PortaSwitch on Multiple Sites

Case study

Improved reliability

Events, such as a data center outage or an IP network routing issue, might render your whole system inaccessible for end users — so your service effectively is down. In order to protect from such infrequent, but catastrophic events and to increase the reliability of your operations, you should deploy a secondary, redundant site.

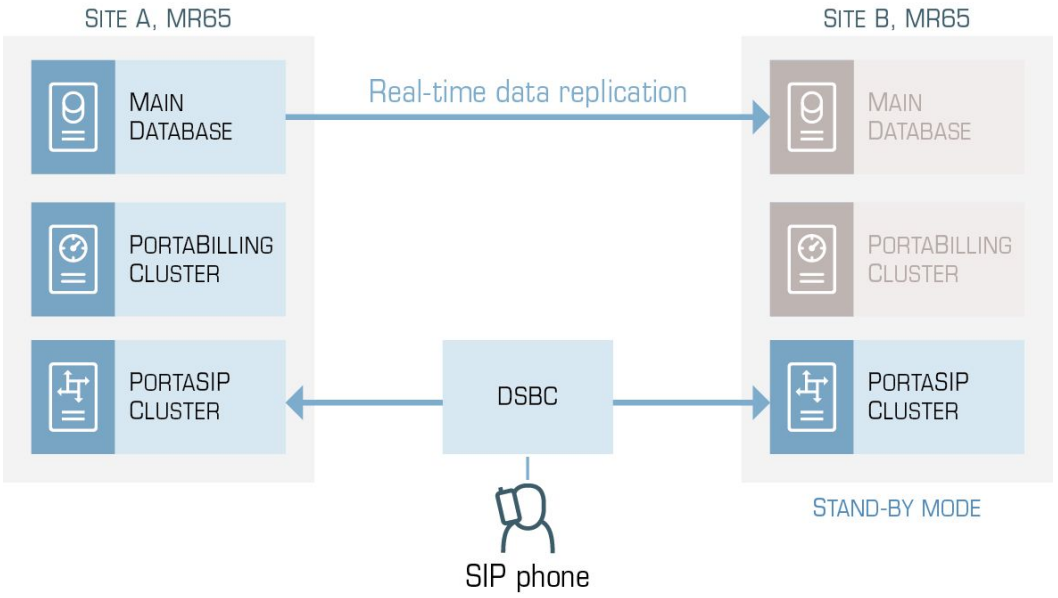

Your redundant site runs the same release as your main site. It has a real-time data replication link with it containing all recent changes in service configuration, up-to-date balances, etc. PortaSIP servers on the secondary site are active, so calls are load-balanced between both sites. Billing servers and database are in stand-by mode. This is done to avoid a “split brain” situation in which the main and secondary site would perform independent, conflicting changes. The same applies to the admin GUI: no access and data modifications are allowed while in stand-alone mode, and the API is available in read-only mode to allow functioning of self-care applications, external portals, etc.

Case 1: Protection against hardware outage on the main site

The main site may become unable to perform the regular activities due to an unlikely (but catastrophic) event such as hardware failure on the master database server or an outage of the internal network on that site. When the secondary site detects that the main site is no longer accessible, it switches to stand-alone mode. Now all activity (calls, messages, etc.) charges are stored in the local database as “deltas”. This allows proper authorization of calls and avoids balance overdraft. When the connection to the main site is restored, the “deltas” are synchronized back to the main database.

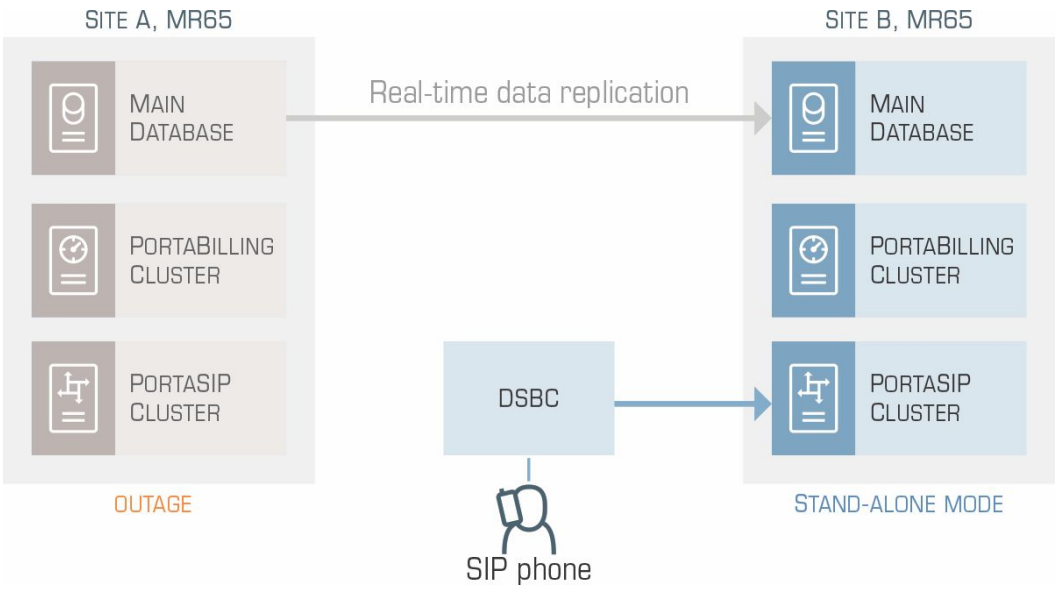

This way your customers do not really notice anything and continue using the service as usual.

Case 2: Protection against lost connectivity between sites

From the perspective of the secondary site, this is identical to Scenario 1 above. The secondary site activates the stand-alone mode. When connectivity is restored, it synchronizes back the charged xDRs and switches back to stand-by mode. To minimize the chance of balance overdraft while synchronization is still in process, transfer of the changes, accumulated on the secondary site, to the main one is done in two steps:

- First, a lock is placed on available funds for total amount consumed while in stand-alone mode.
- Second, detailed xDRs are transferred to the main site (this might take a while).

Case 3: Geographically dispersed sites

Although just one additional redundant site seems to solve the problem of service availability, if you provide service in different countries / regions, you might face network delay issues. For example, if you have 100,000 active users in India (and they make a lot of calls to one another [intra-user]), providing service through a PortaSIP node located in Europe or North America might decrease the sound quality. A better approach is to place additional sites near the areas where your users are congregated. This way, most of the “intra-user” voice media streams will stay local.

Case 4: Zero downtime update

Although you might be lucky and have no data-center outage happen in your life time, there is a definite possibility that the main site needs to be shut down for a planned maintenance (hardware changes, network re-configuration or upgrade to a new software release). Having a secondary site allows you to perform such updates with no downtime noticeable by your users.

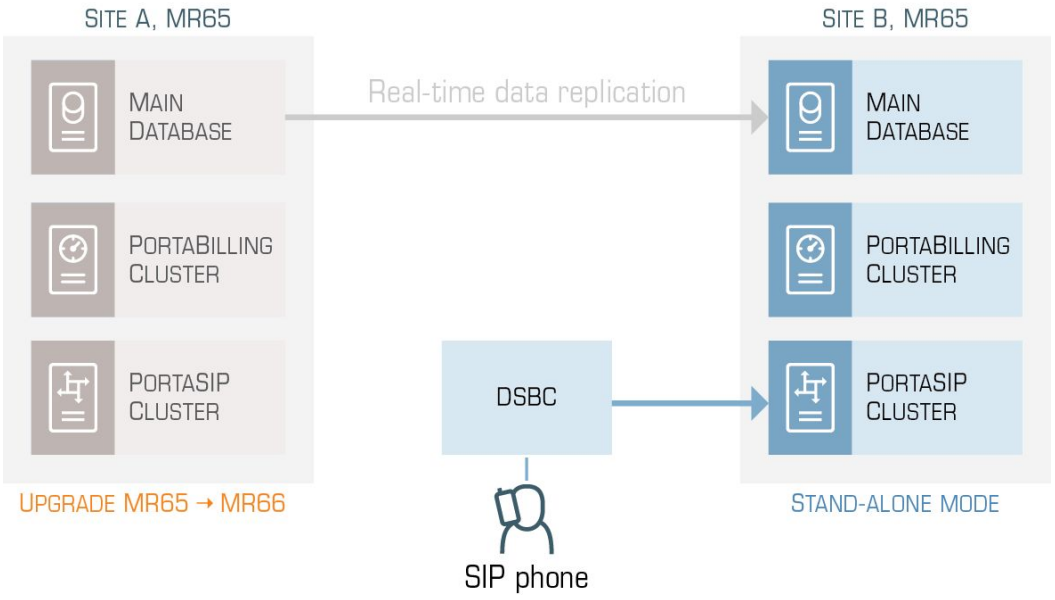

While the main site is being updated, customers continue to use services as usual via the secondary Site B: calls / messages / sessions are charged, balances are updated accordingly and xDRs are written. When the update on the main site is complete, the changes are transmitted to the main site and then the secondary site, in turn, undergoes the update.

Dispatching SBC

If two sites exist totally independently from each other and each has its own IP address, then each SIP endpoint (IP phone, IP PBX, mobile app, etc.) must be able to work with multiple SIP proxies and to make a switch-over if one of them is not available. There are multiple methods to make SIP phones and gateways aware of multiple SIP proxies (DNS, DNS SRV, etc.). Unfortunately, different vendors of endpoints have different particularities regarding implementation, so there is no guaranteed-to-work method to work across all devices.

In addition to that, a normal scenario these days for a hosted IP PBX user is to have his / her IP phone installed on a private network behind a NAT / firewall. If an IP phone registers to the IP address of Site A, the NAT traversal tunnel is open to that IP address. When that site goes down (hardware issue or scheduled software upgrade), the phone cannot receive incoming calls until it re-registers to another "live" site.

Dispatching SBC (DSBC) makes it possible to efficiently overcome these problems. It provides a single entry point for IP phones of end users and hides the actual site topology: even the "dumbest" VoIP device can be configured with a single SIP proxy address. It keeps the NAT traversal tunnel open and available for any site, so phones can receive calls even during a site outage or maintenance.

DSBC should be installed on a separate set of hardware to keep it isolated from the actual "sites". A minimum of two servers are required (three recommended) to create a proper high-availability cluster.

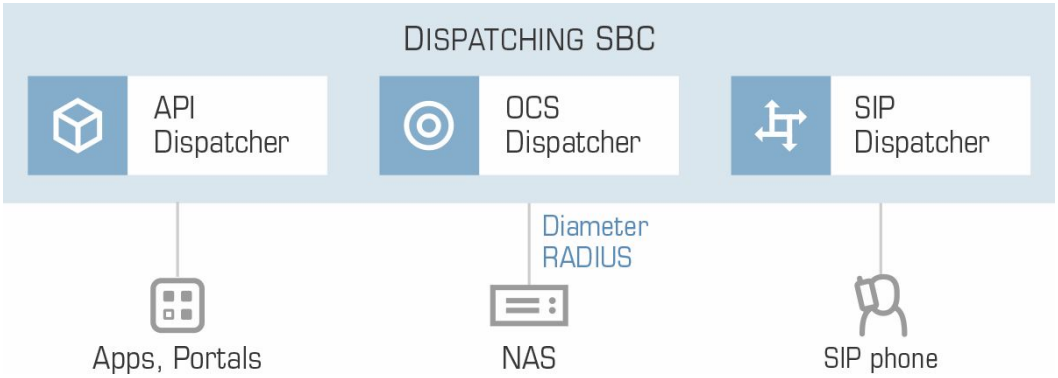

An alternative is to run DSBC in the cloud. Since DSBC only forwards SIP signaling messages (but not the actual voice media), resulting in a miniscule delay (50-100 milliseconds that it takes for an IP packet to travel across the network) that has no impact on call quality and is not noticeable by end users. An actual PortaSIP server on one of the sites handles the call.

Gradual software upgrade with an "alter-ego" site

Dual-version PortaSwitch® helps service providers balance between having access to new features and securing the overall stability of their existing platform. It comprises two sites: one of current release (e.g., MR55) and the other, the "alter-ego" of the new release (e.g., MR70) which operates in parallel. Each site keeps its own set of customers.

Case 5: Gradual migration

The "alter-ego" site, which runs the newest release (e.g. MR70), starts as a full copy of the production site. A snapshot of the current release (e.g. MR55) is copied to the "alter-ego" site and the data is converted to the new release at the database level. Thus, all products, tariffs, rates and other configuration objects required for service provisioning are pre-created and ready-to-use on the new site.

Then a "pilot" batch of customers can be migrated to the "alter-ego" site using the PortaSwitch® migration utility called Porter. Porter copies only customer and account records and their

related service configurations, such as service features and IP Centrex configurations. The xDRs produced for customers on the production site are copied too, thus keeping billing data in sync and preventing money loss during migration. Customer data is transferred automatically, without the involvement of an administrator. During migration, Porter operates directly with the database. This and the already pre-created customer / account dependencies minimize any service disruption time a customer might experience (i.e. the time during which a customer's services are stopped on the production site until resumed on the "alter-ego" one).

To illustrate, we expect a residential customer with 4 accounts and the basic service configuration to be migrated in less than 10 seconds. The exact migration time depends on the amount of data to transfer: the number of accounts, their service configurations and the number of xDRs produced for them once the database copy is made.

All customers connect via Dispatching SBC (DSBC), so migration is seamless for them as their devices don't have to be reconfigured. DSBC transparently redirects requests (SIP, Diameter, API, etc.) to the system on which the customer is currently hosted.

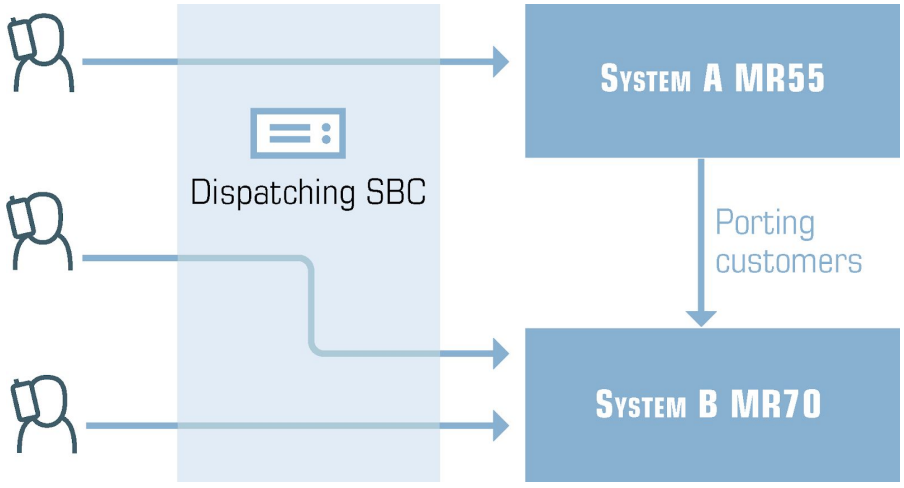

You can observe the migrated customers and verify that everything runs smoothly and that customers are happy with the service (changes, new features, compatibility with various IP phones) on the new system. Ideally, customer accounts should reside on the "alter-ego" site for at least one full billing cycle. After this recommended minimum of one full billing cycle, you can proceed with customer migration in batches (potentially increasing the "batch" size).

Case 6: Dual-version with redundancy

When upgrading via a "dual-version" setup, you control the time and pace of customer migration. Thus, it can take a while (e.g. several months) to migrate all your customers to the new release. During that time, it is important that both "production" and "alter-ego" systems are secure against data-center outages or other unforeseen events discussed earlier.

The best configuration is to build each system as a redundant and part of 2+ sites. What we call a "system" is a combination of 1, 2 or more sites operating in sync; that is, running on the same software version and same configuration, and sharing the same set of customers. Of course, any of those sites may run in the cloud.

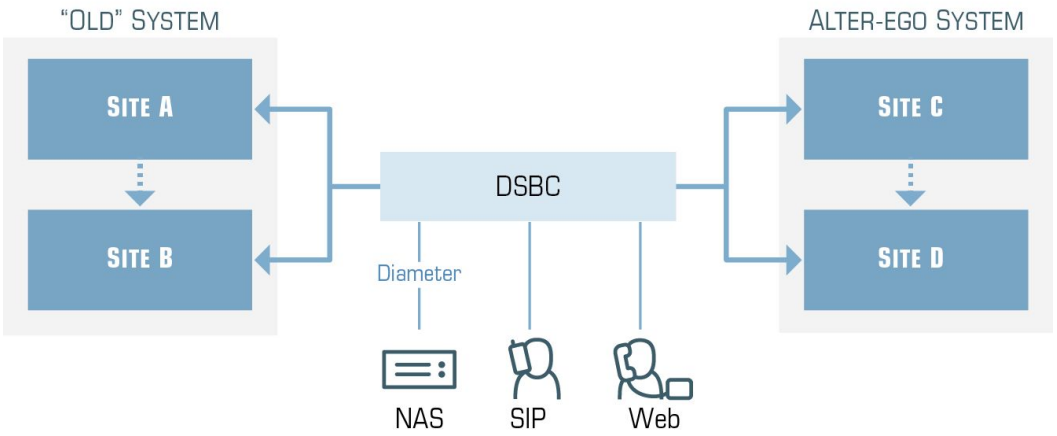

Conclusion

To avoid catastrophic events, such as a data center outage or an IP network routing issue rendering your whole system inaccessible for your end users and your service effectively down, you should deploy a secondary, redundant site.

An extra site for your PortaSwitch installation allows you to either:

- create a high-availability configuration and better network utilization by acting as a redundant site,
- perform a gradual upgrade to a new software version by creating an "alter-ego" (independent) system, which runs on a different software release and has its own subset of customers.

These approaches can be combined creating a "production" – "alter-ego" tandem where each system contains multiple redundant sites. DSBC provides a transparent attachment for the customer to the system on which he / she currently resides.