perspectium

Replicator for ServiceNow

Frequently Asked Questions

Q: What is Replicator for ServiceNow?
A: Replicator the first scalable, real-time Data Replication service for the ServiceNow platform. Replicator uniquely propagates ServiceNow data objects without loss to other ServiceNow instances, other SaaS applications/platforms, databases and middleware. Replicator maximizes data output while minimizing the impact on production instance performance.

Q: What are the typical use cases?
A: To support the software development process, Replicator provides real-time propagation of ServiceNow objects and changes from the ServiceNow PROD instance to the TEST and DEV instances. This ensures that all changes, including update sets, are being applied and tested on the current production data and eliminates the requirement for frequent ServiceNow "cloning".

For installations where reporting negatively impacts production instance performance, Replicator is used to populate and replicate live production data on a Dedicated Reporting Instance.

For enterprise sites with multiple ServiceNow production instances, Replicator provides real-time synchronization of ServiceNow objects bi-directionally between peer instances. Replicator can also simplify the process of instance migration or consolidation.

Replicator provides two-way real-time synchronization of ServiceNow objects with other SaaS platforms, on-premise applications or local databases, including but not limited to: Jira, MySQL, Oracle, Salesforce.com, Verizon or AT&T eBonding and middleware (like Tibco or Oracle Fusion).

Q: Is Replicator Now Certified for Integration with ServiceNow?
A: Yes.

Q: How is Replicator different than other data base replication services?
A: Unlike a typical point-to-point data base replication tool, Replicator is a cloud service that utilizes an enterprise message bus and provides a flexible publish and subscribe framework to propagate ServiceNow data objects. To maximize replication performance and minimize the impact on the production instance, Replicator avoids web interfaces (SOAP, JSON) that are heavily used by other integrations.

Q: How is Replicator different than data export tools?
A: Data export tools integrate to ServiceNow via web services which limits output capacity, are typically batch oriented, don't support differential replication and place a heavy burden on production instance performance. Replicator is a scalable, real-time data replication cloud service that minimizes the impact on ServiceNow performance, and is the only solution to deliver a flexible publish and subscribe interface - publish once, subscribe everywhere.

Q: How does Replicator ensure data isn't lost?
A: The enterprise message bus combined with our store and forward capability ensures data integrity – that is, no data lose. As a monitored replication service, we provide real-time visibility into the replication process (performance and completeness).

Q: How does Replicator ensure data security?
A: Replicator requires two-way authentication to protect replicated data. First, access to the Replicator cloud service requires authentication with a username and password. Second, Replicator utilizes a 24-bit encryption key to encrypt the payload at the source and this key must be known at the target to decrypt the payload.

Q: What are the scalability limits (frequency and volume)?
A: Replicator's enterprise message bus can be clustered to offer virtually unlimited scalability. Throughput is constrained by how many records ServiceNow can push into the message queue and how fast.

Q: How does Replicator handle schema changes (new tables/fields)?
A: Replication integrates with the ServiceNow SCHEMA services to dynamically adjust tables.

Q: Is there control over replication frequency and schedule?
A: Replication timing and frequency is controlled using conventional ServiceNow controls (like the Scheduler).

Q: Does Replicator impact the performance on the production instance?
A: Replicator was specifically designed to provide the highest data transfer rate with the least performance impact. We accomplish this by avoiding web interfaces (like SOAP and JSON) that are heavily used for other integrations. We reduce the number of point-to-point integrations by utilizing a flexible publish/subscribe interface. Furthermore we leverage the ServiceNow background scheduler to ensure we play fairly with resource demands on the platform.

Q: In what way can Replicator enrich the data?
A: The Replicator service can cleanse or add data to the payload at the message bus layer.

Q: Does Replicator support two-way integration?
A: Yes, simply by switching the publish/subscribe roles.

Q: Why is monitoring Replicator important to operations?
A: Replicator monitors availability and throughput, providing answers to questions like:

- What is the current state of the replication process?
- When will it complete?
- Is it meeting our performance criteria?
- Were there any problems detected?
- Why did they occur and what steps should be taken to remediate it?

Q: What types of problems does development experience with inconsistent data?
A: Without consistent data, these primary problems often arise:

- Data structures created in Prod are not represented in Dev (for example, CI relationships with deep tree structures) where recursive scripts can create performance issues
- Meta data created in production can be missed, which impacts the development process and quality
- Unexpected data values imported via data loads often create unexpected results
- Equivalent performance testing cannot be undertaken with equivalent volumes
- End-user UAT testing is more effective if they can replicate situations with live tickets and associated data.

Q: What are the challenges with instance cloning for dev and test?
A: Cloning production systems can help remediate the gap in data, but this will typically require more time built into the development cycle to accommodate these events. Code will generally need to be saved and re-applied following the cloning process. The overhead to keep data consistent is high.

In some cases the cloning of data from Prod to Test requires different parameters than from Prod to Dev systems due to data security considerations. This can represent a further overhead.

Q: Why is data inconsistency a problem for DevOps?
A: Seconds after a production instance is cloned, the cloned instance is out-of-date. Production data by its very nature is changing in real-time; CMDB CI changes, new incidents, new users, configuration changes, etc. There can be drastic differences between dev, test and production data, which often results in additional development (rework) time, operational cost and systems overhead.

Q: Is Replicator any different than the IR sync update set being shared on chard.service-now.com?
A: Replicator is a product that enables publish and subscribe integration technology for ServiceNow, whereas IR sync is a be-spoke web service that integrates 2 instances point-to-point. In addition to the performance and scalability benefits leveraged from an optimized message bus middle ware, Replicator is also a fully supported, extensible integration framework that allows expansion into other non-ServiceNow environments. Lastly, Replicator is integrated with Observer to provide monitored replication in the form of real-time reports on data flow, sustainability, and availability.

Q: How is Replicator licensed and delivered?
A: Replicator for ServiceNow is delivered as a cloud service and licensed under an annual subscription fee.

About Perspectium
Founded by ServiceNow pioneers with a deep insider understanding of the ServiceNow platform, Perspectium provides a unified set of tools/services, deeply integrated with the ServiceNow platform and delivered as-a-service, to optimize ServiceNow platform performance and minimize the challenges with production data access and integration.

To learn more visit us at perspectium.com, email us at info@perspectium.com or call toll free at 888-620-8880.