

How Many Nines Do You Need?

The Site-to-Site High Availability of Novell Business Continuity Clustering

Almost everyone carries some sort of insurance to protect against unforeseen mishaps; but the protection typically provided is reactive and only comes in the form of cash compensation. If you get in a fender bender, insurance will pay to remove the dents in your car and repaint the damaged area, but it doesn't cover hidden costs such as the hassle of getting estimates, arranging alternate transportation while the car is in the shop, let alone having to drive a car that just isn't in the same condition it used to be even after it is repaired.

In the business world those hidden costs can lead to serious consequences for an organization. When a disaster strikes, can you afford to be without mission-critical services for days, hours or even minutes? How will downtime affect your revenue stream? Your customer relationships? Your ability to compete? A cash policy payout provides little, if any, relief in these areas. That's why more and more organizations invest in the proactive protection that Novell Cluster Services and Novell Business Continuity Clustering provide.

> High Availability Basics

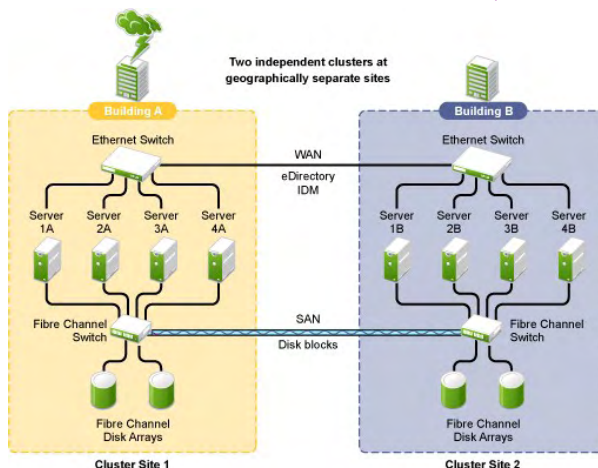
If you use Novell Cluster Services (an entitlement for a two-node cluster is included with Novell Open Enterprise Server), you know that if one of your data center's mission-critical servers fails, the services it hosts will be automatically migrated to another server in the data center within seconds. This provides you and your users with uninterrupted access to critical data and applications. In a Novell Cluster Services cluster, you can have anywhere from two to 32 server nodes participating

in a cluster relationship, sharing the same storage resources and working together to ensure uninterrupted service from that data center. Novell Cluster Services keeps the different cluster nodes in that data center constantly aware of the state of other nodes so that in the event of a server failure it can gracefully move services from one node to another.

The reliability of Novell Open Enterprise Server and Novell Cluster Services delivers many customers an unparalleled 99.999 percent uptime; but there are times when even that five-nines availability isn't enough.

For example, what do you do if disaster strikes your entire data center? That's where Novell Business Continuity Clustering comes in; it lets you migrate your mission-critical services from one data center to another. Built on top of Novell Open Enterprise Server and Novell Cluster Services, Business Continuity Clustering provides site-to-site failover of critical workgroup and networking services. Services running on either NetWare or SUSE Linux Enterprise Server in a Novell Open Enterprise Server environment can easily fail over to another cluster in a completely different geographic location. (See Figure 1.) As a result, even if a major catastrophe affects one of your data centers, you can eliminate downtime and ensure that your critical services remain available.

Figure 1: *Novell Business Continuity Clustering extends the reliability and uptime provided by Novell Open Enterprise Server and Novell Cluster Services with site-to-site failover capabilities.*



> Is Business Continuity Clustering Right for You?

So, when does it make sense to take advantage of the enhanced high availability and disaster recovery that Novell Business Continuity Clustering offers?

If you have multiple data centers that have shared Storage Area Network (SAN) storage and already take advantage of Novell Cluster Services, Novell Business Continuity Clustering is a natural next step in your high-availability setup to even protect you from entire data center disasters.

Novell Business Continuity Clustering is also ideal for government agencies, health care organizations, financial service businesses, and any organization that requires uninterrupted access to mission-critical applications and data. If your organization has data centers in environmentally sensitive locations—such as hurricane,

If you have multiple data centers that have shared Storage Area Network (SAN) storage and already take advantage of Novell Cluster Services, Novell Business Continuity Clustering is a natural next step in your high-availability setup to even protect you from entire data center disasters.

tornado and earthquake zones—you should seriously consider taking advantage of the solution, as should any organization that requires full remote failover in the event of a disaster.

When determining whether Novell Business Continuity Clustering is right for you, you need to first understand the high-availability needs of your business. What data and services are necessary for your business to function? What are the interdependencies for those critical services and data?

Next, you need to calculate the costs of downtime associated with those mission-critical services and data. What's the financial impact of downtime in terms of lost sales, decreased productivity, IT expenses to restore services and any other direct costs? What are the indirect costs, such as reduced revenue because of negative market, customer or partner perceptions? Once you

understand the actual costs of downtime, you can start to recognize the benefits of having better than 99.999 percent uptime with Novell Business Continuity Clustering.

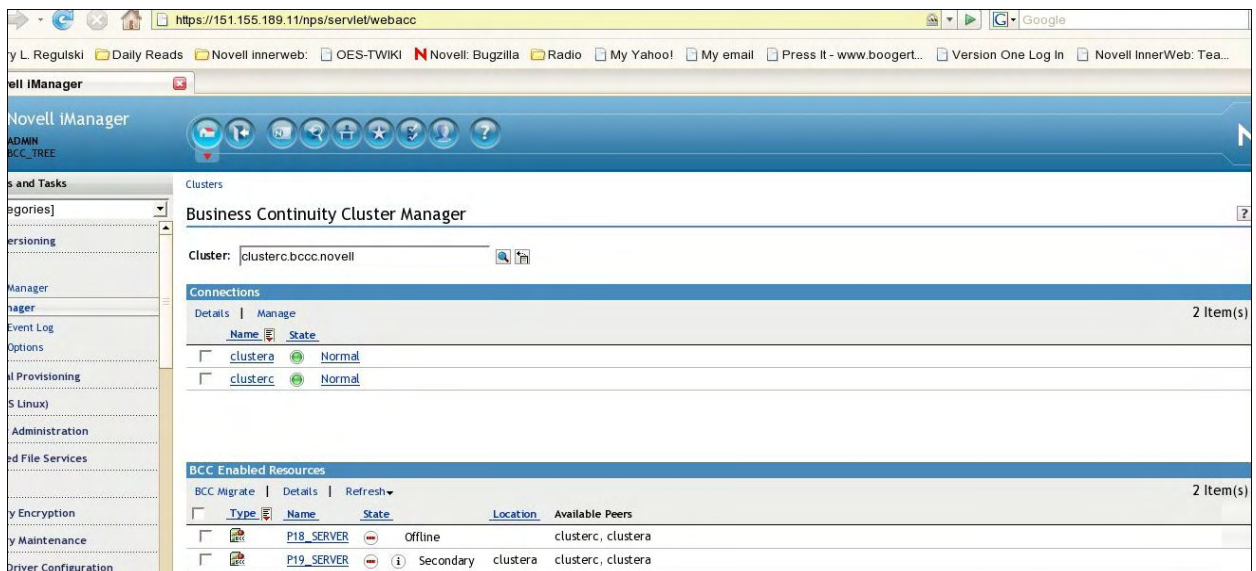
> How Does Novell Business Continuity Clustering Work?

In an effort to protect their mission-critical data and services against potential natural or man-made disasters, many organizations have built and deployed mirrored data centers that are geographically separated. (See *Not Just Disasters*.) Unfortunately, setting up and maintaining mirrored data centers is generally a very manual process that takes a great deal of planning and synchronization. Configuration changes have to be carefully planned and replicated. Any mistake in the administration of a redundant site can prevent it from being able to effectively take over in the event of a disaster. By contrast, Novell Business Continuity Clustering works in conjunction with the mirroring capability of your SAN to automate cluster configuration, maintenance and synchronization. (See Figure 2.)

Novell Business Continuity Clustering utilizes a “cluster of clusters” infrastructure. Each geographically separated data center has its own independent cluster. Each of these independent clusters are treated as “nodes” in a

larger cluster, allowing a whole site to fail over to a different data center site. Novell Business Continuity Clustering automates this failover process by leveraging Novell eDirectory and policy-based management of cluster resources and storage systems.

Figure 2: *Novell Business Continuity Clustering works in conjunction with the mirroring capability of your SAN to automate cluster configuration, maintenance and synchronization of your geographically separated data center clusters.*



For example, say you have a clustered data center in New York and another in Boston. Your data center SAN in New

Once you understand the actual costs of downtime, you can start to recognize the benefits of having better than 99.999 percent uptime with Novell Business Continuity Clustering.

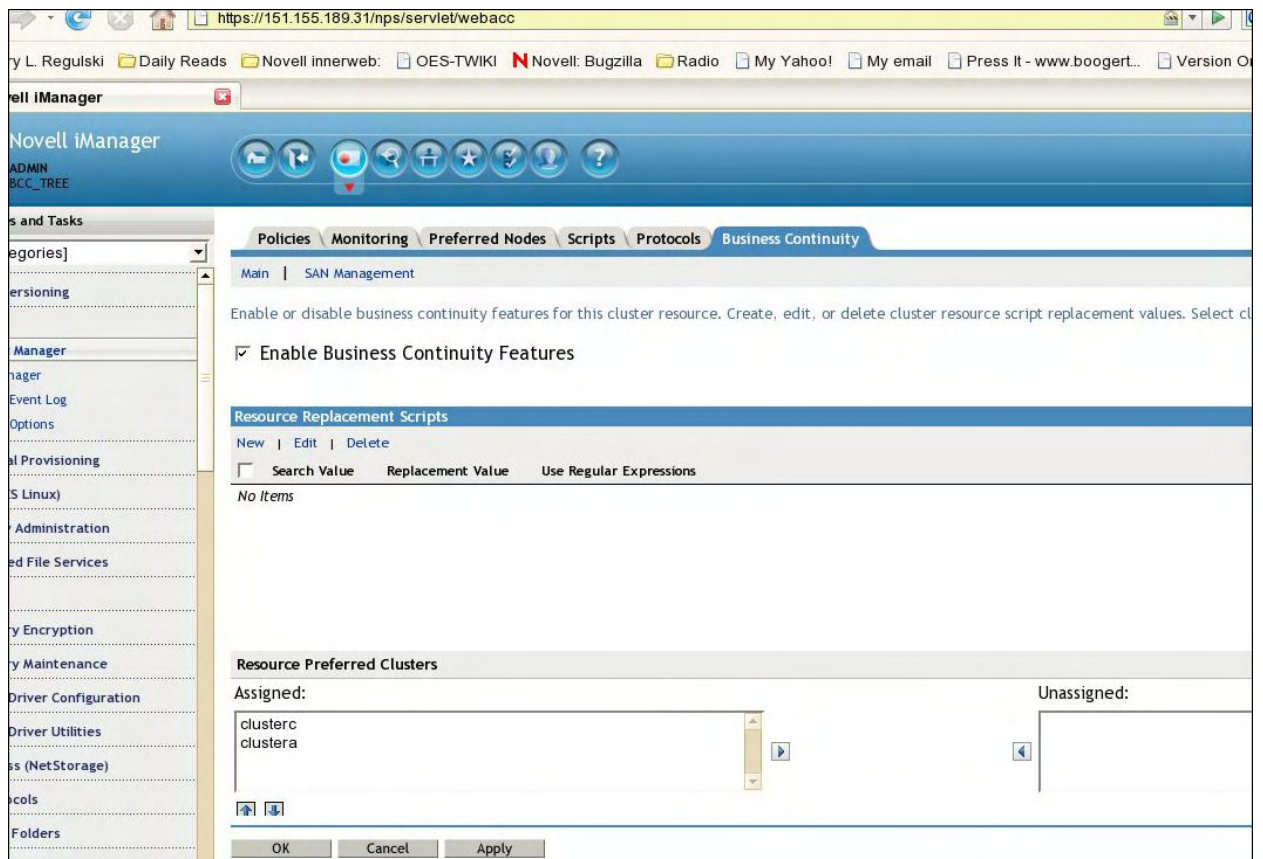
York keeps a mirrored copy of all your business data stored on the Boston SAN, and vice versa. These mirrors are kept in synch using your SAN mirroring software. The role of Novell Business Continuity Clustering is to keep each cluster aware of the status of the others. It uses eDirectory to store resource and configuration information of each cluster, automatically transferring that information between or within directory trees as needed. So, because the New York cluster knows everything about

the Boston cluster, if the entire Boston cluster goes down, the New York cluster will know how to properly fail over all the Boston resources. When needed, a single click will automatically migrate and load these resources onto your New York cluster. (See Figure 3.) (See [Coming Soon: Preferred Clusters](#).) It's important to note that this process does not migrate the business data. That data should already exist as mirrored secondary storage on the SAN. Instead, it's migrating the cluster resources, such as the applications or services that were being hosted by the cluster in the Boston data center.

> Novell Business Continuity Clustering Competitive Advantages

Rather than taking the "cluster of clusters" approach employed by Novell Business Continuity Clustering, most competing offerings use stretch clusters. A stretch cluster is basically one large cluster that has nodes participating in varied geographical locations. So, instead of having a cluster of clusters comprised of a four-node cluster in New York and a separate four-node cluster in Boston, a stretch cluster would simply be one eight-node cluster with four nodes in New York and four nodes in Boston.

Figure 3: With a simple click of a button, Novell Business Continuity Clustering gives you the power to automatically migrate your mission critical services from one data center to a data center at another location.



While this might appear simpler for migrating resources between geographical locations, it introduces some significant problems due to the inherent latency that will exist between the geographically separated cluster nodes.

The first casualty of stretch clusters is reliability. Due to the latency between clustered nodes, your cluster cannot detect unavailability of a resource as quickly. Also due to the distances, split-brain scenarios are more likely to occur. You simply cannot get the cluster reliability like you can in a Novell Cluster Services and Novell Business Continuity Clustering environment. The stretch cluster only provides disaster recovery benefits and cannot deliver productivity enhancements. In fact, in most cases it hampers the performance capabilities and reliability of the cluster.

Not Just Disasters

While Novell Business Continuity Clustering is primarily considered enhanced high-availability protection against data center disasters, it can also deliver unexpected benefits in the areas of infrastructure maintenance and regulatory compliance. These benefits are similar to those experienced by Novell Cluster Services customers. Even though Novell Cluster Services is focused on high availability, administrators have found many ways in which the solution makes their lives easier. For example, while there might not be a major impact on your business if your printers go down over the weekend, it can be a hassle when you get a phone call on a Saturday evening telling you that you need to get the printers back up that night. That's why many customers, in the name of sheer convenience, put Novell iPrint on their clusters. Likewise, Novell Cluster Services allows you to perform server maintenance on a cluster node without bringing down the services your users need. In other words, you can perform maintenance on server hardware during regular hours, rather than waiting for the weekends. Customers will discover similar benefits as they take advantage of Novell Business Continuity Clustering. A customer at BrainShare stumbled upon one such benefit as Kent Boogert, Development Manager at Novell, demonstrated how easy it is to extend Novell Cluster Services with Novell Business Continuity Clustering. Because of compliance requirements, this customer has to show on an annual basis that his disaster recovery process works as required. Demonstrating this compliance takes a full week or more of his team's time to simulate the disaster, power down the data center and recover data and services. After seeing the demo, this customer realized that not only would his organization improve actual business continuity with Novell Business Continuity Clustering, they could also demonstrate their required compliance by migrating data center resources with a simple click of a button.

The other major disadvantage associated with stretch cluster latency is the effect on heartbeat responses and maintenance. For example, you might configure your Novell Cluster Services environment to automatically fail over a node's resources if it goes for eight seconds without issuing a heartbeat. If you configure a stretch cluster with an eight-second heartbeat response, you'll have servers frequently and unnecessarily failing over due to the normal amount of packets that get dropped over WAN connections.

As a result, you're more likely to lengthen your heartbeat settings to account for the inherent latency. This means that when you do have a failed server, it will take more time for the cluster to recognize it and begin the failover process. Not only does this lengthen your downtime, but the delay can create split-brain situations where more than one node believes it is the primary server, leading to divergent sets of data that require significant effort to correct.

When comparing the different approaches, remember that Novell Business Continuity Clustering not only protects your business systems against disaster, but it also simplifies cluster maintenance and enables you to fully utilize the processing power of your cluster investment, so you can also utilize your mirrored remote site clusters for regular production usage. While mirrored clusters in other vendors' solutions are only used in the event of a disaster, Novell enables you to use the latent processing power of your investment, even if a disaster is never encountered. Also, Novell Business Continuity Clustering is the only solution to provide you a complete site-to-site failover solution for your entire Novell workgroup infrastructure.

> Preparing to Implement Novell Business Continuity Clustering

The first step in ensuring that your critical data and services remain available in the event of a disaster is to design your infrastructure based on your business needs. This means identifying the key system factors that drive your business. You need to determine which of your services are most critical to your operations, where those services currently run, and where they need to be running to ensure business continuity.

One of the keys to business continuity is to make sure your individual data center clusters are rock solid; however, your business continuity plans also need to take into consideration Local Area Network (LAN) connectivity, SAN connectivity, storage design and eDirectory design.

Your main goal for LAN connectivity in your clusters is to protect your heartbeat process to avoid false split-brain scenarios. The heartbeat process basically consists of a node sending out a ping to let other cluster members know that it is running as expected. If connectivity problems prevent a node's heartbeat ping from reaching other nodes in the cluster, the cluster must decide that the node is not functioning and must cast that node out of the

cluster, moving its cluster resources to another functioning node. But if both cluster nodes happen to stay alive and try to assume responsibilities of all the cluster resources, it results in diverging sets of data on those servers and even data corruption on the shared disk.

To ensure that heartbeat pings reach other cluster members and avoid unnecessary resource migrations caused by split-brain scenarios, Novell employs a patented split-brain detection method that uses both LAN- and SAN-based communication to determine the true state of cluster nodes. So, even if a server loses LAN connectivity, Novell Cluster Services can still receive heartbeat pings via SAN-based communications and vice versa.

Even though Novell Cluster Services provides multiple paths for heartbeat communications, it still is a good idea to have redundant LAN communication paths between clients and cluster nodes. Also, it's recommended that you have a dedicated virtual LAN (VLAN) and a dedicated IP address range for each cluster. You should also have redundant links between your data center sites. These steps provide additional protection against unnecessary resource migrations and divergent data sets.

In terms of SAN connectivity, you need to ensure redundant access to the Split Brain Detection partition to avoid false SAN device alerts. You also need to ensure redundant connections to each data disk. This might mean connecting cluster nodes and storage systems via two independent fabrics, configuring two paths between cluster nodes and storage systems using native multi-pathing technology or vendor-specific solutions, or having a minimum of two mirror connections between storage systems over different fabrics and Wide Area Networks (WANs).

For storage design, you need to have independent resources for your failover. In a Business Continuity Clustering environment, the Logical Unit Number (LUN) is the failover unit. While it is possible to have multiple storage pools or partitions for each LUN, it is not recommended. The primary storage design principle is to have one storage pool per LUN. Also, as mentioned previously, data must be mirrored between data centers. You can implement host-based mirroring, but storage-based mirroring is recommended. If you use host-based mirroring, make sure that mirrored partitions are only accessible for the nodes of one of the member clusters at any given time.

A major part of preparing for a Novell Business Continuity Clustering implementation will be figuring out what data to replicate. While you might want to mirror all the data from one data center SAN onto another data center SAN, it's typically not feasible financially or operationally. You need to determine how much business data you need to replicate between your data centers, and that begins with an assessment of your various data sets.

You also need to determine the frequency of data replication. Do you need your data mirrored in real time, or is some level of latency acceptable? Your SAN vendor can be a valuable resource in helping you determine the best scenarios for your unique data replication needs.

Basic rules of thumb for real-time mirroring include having link speeds of 1 GB or better, fibre-channel cable lengths less than 200 kilometers between sites, and dedicated links. For distances greater than 200 kilometers, factors to consider include the amount of data being transferred, the bandwidth of the link, and whether or not snapshot technology is being used.

In terms of eDirectory design, while you can have clusters in separate eDirectory trees, the recommended configuration is to have each cluster in the same eDirectory tree. You can greatly simplify cluster administration if you have a separate Organizational Unit (OU) for each geographical cluster (or cluster OU). You should also install all nodes of a cluster into the same container (cluster OU) and place the cluster object in this container as well.

One of your main goals for eDirectory design is to avoid a Novell Directory Services (NDS) Sync state by ensuring direct access to cluster configuration information for each cluster node. You can best accomplish this by partitioning the cluster OU and replicating it to eDirectory servers that hold a replica of the parent partition, and to all the cluster nodes. This will help prevent resources from staying in a Novell Directory Services Sync state when modifications are made to their configurations.

When designing clusters in a Novell Business Continuity Clustering environment, you'll need a unique configuration for each cluster resource. This means making sure all possible IP addresses and volume IDs of a Business Continuity Clustering-enabled resource are unique across all Business Continuity Clustering peer clusters. Each cluster must also have a unique name, even if the clusters reside in different eDirectory trees. Note that clusters cannot have the same name as any of the eDirectory trees in the business continuity cluster.

> **Time to Protect Against Downtime**

Every organization's tolerance for downtime is different, but whether it's thousands or millions of dollars at stake, every second you're down costs you. If you can't risk the lost productivity, sales, customers, partnerships and overall business viability that data center downtime creates, it's time to take a look at the proactive protection Novell Business Continuity Clustering provides against data center-wide failures. **N**

Coming Soon: Preferred Clusters

The next release of Novell Business Continuity Clustering will include a new feature—known as preferred clusters—that allows you to pick and choose where your data center cluster resources fail over. In the current version of

the product, if your data center fails, Novell Business Continuity Clustering simply migrates all of its cluster resources to the next known available cluster. This method works well for most organizations, which typically rely on two-site clusters. But when organizations deploy Novell Business Continuity Clustering in three or more sites, preferred clusters gives them more flexibility. If you have data centers in New York, Boston and Milan and one goes down, preferred clusters lets you decide how and where those cluster resources migrate. For example, if your New York data center goes down, you

can choose to have all the resources go to either Boston or Milan. You can also specify that a specific set of resources go to Boston, while the rest go to Milan. This is especially helpful if your Boston and Milan data centers do not have the same capacity as your New York site. The next release of Novell Business Continuity Clustering will also introduce 64-bit support for x86-64 CPUs. While this delivers increased performance, the biggest benefit is that it enables support for Linux servers running the x86-64 distribution in Novell Open Enterprise Server 2 environments.