

Distributed Livelink

ONE INSTANCE, TWO INSTANCES, THREE INSTANCES OR MORE

If you use OpenText Livelink, you may have more than one instance and possibly several instances depending on your organization’s business needs. Are you getting the most efficiency out of your OpenText environment and your Wide Area Network (WAN)? Your business can be more agile and cost effective by utilizing Distributed Livelink.

Distributed Livelink refers to OpenText Livelink environments that span multiple instances – located in either a single location or spanning multiple locations. An instance is defined as one set of database tables used by one implementation of Livelink. This article describes the common application scenarios where a Distributed Livelink solution is warranted and why a Distributed Livelink solution is the answer.

The easiest way to understand where a Distributed Livelink solution is the best answer is to consider the following application scenarios.



Application Scenario

Description

Geo-Replication

Local, national or global deployments of multiple Livelink instances with a single server or multiple load-balanced servers in each implementation.

Extranet with same Domain or Cross-Domain Synchronization

Multiple Livelink instances deployed in the same location to support extranet or other application-specific scenarios.

Real-Time Active-Active Disaster Recovery

Secondary Livelink instances deployed to support disaster recovery and COOP (Continuity of Operation) scenarios

Content Aggregation and Syndication

Collection and/or distribution of Livelink content between two or more web applications (including Livelink workflows and distributed web content management).

Distributed Livelink | ONE INSTANCE, TWO INSTANCES, THREE INSTANCES OR MORE

It is common for most OpenText customers to have a need for one or more of these scenarios in their organization. Let's look at each of these in turn.

GEO-REPLICATION

Why is geo-replication an important application scenario? The business driver for implementing a geo-replication solution is the need for organizations with multiple Livelink instances to synchronize all or part of their Livelink content across some or all of their Livelink instances. This turns out to be a common requirement. The requirement for multiple local Livelink instances can result from several scenarios:

- a merger or acquisition and the need to consolidate or synchronize Livelink content
- a requirement to address WAN bandwidth, latency and reliability issues by deploying additional local or regional Livelink instances
- Smaller deployments of Livelink (departmental or project oriented) that need to be synchronized with or migrated to a company's primary Livelink implementation

The need to address WAN bandwidth, latency and reliability issues deserves special mention. This is the most frequent driver for a Distributed Livelink environment. Here are some examples:

- Lawyers in a global firm will want immediate access to case files
- An engineering team at a remote construction site needs fast access to a large variety of business process documents and potentially very large drawing files
- The crew of a cruise ship with slow or unreliable network connections needs fast, reliable access to up-to-date procedures and information needed to run the ship
- A battlefield scenario where communications between command headquarters and remote Livelink instances may be limited by slow or unreliable WAN connections

For this Distributed Livelink scenario to be effective, it needs to efficiently support near-real-time (NRT) document-level replication.

LIVELINK EXTRANET

For Livelink Extranet scenarios, a common requirement is the ability to publish content authored internally using the corporate intranet to a Livelink extranet that is isolated from the internal network using separate web applications, Forms Based Authentication (FBA) and/or different Active Directory domains. For this scenario, the ability to replicate Livelink content across different authentication protocols (Windows NTLM and FBA) or multiple Active Directory domains are the key requirements.



Distributed Livelink | ONE INSTANCE, TWO INSTANCES, THREE INSTANCES OR MORE

REAL-TIME ACTIVE-ACTIVE DISASTER RECOVERY

There are numerous Livelink disaster recovery strategies and technology solutions, but one of the simplest to deploy and support is near-real-time (NRT) replication of a Distributed Livelink implementation to an online, active disaster recovery instance. This is easily enabled using a NRT replication of document or item-level changes to a disaster recovery instance that can act as a full or partial replica of the source environment. The replication requirements are similar to the previous application scenarios:

- NRT replication.
- A Distributed Livelink replication solution is ideal when the hit disaster recovery instance is located in a different city or a different part of the world where WAN efficiency is an important factor.

CONTENT AGGREGATION AND SYNDICATION

Content aggregation in a Livelink context refers to the collection and replication of Livelink content (almost all Livelink object types are supported including work-flows) from one or more Livelink instances to another instance. A common example is the replication of work products produced in a remote or regional office to a central headquarters farm. Content syndication refers to the distribution or broadcasting of content from a central Livelink instance to one or more remote instances. In addition, there are hybrid scenarios where multiple source instances broadcast content to one or more target instances. The key replication requirements for these scenarios include the above (NRT document-level replication) plus the ability to replicate objects to locations in a target environment that is structured differently from the first.

SYNERGY REPLICATOR FOR LIVELINK

Synergy Replicator for Livelink is an enterprise Livelink replication solution that supports NRT, cross-domain, document-level synchronization of any network of Distributed Livelink instances. Replicator installs and is managed in the same way that other Livelink modules are installed and managed. Replicator detects changes in the Livelink content as they occur and batches these changes into replication packages. A Java based Transfer Service (Java Assistant) is employed by Replicator as the highly-reliable, re-startable transport for downloading replication packages over fast, slow and unreliable network connections. Replicator uses Java compression capability to compress packages making for efficient use of the WAN bandwidth.



About the Author:
Greg Worden has over 13 years experience in all phases of Livelink including designing, implementing and tuning Livelink implementations. Greg is a principal consultant at Synergy, a global provider of enterprise data replication, and operations management solutions for OpenText Livelink Products and Technologies