Maximizing the Robustness of Your Data Movement Infrastructure

In response to end-user expectations for real-time information, regulatory pressure and a desire to maintain business continuity, companies are setting increasingly aggressive goals in the areas of minimizing downtime in their IT systems (sometimes called recovery time objectives or RTOs) and reducing the risk of data loss in the event of system failures (sometimes called recovery point objectives or RPOs).

There are three key challenges companies need to address when striving to meet such goals in the systems that move real-time data throughout their enterprise:

- **High Availability:** Real-time data movement infrastructure needs to be extremely reliable for continuous operation, and fault tolerant so the flow of data isn't interrupted by problems with individual applications, connections or devices.
- **Disaster Recovery:** Sometimes unpredictable disasters cause unavoidable outages at the datacenter level, and to quickly resume operations companies need their backup data movement infrastructure to pick up where the primary system left off, as quickly as possible and with no loss of information that was in transit.
- **Slow Consumer Handling:** These days most applications send information to a wide range of consumers over different kinds of networks, many of which aren't always able to consume information as quickly as it's sent. Data movement platforms must be able to accommodate these "slow consumers" by delaying delivery until they're able to catch up, in such a way that high-speed publishers and other consumers aren't impacted.

Solace message routers are inherently robust and uniquely capable of addressing all of these challenges. This paper describes how Solace can help you implement the most robust real-time data movement infrastructure possible.

Solace Systems®

Fault Tolerance for High Availability

Most messaging systems claim they provide high availability via fault tolerance, but it's important to consider the real meaning and ramifications of "availability" in your data movement infrastructure, and potential impact of problems.

The exact mechanism by which fault tolerance and high availability are achieved, and the amount of time that failover takes, can have serious ramifications on your system in terms of downtime, delayed delivery and data loss. Fault tolerance that suffers from slow or inconsistent failover frequently fails to meet key HA goals.

Solace message routers support high availability and robustness with architecture and features such as redundant components, automatic fail-over of paired devices, isolation of control and data planes, and per-client queue management.

Fault tolerance that suffers from slow or inconsistent failover behavior doesn't meet most companies' HA objectives.

HA Architecture and Failover Process

Solace provides full appliance redundancy within the datacenter, eliminating all potential single points of failure.

- Solace message routers feature redundant components (e.g. hard drives, power supplies, fans, Ethernet ports and fiber links) so they can tolerate and recover from component failures without resorting to an HA failover or impacting applications.
- To support guaranteed messaging, Solace message routers are equipped with an optional card called the Assured Delivery Blade (ADB) that implements a non-volatile message store using FPGAs, DRAM and proprietary hardware.
- Two Solace message routers work together in an active-hot standby configuration, with the ADBs of each appliance connected to their mate via a pair of redundant 10 Gbps fiber links.
- These message routers also have shared access to an external disk array for use in storing messages for slow or offline subscribers.
- Connectivity to the SAN is through a dual ported HBA card installed in each message router, so each one has redundant paths to the SAN.

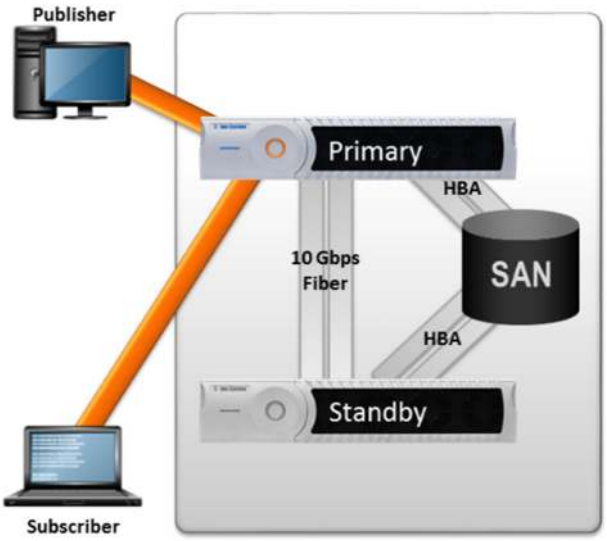

The Importance of Fast Failover

Solace uses a patented approach to maintain real-time state on both active and standby message routers so activity switches happen in around 15 seconds no matter how much data has accumulated in the spool. Software-based messaging middleware products, on the other hand, persist messages and their state on disk, so in the event of a failure the standby broker must retrieve all state from disk before it can resume providing service.

As such, to understand the impact an outage can have you need to consider the potential of large message spool. For example, assume a software-based message broker takes around 30 seconds to fail over with 1 GB of data in the message spool. Since failover is contingent on reading all data from disk before service can be started up, failover time scales linearly so with 10 GB in the message spool it would take several minutes for the backup system to even *start* processing messages. Extrapolating to message spools with hundreds of gigabytes in them, the software broker would be down for several hours. Meanwhile, Solace failover times remain consistent at around 15 seconds.

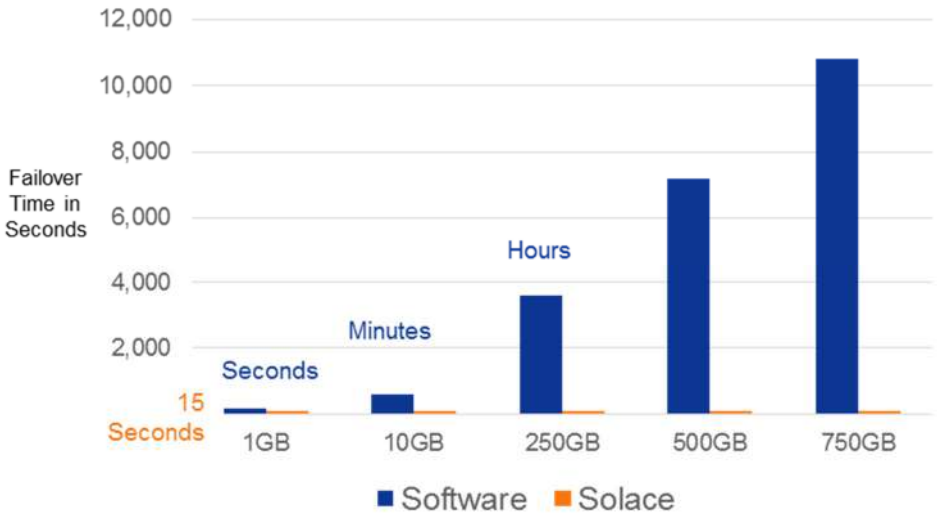

Video Demonstration

Solace has created two videos that demonstrate the automatic failover of Solace appliances, with typical and very large message spools, that you can watch at http://solacesystems.com/resources/high-availability-video

Integrated and Automated Disaster Recovery

Ensuring business continuity when disaster strikes is mission critical. Most software-based data movement solutions rely on storage replication products to transfer in-flight messages to a backup datacenter – a complicated approach that involves the coordination of multiple teams, takes a long time, and introduces the risk of losing messages.

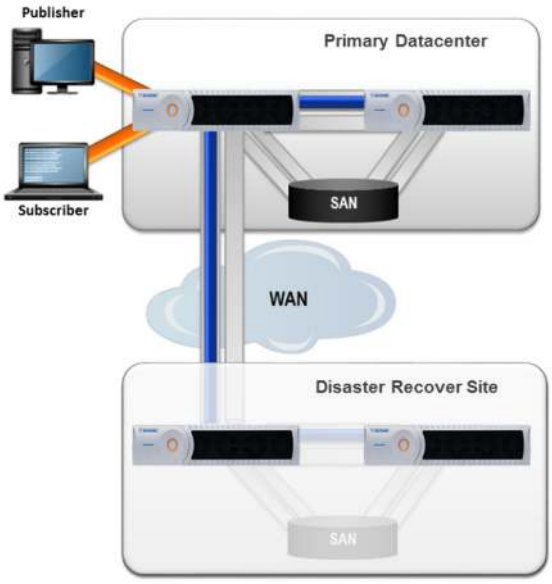

Solace's replication architecture keeps remote appliances continuously synchronized so they can pick up message processing in just seconds once an administrator decides to fail-over to the backup site.

Solace handles disaster recovery without third-party products by propagating messages directly from message routers in the active datacenter to appliances in the DR site. Solace lets administrators configure if and how messages are replicated on a per-topic basis, giving them the option of synchronous and asynchronous replication.

Synchronous Replication

With synchronous replication, which is used for ensuring the delivery of the most mission-critical messages, messages must be successfully replicated to the backup site before acknowledgements are sent to publishers. When replicating messages synchronously, software-based systems are limited by the performance of storage-replication tools that were designed for read-centric uses cases, versus write-centric applications like real-time data movement.

Solace's synchronous replication solution, on the other hand, leverages very high performance inter-device communications so performance is limited only by the available bandwidth and round trip time of the WAN link carrying the data. Additionally, the Solace APIs provide the ability to stream published messages into the appliance. In this mode of publishing, applications register call-back acknowledgement functions that are invoked as soon as the message hits the backup site. This lets applications achieve synchronously replicated publishing rates comparable to asynchronous replication, with the absolute delivery guarantee of synchronous replication.

With Solace's synchronous replication solution performance is limited only by the bandwidth between the sites, and the round-trip-time over the WAN link.

Another downside of storage-replication tools is that a single synchronously replicated publisher will head-of-line block all other publishers, preventing subsequent writes to the SAN until the message has been written.

Solace doesn't block the publishing of messages while one message is being synchronously replicated, so an application can synchronously replicate the most mission-critical messages without impacting other publishers whose messages are

being replicated synchronously, asynchronously or not at all. This means many publishers can send synchronously replicated messages at the same time for much higher aggregate throughput.

Asynchronous Replication

With asynchronous replication, publishers are sent acknowledgements as soon as a message is persisted in the primary datacenter. It's used for scenarios in which performance requirements trump the need to completely eliminate the risk of some message loss in the event of a data center failure.

Storage-replication tools asynchronously replicate messages in timer-driven block transfers, which means the risk window for message loss can be seconds long and entail the loss of thousands of messages. Furthermore, storage systems don't understand message and data structure formats, so a failure that occurs while replication is in progress can corrupt the message spool file at the backup site.

Solace's asynchronous replication solution, on the other hand, forwards the message to the backup site as soon as the message is published. This means the risk window for Solace's messaging-centric asynchronous replication is very brief and limited to messages in transit at the time of the failure, and there's no risk of corrupting the backup message spool.

Solace's approach lets applications designed around a blocking-send paradigm (that can't take advantage of the Solace streaming publishing option to efficiently replicate messages synchronously, especially over WANs with high round trip times) to achieve high publishing rates through asynchronous replication, with minimal risk of message loss in the event of a data center failure.

With Solace there is never any risk of corrupting the message spool in backup message routers in the event of a primary data center failure.

Accommodation and Isolation of Slow Consumers

Failures in your network, application or server can cause "slow consumers" that can't keep up with the flow of messages being sent to them, whether they are disconnected, underperforming or overwhelmed by a burst of messages.

Slow consumers burden the messaging system, and in peer-to-peer systems can cause systemic collapse. For software brokers this typically results in severe performance degradation as messages accumulate and need to be read back from disk before they can be delivered. That can cause a cascading that effects the performance of publishers, fast consumers and other brokers. Pretty soon, the whole platform is suffering from slow performance because of one slow consumer.

Solace message routers have been specifically designed to handle slow consumers in such a way that they do not impact the performance of publishers or fast consumers.

Typical Slow Consumer Handling

The chart to the right illustrates how software-based messaging systems behave in the face of slow or disconnected consumers. It illustrates message rate from the perspective of a message broker, with the blue line representing aggregate input rate and the orange line aggregate output rate. The green line represents the message input of one subscribing application.

In Section 1 of the chart, the system is in steady state, where all publishers and subscribers are keeping up.

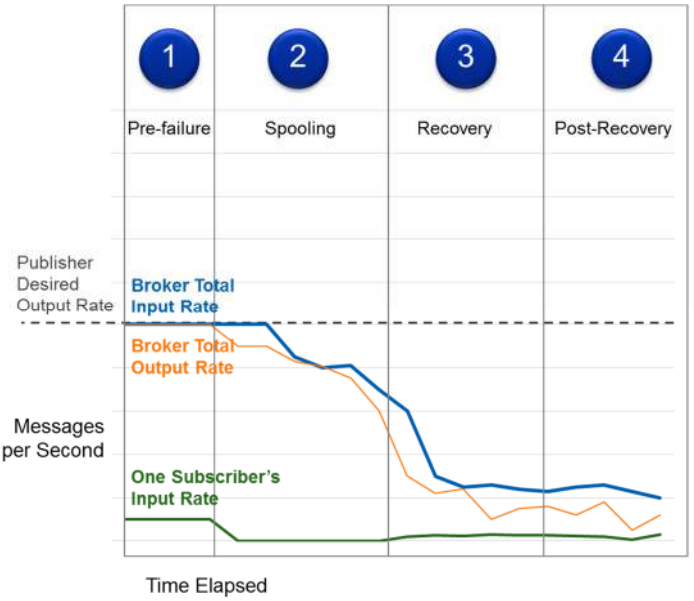

Subscriber Failure and Spooling

Section 2 shows what happens when one subscribing application encounters a fault. At first the message broker may be able to keep up but as the amount of messages and message state grows the performance will eventually degrade. This forces the message broker to push back on publishers and effectively throttle the publish rate. Often the message broker's delivery to fast consumers becomes erratic during these periods and fails to even keep up with the degraded input rate from the publishers.

Recovery

During recovery, operations teams frequently face a dilemma illustrated by section 3. If a failed application is brought online it will drain resources from the message broker as unspooling competes with the processing of fresh input messages. This coupled with the OS prioritization of disk reads and writes further degrades the performance of the message broker, compounding the problem. The situation can get even worse if the message broker starts missing the OS disk cache due to large message spool accumulation.

Resumption of Normal Operation

Section 4 illustrates that the system may never fully recover from this consumer failure, causing consumers to continue to fall further behind, thereby creating larger and larger queues.

Solace Slow Consumer Handling

Solace message routers have been specifically designed to handle slow consumers in such a way that they do not impact the performance of publishers or fast consumers. For guaranteed messaging, publishers and fast consumers are identified and prioritized over slow or recovering consumers even as message storage to slow consumers continues to increase. Our patented approach ensures complete performance isolation, even at rates 20-30x higher than software-based message brokers can handle.

The chart to the right shows how Solace handles slow consumer scenarios. Again, section 1 shows the aggregate input and output of a message broker, along with the input rate of a single subscriber.

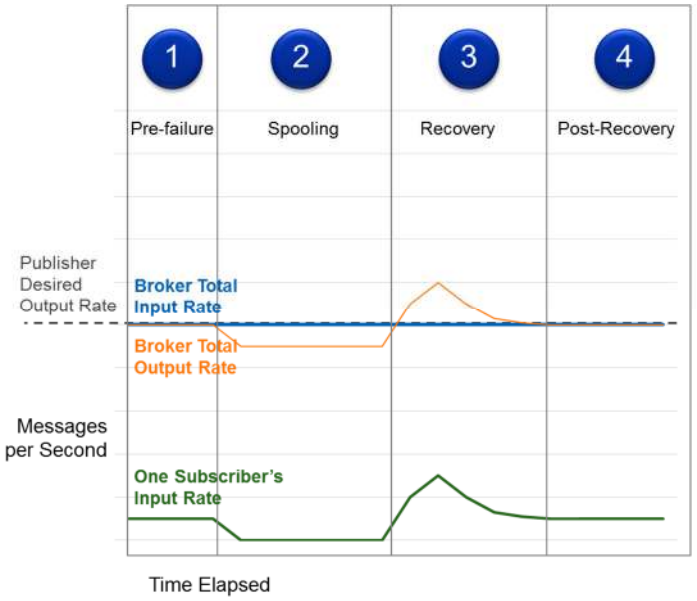

Subscriber Failure and Spooling

Section 2 of the chart shows that the publisher's output rate is unaffected, as is the input rate of consumers that are able to keep up with the flow of messages.

Recovery

Section 3 shows how the message router's high capacity allows it to unspool and deliver messages to recovered subscribers so they can catch up to the real-time flow. The Solace appliance prioritizes input traffic and fast consumers, so even recovery does not affect publishers and fast consumers. This alleviates another key concern for operations teams, as they can recover clients and let them catch up at any time.

Resumption of Normal Operation

Section 4 shows how message flow quickly returns to pre-failure conditions instead of causing a cascading systemic collapse.

Summary

Data movement within modern IT systems demands sophisticated fault tolerance and disaster recovery capabilities that keep information flowing no matter what. By layering innovative functionality on top of its inherently advantageous hardware solution, Solace makes it easier than ever to harden your system so it can keep supporting your business even through the inevitable failures and disasters that will affect your IT systems.

To learn more visit solacesystems.com or call +1 613-271-1010.