

QUICK TAKE



September 2, 2005

Don't Rely On Just One Method For Distributed Access To Development Assets

by **Carey Schwaber**
with Liz Barnett

EXECUTIVE SUMMARY

Providing access to the most recent versions of development assets like source code and test scripts is far more challenging for distributed development organizations. Today's software configuration management (SCM) solutions support a variety of approaches to distributed access. The suitability of each approach depends principally on the nature of the development shop's distribution — a main site with a few smaller sites, equal distribution across several sites, many sites with just a few people, or some combination thereof. Heterogeneously distributed shops will require multiple kinds of support for distributed access. And because the distribution of development shops usually becomes more diverse with time, forward-looking shops will adopt SCM tools that solve the problem of distributed access in multiple ways.

DISTRIBUTED ACCESS — LIKE DISTRIBUTION ITSELF — TAKES MANY FORMS

For many development organizations, geographical distribution — including but certainly not limited to outsourced development and joint development efforts — is a fait accompli. The question is how, not whether, to manage multisite development efforts. One of the many challenges associated with distributed development is providing access to the most recent versions of development artifacts like requirements, models, source code, test cases, documentation, and associated metadata. This is obviously problematic when a single development team is split across multiple sites. But even team collocation doesn't obviate the need for distributed access; cross-platform development, asset reuse, and service-oriented architecture (SOA) initiatives all increase the need to manage assets across teams.

Today's SCM solutions offer better support for distributed development teams than ever before. In addition to serving as asset repositories, they now include the change management capabilities necessary to coordinate work performed in multiple locations.¹ While each SCM solution previously supported only one approach to distributed access, market-leading SCM solutions are increasingly enabling users to tackle the problem from multiple angles. There are three general approaches to distributed access:

- **Centralized repositories with remote access.** Most SCM tools let development organizations store shared assets in a single, centralized repository and provide other sites with remote access to that repository (see Figure 1-1). When a single developer relocates, this is the easiest and most economical way to provide him or her with access to up-to-date development assets. The primary drawbacks associated with this approach are network performance and availability. Both factors vary by location, although they become less significant obstacles with every passing year.² Nevertheless, tools that rely on this approach must include strong support for disconnected usage in the event of network downtime or brownouts. To cut down on administrative costs, companies can also engage a third party to host their repositories.³

Products that support this approach include: CollabNet Enterprise Edition, IBM Rational ClearCase Change Management Solution, Serena ChangeMan Dimensions, and VA Software SourceForge Enterprise Edition.

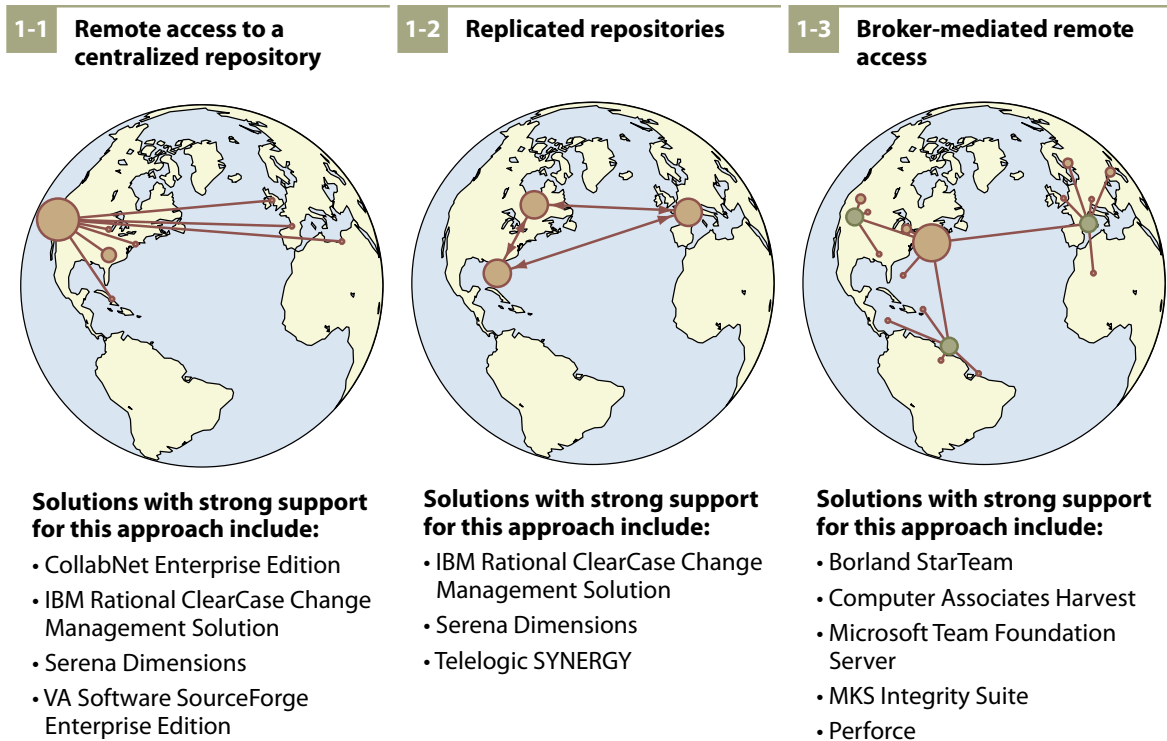
- **Replicated repositories.** Another popular option is the replication of repositories across locations (see Figure 1-2). Changes are stored in repositories at various sites, and these repositories are synchronized at regular intervals. The downsides to replication include the lack of real-time data about changes made in other sites and the administrative burden associated with reconciling changes upon synchronization. But replication is widely used by some of the world's most distributed development organizations — for example, a major communications company with 95 replicated development sites. No other technique has been successfully used on such a large scale. However, while replication can be an efficient solution for larger development sites, it's less appropriate for smaller sites.

Products that support this approach include: IBM Rational ClearCase Change Management Solution, Serena Dimensions, and Telelogic SYNERGY.

- **Broker-mediated remote access.** Another way to overcome the performance limitations of a centralized repository strategy is to install message brokers in network-near locations (see Figure 1-3). Installing such brokers in a hub-and-spoke model smoothes traffic peaks and offers far better performance than unbrokered remote access. Remote access mediated by brokers provides all sites with access to real-time data about changes in all sites and requires less administration than replication does. But this approach still requires strong local network performance and availability, so it's important that shops adopting this strategy make sure they have strong support for disconnected usage. End users report high levels of satisfaction with this strategy, but it is less road-tested than other strategies for distributed access.

Products that support this approach include: Borland StarTeam, Computer Associates (CA) AllFusion Harvest Change Manager, Microsoft Visual Studio Team Foundation Server, MKS Integrity Suite, and Perforce.

The best SCM solutions support more than one of these methods. These solutions are the right choice for development organizations with multiple types of distribution. For example, a shop with three main sites and a few dozen work-at-home developers might want to replicate its three main sites' repositories, while providing remote access for its home workers. And even shops that rely on remote access to a centralized repository might want to replicate their repository for backup purposes. Enterprises should adopt SCM solutions that are flexible enough to grow with their development organization — not just in terms of number of users but also in terms of number and size of locations. After all, it's difficult to predict how a shop will be distributed in five or 10 years, and it's not easy to switch SCM solutions.

Figure 1 Three Strategies For Distributed Access To Development Assets

Source: Forrester Research, Inc.

Market-leading SCM solutions that offer strong support for both remote access to a centralized repository and repository replication include IBM Rational ClearCase Change Management Solution and Serena Dimensions.⁴ And both Borland StarTeam and MKS Integrity Suite support direct remote access to their centralized repositories as well as remote access mediated by brokers. Other vendors will improve their support for additional means of access within the next few years. For example, CA is investigating the possibility of supporting repository replication in future releases of Harvest.

ENDNOTES

¹ Standalone SCM is rapidly becoming a thing of the past. In response to the commoditization of version control tools and even base SCM tools, vendors have transformed their SCM tools into process-centric change management solutions. And in recent years, process-centric SCM solutions have come to serve as the infrastructure of development life-cycle management suites, managing changes to life-cycle artifacts like requirements, models, code, test scripts, and process documentation. See the July 22, 2005, Trends “The Expanding Purview Of Software Configuration Management.”

- ² In general, bandwidth costs have continued to decline. The overall level of service and reliability is such that remote offices may be operated reliably with this new generation of “business Internet” services at a fraction of the cost of traditional approaches. Latency and bandwidth are sufficient to support a reasonable mix of both transactional and miscellaneous IP traffic. See the March 24, 2005, Trends “Remote Offices: Critical Links In Enterprise Architecture.”
- ³ The only vendor currently offering such services for its own SCM solution is CollabNet. CollabNet Enterprise Edition, the product in question, is more than an SCM tool; it’s a collaborative development environment that happens to include a Subversion repository. Telelogic recently acquired FocalPoint, a portfolio management tool with extensive decision-making support; this product is sold on a subscription model and is offered as a hosted service. However, Telelogic does not provide hosting services for its SYNERGY SCM products. No other market-leading SCM vendor offers hosted SCM or partners with third parties to provide hosting of their SCM solutions. See the April 29, 2005, Quick Take “For Simple Version Control, Subversion Beats CVS,” and see the June 14, 2005, Tech Choices “For The Short Term, Keep Using Specialized Collaborative App Dev Software.”
- ⁴ IBM Rational ClearCase has been known for its sluggish performance over WAN, but the introduction of the ClearCase Remote Client, which is optimized to operate through firewalls and in low-bandwidth environments, is a step in the right direction. Rational will enhance the functionality of the ClearCase Remote Client in upcoming releases.