWeb Content Deployment Platform (R-1) is a cost effective, feature-rich web content and application deployment automation platform trusted by leading Global 2000 companies. Hundreds of enterprises and government organizations put R-1 through the paces to meet their mission-critical deployment needs. Ensuring local and distributed servers, as well as complete datacenters, are synchronized with the most up-to-date content and code R-1 is the single source solution for organizations of any size.

Server administrators trust R-1 with all the critical technical details to automate website publishing: Files, web content, IIS settings, .NET assemblies, and COM+ components can roll out to one, dozens, or hundreds of physical or virtual servers. Updates occur behind the scenes, with no service interruptions, down-time, or impact on end users' browsing experiences. Even large enterprises enjoy significant efficiencies, as R-1 can scale to handle infinitely large content repositories without compromising performance or stability. R-1's cross platform support, open architecture, and extensive configuration capabilities allow for seamless integration into existing content-based file workflows and content lifecycles.

Organizations in need of real-time content deployment, distribution, or synchronization can employ an I/O-level or file-level solution. The former system is based on capturing disk I/O in real-time and replicating the changes to a remote-server in the same sequence they occurred on the source. The latter is based on capturing events on the file level, and subsequently replicating the changes to a remote server. While I/O level replication is necessary in database replication, it can introduce significant issues when employed to vast content depositories (web sites, corporate data, etc.)

Enterprise projects that do not require write order replications will see significant improvements by using a file-based event-driven product such as R-1 over an I/O based solution. Unlike I/O based solutions that only monitor the source system, R-1 can integrate the event-driven replication with pre / post command processing, repair corrupt or missing data on the target systems, recover from system and communication errors without consuming significant amount of system resources, transfer data in multiple parallel streams, and is scalable to hundreds of target systems. All distribution jobs are available for monitoring through a single monitoring interface, and include the ability to incorporate e-mail notification and database entry of each file transferred and deleted. Furthermore, all enterprise-wide distribution policies run independently of one another and benefit from R-1's market leading automatic recovery and resilience – scalable to hundreds of targets without sacrificing reliability.

2. Introduction and Background

Note: This white paper assumes the reader is familiar with RepliWeb R-1 in a scheduled data-distribution context, and is evaluating the costs and benefits of employing a real-time solution.

RepliWeb R-1 is a high-level, one-to-many file distribution system, engineered to perform in complex and dynamic IP networks (LAN, WAN, Internet, eCDN, satellite based, over Unicast and Multicast). It provides unattended, Center-to-Edge content deployment and synchronization that scales in volume of data, number of processes and number of target machines, employing heterogeneous (cross-platform) server environments (Microsoft Windows™, UNIX). Each content deployment / synchronization job resolves the absolute minimum amount of data to be transferred or deleted between the distribution source and targets (See Comparative Snapshot Matrix (CSM) technology).

In addition to determining the topology of the content distribution, a system administrator must determine when the data consolidation will take place – a factor unique to each enterprise and project. RepliWeb R-1 uses a highly configurable job scheduler to initiate the content distributions according to user-determined policies (at specific times, on the creation of specific files, etc.). The nature of scheduled content distribution and synchronization can result in the source and target systems falling out of sync between the scheduled distributions. While some enterprises are satisfied with weekly, daily, or hourly distributions to synchronize the data between their source and target systems, others require more immediate data consolidation.

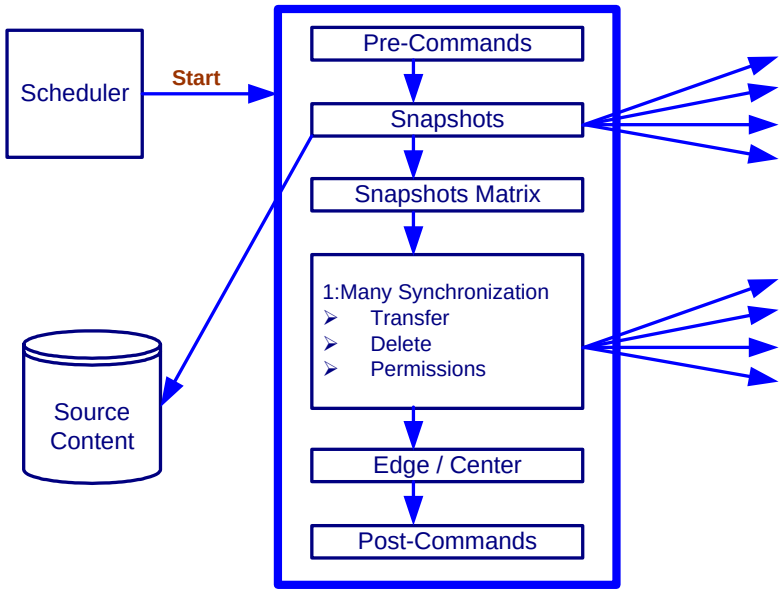

Figure 1 – Distribution Flow.

In scheduled content distribution this can be accomplished by scheduling distributions as close together as possible - at one (1) minute intervals. Due to the time and system resources needed to generate the system snapshots (CSM) and perform the subsequent data synchronization, a one (1) minute scheduling time can become resource intensive, resulting in a continuous loop of distribution after distribution. The only way of addressing this issue is to introduce a real-time replication solution in place of a scheduled one. As we will see in the following section, real-time solutions introduce their own synchronization issues, the severity of which depend on the project's requirements.

This white paper discusses the RepliWeb R-1 **Continuous Update** feature (for Microsoft Windows NT/2000/XP) and how it eludes real-time synchronization issues by introducing an **event-driven** data consolidation stage within a scheduled data distribution framework.

3. R-1 Continuous Update Architecture and Distribution Flow

An R-1 distribution goes through a series of independent stages designed to be completely stateless. This means that nothing in the system is taken for granted, and each stage of any replication process is fully fault tolerant and recoverable. The RepliWeb Event Driven Agent (REDA) is an event driven process - not a device driver - that detects, collects, and filters file level events on the source data that are relevant to each of the distribution targets.

A visual representation of the distribution flow can be seen in **Figure 1 – Distribution Flow**.

The source machine sends instructions to the target(s) requesting a snapshot of a given directory, while in turn building its own. The respective machines do all of the state processing locally, without use of the network or any conversational protocols. Once the targets have returned snapshots to the source, the source machine builds a matrix of which objects need to go to which machines, and which need to be deleted –subsequently performing the necessary operations.

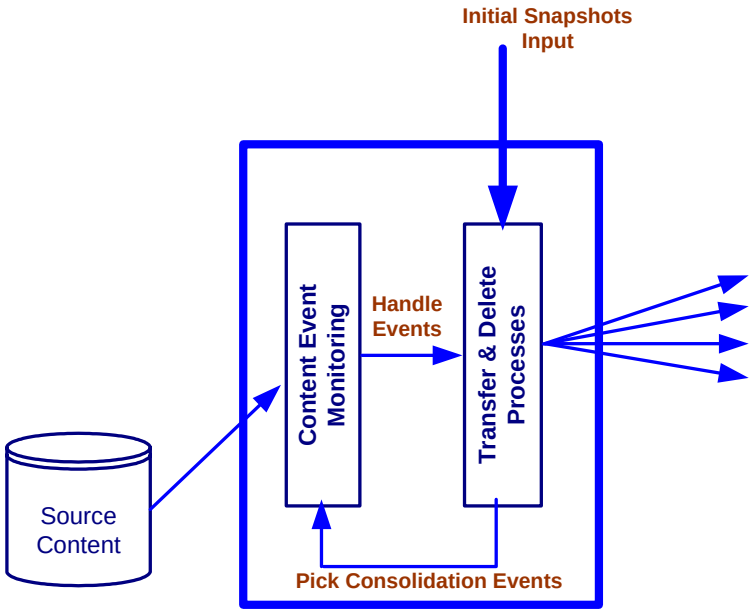

Figure 2 – Continuous Synchronization

From the moment the CSM generation is instigated, the REDA begins detecting and collecting any relevant changes made to the source directory (relevant with respect to the source file and 'exclude' specifications). The REDA then continuously instructs the R-1 distribution to propagate the relevant file / directory creation and deletion events to the target systems. The continuous update distribution will autonomously force a full

synchronization on a pre-determined periodic basis as well as under special circumstances (for example, if a flood of events results in a ‘dropped’ event).

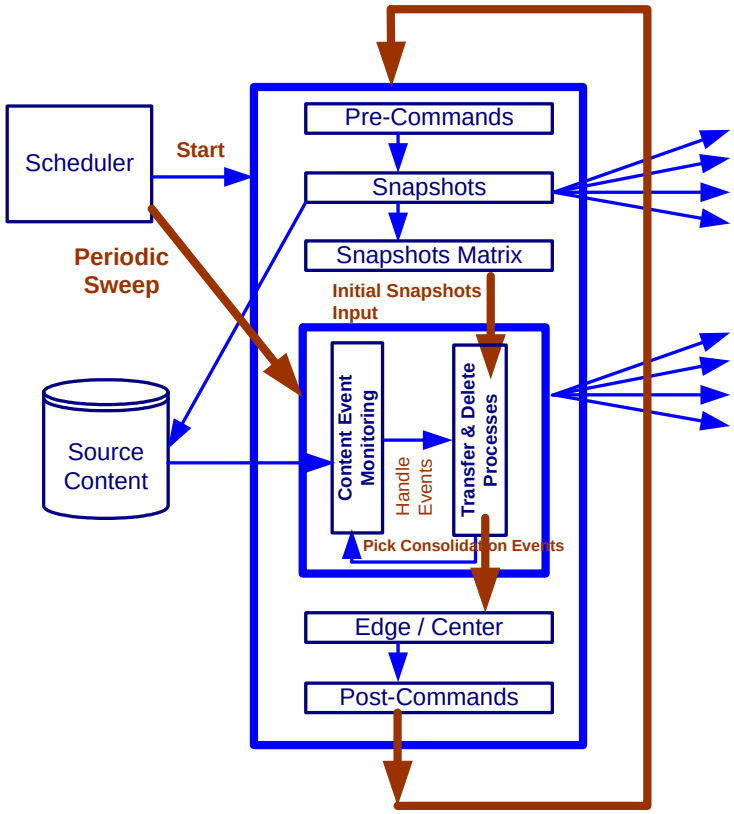

Figure 3 – Combined Operation

The forced synchronization ensures that at fixed intervals, any damaged, missing, or corrupt files on the target are corrected according to the latest source data, in addition to the execution of all distribution level tasks (On-exit commands, e-mail reporting, etc…)

4. I/O vs. File-level Real-Time

Real-time replication solutions can be divided into I/O-level and file-level solutions. I/O level solutions monitor disk write operations on the source system and propagate them to the target system. By replicating the disk write sequence, a critical aspect of database replication, I/O-level solutions sacrifice performance under many conditions and introduce device drivers on the source system. File-level event driven solutions such as R-1 monitor events on the files and directories themselves and enjoy the benefits of being incorporated into a larger distribution framework. For projects where both satisfy the functional requirements equally well, there are a few key demarcation points:

1. **Production issues** - All real-time replication systems assume by design that both source and target systems begin from a synchronized state before the real-time system is activated. This is the single largest hurdle to implementing a real-time solution in production environments. Each and every system failure requires the halting and manual re-synchronization of the systems involved before the real-time system can continue to operate as designed. Production systems are designed to be both fault-tolerant and highly available; properties that are difficult to maintain with a real-time solution. Furthermore, production systems often require the integration of pre/post replication commands into the replication process – a feature I/O replication by nature cannot offer.

   R-1 Solution - RepliWeb solutions are stateless, holding nothing in RAM. Even if the source system were to fail and the collected events lost before or during propagation, the recovery process ensures that any source data changes that occurred during the downtime are propagated from the exact point of failure when conditions allow. Pre/Post command processing can be integrated into the real-time distribution process on forced and periodic full synchronizations, allowing for production integration of the distribution process.
   High availability systems using cluster technology can exploit this recovery method as control is automatically passed to the fail over node if a node does go down.

2. **Target-side considerations** – A real-time system only monitors one side of the equation - the source machine. If a file becomes corrupted or removed from the target machine(s), there is no way for the source machine to detect the problem and re-transfer the damaged file. The file on the target machine will remain unavailable indefinitely. Furthermore, if a file on the target cannot be updated (for example, if the target system rebooted, or failed), there is no guarantee that a memory resident real-time system will be able to propagate the changes if there is a subsequent failure on the source machine.

   R-1 Solution - If a file is damaged, lost or corrupted on the target machine(s), or locked in use, it is repaired on the next forced or periodic synchronization. This option is not available to real-time solutions, as they have no way of knowing what the status of a file is on the target machine. R-1 avoids these risks by forcing a full synchronization upon recovery from a system failure.

3. **Responding to communications failures –** When a communications failure does not permit the propagation of real-time changes, I/O monitoring systems continue to record all changes on the source system. Since these changes cannot be propagated, system memory and subsequently system disk space are exhausted. Each I/O write is accompanied with all the information necessary to write the changes to the target, usually four times more data than the changes themselves. This significantly degrades the performance of the source system during the communications outage, and can even result in a production system exhausting its resources, crippling its production responsibilities. These effects are multiplied in the event of multiple target systems.

R-1 Solution – During communication outages, R-1 continues to track file-level events. Once communication between the source and target systems is restored, the changes are propagated to the target systems in multiple parallel streams. System resources are never exhausted in the event of a system or communication failure. The distribution will attempt to recover through the active data transfer, however, if communication has not been restored after successive recovery attempts, the distribution pauses real-time replication. Once communication has been restored, a full synchronization is automatically carried out that by its nature integrates any changes that occurred during the downtime.

4. **Performance Considerations** – I/O level real-time systems operate on a queuing mechanism, propagating changes on the target systems in the same order as they occurred on the source systems. While this is necessary in database synchronization, it is inefficient and time consuming in file synchronization, especially at high I/O write rates. The performance of I/O level real-time systems degrades as the rate of event generation increases as the events must be queued until they can be replicated to the target systems one-by-one. If the queue becomes too long, the events are moved from memory to the hard disk. This results in a significant amount of disk space being used, as well as further performance degradation as the events need to be reloaded into memory before they can be propagated.

   R-1 Solution – Since file level changes do not need to be propagated in I/O sequence, all file-level changes are transferred to the target system(s) using multiple parallel streams. This significantly improves performance in both LAN and WAN environments. Furthermore, performance is independent of the rate of event generation. If a break-even is met whereby a flood of events overwhelms the source system, a full synchronization based on RepliWeb's CSM technology results in the most efficient propagation of the high number of changes. Only the bare minimum needed to bring the source and target systems is transferred over the network, whether it be permissions, security, or changes to the data itself.

5. Conclusion

The introduction of the RepliWeb R-1 Continuous Update feature allows system administrators the reliability and resilience of a RepliWeb R-1 scheduled distribution, while significantly improving performance in projects requiring immediate data consolidation.

Through the implementation of the REDA over a scheduled distribution solution, R-1 Continuous Update overcomes the pitfalls of standard real-time replication solutions by

regularly performing a full analysis and repair of the target system(s) involved at fixed intervals. The REDA framework allows it to be aware of the number of events generated, ensuring that a flood of events does not overload the system - If such a scenario does present itself, the REDA forces a full synchronization ensuring that a system is always up to date. R-1 Continuous Update has significant benefits over I/O level real-time solutions by allowing integration of pre/post command processing into the real-time distribution process, recovering from system and communication errors transparently and without affecting the source systems performance and stability, and repairing the target file-system(s) periodically and after a system / communication failure.

By integrating the REDA within existing R-1 distribution jobs, all data distribution activity performed through a forced synchronization or through the REDA are accessible through a single monitoring interface. This integration, when coupled with RepliWeb R-1's market leading resiliency and reporting functionality (e-mail notification, database entry, etc.) results in a complete, enterprise-level solution that is capable of handling even the heaviest loads – without sacrificing reliability.

###

For any additional information, please contact us at support.repliweb.com