

ATLAS

A High Availability Server Designed Especially for Small Business

In most businesses, workflow processes rely on one or more computer servers to provide essential services such as document management, database access, or email. When a server fails or is otherwise unavailable, workflow processes fail, causing reduced productivity and lost profits. For these reasons, large businesses often employ High Availability Servers (HAS) to minimize server downtime.

These HAS solutions fall into two categories, (1) high availability clusters (HA clusters) and (2) non-stop computers – both of which require resources typically unavailable to small businesses. HA clusters are much less expensive than non-stop computers to purchase, but require substantial ongoing IT resources to ensure that they will work correctly at “failover.” Non-stop computers require minimal ongoing IT resources to ensure correct operation at failover, but are much more costly to purchase. These costs multiply when more than one server is required to support the business services.

What has been missing is an HAS solution that is cost effective for small business, one that is reliable, affordable up front, and requires minimal ongoing IT resources.

ATLAS is Designed Specifically to Fill This Need

ATLAS (patent pending) is designed to achieve three essential requirements:

- ❑ Minimize cost.
- ❑ Minimize required IT support.
- ❑ Maximize confidence that the system will work correctly at failover.

These requirements are met through two fundamental design goals:

1. Remove all unnecessary complexity.
2. Provide the necessary mechanisms to ensure correct operation at failover.

Goal #1 is realized through a “reduced complexity architecture” with manual failover, which requires a shutdown and reboot. This approach requires fewer special purpose software modules, with fewer aspects to configure, test, and maintain over time. This reduces the required ongoing IT support, while increasing confidence at failover, since there are fewer things (and permutations) that can go wrong.

Instead of this *unnecessary* complexity, ATLAS provides the *necessary* mechanisms to ensure that the system will work correctly at failover – goal #2 -- through the ATLAS Sentry software suite. Sentry Synchronizer maintains compatible environments on all servers while Sentry Verifier continuously exercises and verifies all services on all servers.

Goal #2 must be addressed with *any* HA cluster. If not, one cannot possibly know what will happen at failover. With existing HA clusters, this is left to the end user, which mandates ongoing, periodic reevaluation and testing during off-hours. With ATLAS, this is handled automatically, relieving the end user of this substantial, continuing burden. The result is a reliable HAS solution that is both affordable and that requires minimal IT support.


An Overview of ATLAS

ATLAS (patent pending) is a cluster of two commodity servers (a primary and a secondary) with front mounted swappable RAID disks, connected to the existing business network. The primary runs all critical services (e.g., web server, email server, database server, file server) and stores all critical data.

If a problem occurs, one simply shuts down both servers, swaps the RAID disks from the primary to the secondary, then boots the secondary. The secondary now boots up as the primary, running all services and utilizing the same data and configurations as the original primary. We refer to this as “manual-failover.” The primary can then be serviced without the drawback of intense time pressure to restore the system. This procedure can also be used to perform planned maintenance without requiring off-hour service or downtime.

Fundamental to ATLAS is its “reduced complexity architecture” with manual-failover. Compared with existing HA clusters, this has fewer options and permutations to configure, maintain, and test over time, which reduces the required ongoing IT support. This also decreases the likelihood of problems at failover, since there are simply fewer things that can go wrong. Moreover, the cold restart ensures that all services start in the same manner and with the same system state as they did originally.

Also fundamental to ATLAS is its Sentry software suite. This ensures correct operation at failover by performing two critical tasks: (1) Sentry Synchronizer maintains synchronization between the primary and secondary, and (2) Sentry Verifier continuously verifies that all services run correctly on the secondary, as well as on the primary. These essential mechanisms relieve the end user of this substantial, relentless burden that exists with existing HA clusters. That is, ensuring that it will work correctly at failover. This significantly reduces the amount of ongoing IT support required, while increasing confidence at failover.



ATLAS Reduced Complexity Architecture

Manual failover provides a safer and simpler solution:

- ❑ It eliminates the need for (1) heartbeat software, (2) multiple redundant heartbeat links, and (3) auto-failover software that would react to the heartbeat failure.
- ❑ The cold restart feature ensures that, at failover, all services start in the same manner and with the same system state as they did originally.
- ❑ No shared disks are used. This eliminates the possibility of user data loss due to “split brain”, and the need for software to guard against it.
- ❑ It reduces the possibility of transient problems and increases confidence at failover.

Minimal data mirroring is required:

- ❑ No user data is mirrored. Only configurations and Sentry Verifier data are mirrored.
- ❑ This allows the use of simple, robust mirroring techniques, as opposed to more complex methods required for large amounts of user data that change rapidly.

Load sharing is disallowed. Thus, at failover, all services start in the same manner and with the same system state as they did originally.

ATLAS ensures that software updates do not subvert the goals of ATLAS. For example, an update to the Secondary will not cause the Primary to malfunction. If a problem does occur somewhere in the system, it can be quickly restored to a previously working state, since:

- ❑ Multiple versions of ATLAS software modules can be simultaneously installed.
- ❑ Any installed version of a module can be easily selected to run.
- ❑ All module versions are both backward and forward compatible.

ATLAS Sentry Synchronizer

A fundamental problem with any HA cluster is that the server environments will diverge over time, resulting in different combinations of hardware, firmware, device drivers, patches, application software, or even operating systems. How can one know if the services running on the primary will run correctly in the different secondary environment, or if they will even run at all? Unless mechanisms exist to (1) maintain compatibility between the primary and secondary, and (2) continuously verify that services will run correctly at failover, one cannot have confidence in the HA cluster. With existing HA clusters, this is left to the end user, who must undertake relentless periodic reevaluation and testing during off-hours. This may not be possible or practical. Sentry Synchronizer provides for this, automatically:

- ❑ It ensures that the primary and secondary environments are tracking, even across different versions of software and configurations.
- ❑ It ensures that Sentry Verifier test environment correctly exists on all servers.
- ❑ It manages all configuration changes in a version control system.
- ❑ The user manages only the Primary. Secondaries are automatically synchronized.

ATLAS Sentry Verifier

Sentry Verifier is responsible for continuously exercising and verifying that all services run correctly on all servers. Anything that would cause a service to run incorrectly will be detected and reported, regardless of the cause. This includes hardware or network faults, software incompatibilities, configuration problems, or even a malfunctioning Synchronizer.

The ATLAS Design ensures that, after failover, all services start the same way, and run in the same environment as they did before failover -- *the same environment in which they were exercised and verified until failover*. This is a very significant point. The continuous exercising and verification by Sentry Verifier would be of diminished value if, after failover, services had to start, restart, or run in a different environment – as is the case with warm auto-failover or load sharing.

Comparison of ATLAS and Existing Products in terms of Failure Modes, Cost, and Function

Comparison of Failure Modes for All Options

Result Failure	Hot Swap Server	Non-stop Computer	HA Cluster	ATLAS
Power supply fails	UP	UP	UP	UP
Fan fails	UP	UP	UP	UP
Disk fails	UP	UP	UP	UP
CPU fails	DOWN	UP	UP	UP
Memory fails	DOWN	UP	UP	UP
Motherboard fails	DOWN	UP	UP	UP
I/O Bus fails	DOWN	UP	UP	UP
Disk controller fails	DOWN	UP	UP	UP
Net interface fails	DOWN ¹	UP	UP	UP
Other hardware fails	DOWN	UP	UP	UP
Planned maintenance	DOWN	DOWN	UP	UP

Note 1: Can be UP if it is possible to configure with multiple NICs and failover.

Comparison of Cost and Function of High Availability Options

Result Cost, Function	Non-stop Computer	HA Cluster	ATLAS
Cost (system)	High	Low	Low
Cost (IT resources)	Low	High	Low
Type of failover	Auto	Auto	Manual ³
Danger of split brain?	No	Yes ²	No
Confidence at failover?	Yes	No ¹	Yes
Verification of services?	No ¹	No ¹	Yes
Synchronization?	Yes	No ¹	Yes
Hardware failures	UP	UP	UP
Planned maintenance	DOWN	UP	UP
Load sharing possible	No	Yes	No ³

Note 1: Possible only if the end user provides for this.

Note 2: Potential loss of data.

Note 3: These increase confidence at failover.

ATLAS continuously synchronizes the primary and secondary environments, and verifies that services will run correctly at failover, thereby removing that substantial, ongoing IT support burden. In addition, the ATLAS “reduced complexity architecture” with manual-failover increases confidence at failover, while further reducing the IT support requirements. As a result, ATLAS provides the high confidence level and low IT overhead of a non-stop computer, while also supporting planned maintenance, at an affordable price.