Catalogic® ECX™: Copy Data Management for SQL Server Databases

Catalog. Automate. Transform.

ECX for SQL Server Highlights

- Automate the creation and use of SQL Server data copies — snapshots, clones and replicas – on existing enterprise storage infrastructure.

- Automate SQL Server test and development infrastructure provisioning, reducing deployment time as much as 99%.

- Supports both SQL Server standalone and AlwaysON modes.

- Agentless model simplifies management while still providing application awareness.

- APIs allow integration with popular DevOps tools.

- SQL Server log capture provides point-in-time recovery.

- Simple licensing based on storage controllers.

- Supports SQL Server running on either physical servers or VMware virtual machines.

- Supports a multi-vendor selection of storage arrays, including systems from Pure Storage, IBM, NetApp and DellEMC.

Microsoft SQL Server databases are a fundamental component of many organization's most critical business processes. As such, copies of SQL Server data are needed for many IT functions, including local and remote recovery, development and test, reporting, compliance and so on. According to IDC, over 80% of organizations maintain more than ten copies of every SQL Server database instance at any given time. Not only does this create a massive infrastructure cost penalty, but the ongoing demand for copies also consumes valuable IT staff time and effort.

For this reason, Copy Data Management (CDM) is fast-becoming a must-have solution for any enterprise SQL Server environment. ECX integrates storage array copy processes (snapshots, replication, clones) with SQL Server, allowing customers to leverage their existing storage infrastructure to eliminate copy sprawl, while dramatically improving their capabilities to deliver fresh copies of key databases to those functions and business units that demand them.

The First Step: Automating the Copy Process

Because of the unique characteristics of most SQL Server environments (high I/O, need for continuous uptime, strict protection SLAs, etc.), performing copy creation and management requires knowledge of the database management system and specific points of integration to ensure that the copy process has minimal impact, completes successfully, and results in application consistent copies every time. The bi-product of proper CDM is the end of copy data sprawl, as the ad-hoc and unmanaged processes for creating copies are eliminated in favor of a centralized, holistic approach.

With ECX, users create copy policies that define key copy metrics: how often to make copies, where the copies will reside (local, remote or both), how long copies are retained and so on. Once a policy is created, it can easily be applied to any SQL Server workload. No more complex scripting or working with multiple tools; just click-and-apply ease of use.

The ECX copy process is agentless. Rather than having to deploy and maintain agents across all SQL Server hosts, ECX uses a run-time code injection method that gives you the benefits of an agent – true application awareness – without the maintenance headaches. It supports SQL Server running on either physical servers or as VMware virtual machines.

The meta-data catalog at the core of ECX tracks all copies, making them searchable and reportable. ECX also handles the deletion of copies, ensuring that you only retain the data you really need.

Using Data Copies: Automation and Self-Service

While copy creation is the starting function of ECX, the true value shows itself in using SQL Server data copies. It is not uncommon for IT organizations to spend significant staff time and effort creating, moving and delivering copies of SQL Server data, whether as mounted data copies or by bringing up new virtual machines. In many cases, each copy requires coordinating multiple teams, seeking layers of approvals, and then finally waiting for a slow copy process to complete. End-to-end this process can take several days (considered fast!) to several weeks (more typical). The result is often data that is stale when it arrives, with subsequent impacts to project quality and timeliness.

Imagine if you could automate the entire process, or hand it off via a user self-service portal, all while being able to deliver fresh copies of data in only minutes? With ECX you can! Application consistent copies of SQL Server can be automatically mapped to host servers every day (via iSCSI or Fibre Channel), ready to go when the work day begins. Ad hoc requests can be satisfied with a few clicks.

If SQL Server is virtualized, ECX can spin-up the entire workload of one or more VMs, changing IP addressing and security on the fly to fit the recovery requirements (e.g. whether restoring a lost production system, or bringing up a test environment).

Self-service access means that IT can relieve itself of the daily grind of copy provisioning while still retaining ultimate control and not giving direct access to storage systems. For example, using the ECX self-service portal, a developer could spin up their own SQL Server copies whenever they needed them, but they would be restricted to a pre-defined set of storage volumes or VMs. This prevents abuse of infrastructure and maintains security.

The benefits of ECX copy delivery are even more pronounced when used in conjunction with all-flash storage arrays. The high I/O of these systems allows

zero-footprint snapshot copies to run at production speeds, even when servicing multiple workloads from the same copy.

Data Protection and Disaster Recovery

Through its template-based management and orchestration of application-aware SQL Server copies, ECX becomes a powerful solution for next generation data protection and recovery. You can skip traditional backup completely and move to a snap-and-replicate model that provides near instant protection and rapid recovery. Compared to the slow restores of traditional backup, ECX allows IT to mount and instantly access copies that are already in the production storage environment. ECX catalogs all snapshots and replicas and alerts you if a snap or replication job was missed or failed. Disaster recovery can be fully automated and tested non-disruptively.

Flexible recovery options allow SQL databases to be restored to the original location or to an alternate server, using the original database name or an alternate.

Complete Dev-Test Automation

With ECX, data copies or virtual machines can be delivered automatically to Dev-Test teams, allowing them to work with true production data sets that contain fresh data. Experience shows that using fresh production data leads to fewer bugs and faster development cycles than using outdated or "synthetic" data sets.

When required, older data sets can also be delivered, allowing test and development teams to look back at previous data sets that may be needed for troubleshooting.

Because Catalogic uses an in-place copy model, all development work is done on the same storage stack used for production. This ensures there are no anomalies introduced into the development process, as may happen when using third party storage that is different from the production storage. Catalogic also allows for proper performance testing because it runs as an out-of-band control plane and doesn't sit in the data path between servers and storage.

When needed, pre- and post-scripts can be used as part of data delivery workflows to produce a secure (masked) copy of SQL databases using the Dynamic Data Masking feature of SQL Server 2016.

Dev-Ops Tool Support via APIs

All ECX features can be controlled via the ECX RESTful API. This means that ECX offers true "infrastructure as code" to agile, DevOps environments that require copies of SQL Server data. With a single line of code from a DevOps tool, developers can bring up SQL Server data copies or even full working systems (storage, networking and compute, defined as needed). Catalogic supports popular DevOps tools such as Chef, Puppet, Ansible, IBM Bluemix, IBM UrbanCode and more. In addition, Catalogic offers pre-built scripts to help with DevOps integration, as well as plug-in tools.

SQL Server Log Management

Array snapshots are an excellent way to capture large data sets quickly, but SQL Server transactions will continue to take place in-between snapshots. For this reason, ECX includes SQL Server log capturing. By keeping track of logs, ECX can provide point-in-time recovery to transactional points that fall between snapshot captures. This also allows rolling forward to a point more recent than the last snapshot.

Simple Licensing

Unlike other solutions that require complex licensing based on data size, CPU cores, SQL Server instances or other metrics subject to continual change and increase, Catalogic uses a simple storage controller-based licensing system. License the storage controllers you wish to use with Catalogic copy data management, and there are no concerns about data size, number of SQL Server instances, etc.

VADP Backup

For SQL databases that are not on supported storage arrays, Catalogic ECX provides a VADP backup mechanism that uses block-level incremental copies to move data to supported storage arrays. From there, ECX supports the same set of snapshot and replication features.

Support

For virtual environments, Catalogic ECX supports SQL Server 2012, SQL Server 2014, and SQL Server 2016 standalone and AlwaysON, installed to a virtual machine using a VMDK configuration, running Microsoft Windows 2008 R2, Windows 2012 R2, and Windows 2016.

For physical environments, ECX supports SQL Server 2012, SQL Server 2014, and SQL Server 2016 running on Windows 2012R2 and Windows 2016, in standalone mode, Windows Server Failover Cluster (WSFC) or AlwaysON with WSFC. Additional OS support will be added in future releases. Initially, physical support will be limited to Pure Storage and IBM storage arrays supported by Catalogic. Note that physical server support will be released with Catalogic ECX version 2.6.

Alliance Partner