Redundancy Recommendation

T. Andrew Black

Writer's comment: One of the assignments in English 104A was to write a recommendation for improving a product or service. I had originally co-authored a document with a similar purpose for my job last year, in which I assessed an e-commerce client's website uptime vulnerabilities and gave recommendations on improving reliability. Since I wanted the assignment to have practical applicability, I saw this assignment as an opportunity to try to improve upon my previous effort and sharpen my technical writing skills. I started with a clean slate, and with Dr. Clarke's guidance, the result was a much clearer, concise document describing a technical solution to a real-world problem.

—Andrew Black

Instructor's comment: When I asked my English 104A students to write a feasibility / recommendation report, many wrote fine papers evaluating pieces of computer equipment, wireless phone services, and copy machines. Andrew Black, though, took a different tack. Using his considerable technical expertise, he chose to recommend not a product but a design for improved server reliability. What struck me as I read his report, first in draft and then in its final form, was the care with which Andrew addresses his client's needs. His careful assessment of the risks of server and computer failure, his cost analysis, and his solutions to the technical problems he describes make it clear that he is interested in serving his client as much as in making a sale. But his interest in his audience's needs goes beyond an interest in their technical needs. I knew when Andrew asked how he might improve on an already-excellent draft—and when he refused to accept my repeated assurances that his draft was already very, very good—that he would bring the same professionalism and care to his writing that he does to his more technical endeavors. When you read the report that follows, I'm sure you'll agree.

—Marlene Clarke, English Department

SpecialtyGoods.com Redundancy Recommendation

Evaluation and Recommendation for Improving the Reliability of SpecialtyGoods.com

PREPARED BY T. ANDREW BLACK, 10/26/02

1.0 Background

SpecialtyGoods.com, an online retailer of fine specialty goods, has been operating since Spring 2000. In Fall 2001, SpecialtyGoods replaced its web site servers to improve performance. Although the web site has not experienced performance or reliability problems since the upgrade, CEO John Silverman has requested that our company, Technology Analysts, Inc., investigate the site's current risk of downtime (time that the server is off due to malfunction) and provide recommendations to lower the risk. He has made it clear that any significant downtime could significantly threaten the young company.

This document provides an overview of the current risk of downtime, options for improving reliability, and our final recommendation.

2.0 Summary of Findings

The current level of risk of downtime is not high. It is on par with the risks experienced by similarly-sized e-commerce enterprises. However, were a problem to occur that affected one of the servers, the web site would go down, and depending on the problem source, it could take hours or days to restore.

Our recommendation is to install a new web and database server that would be on standby in the server rack (Solution One below). This solution is the most cost effective option. It also minimizes costly downtime. If a problem were to occur with one of the live servers, the appropriate backup server could be started, information could be restored, and the web site could be brought live within 6-8 hours.

The total initial investment for this solution is $70,927–$74,927. Annual recurring costs would rise by only $377 (about 1.5%). The entire solution can be implemented in two to three weeks.

Some of the concepts and information in this document were adapted from a document co-authored by T. Andrew Black and Robert Walling in August 2002. SpecialtyGoods.com and Technology Analysts, Inc. are fictitious names representing real companies.

3.0 Current Infrastructure and Risk Assessment

3.1 Current Infrastructure

The current infrastructure consists of a router, one live web server, one staging web server, and one database server (seeFigure 1). The router connects the internal servers to the Internet Service Provider's network (Verio), which provides multiple redundant connections to the Internet. The web server hosts the web pages, objects and web site software, while the database server hosts the database and the database software. The staging server is a low-end PC that operates as a test web server. When changes need to be made to the web site, they are first posted to the staging server, then tested and transferred to the live web server.

3.2 Current Infrastructure Risk Assessment

Risk of Downtime. The current web server infrastructure is relatively robust by small online merchant standards. All hardware is brand-name and is served under 3-year on-site warranties. The router is a very reliable Cisco unit and is very unlikely to fail. In the case that it does fail, however, it is served under the Cisco support agreement with a 4-hour maximum response time.

Both the web and database servers are manufactured by Dell, which is one of the most reliable brands in the industry. These servers are equally unlikely to fail. The most likely cause of server failure, however, is hard drive failure.[*] Each server has two hard drives, which are arranged in a RAID configuration. Under this scenario, if one hard drive fails, the other takes over automatically with no data loss, requiring no server downtime. Memory failure and other component failures are possible but are unlikely. Therefore, the possibilities are slim that one of the servers could go down at any point due to component failure. However, if the memory or other component besides the hard drives were to fail, the entire web site would go down.

Software can also become corrupted and fail. However these glitches are relatively easy to diagnose and fix remotely. We also now reboot the servers weekly and update all software with the most recent upgrades and patches to help alleviate this risk.

[*]Although manufacturers do not provide data on failure rates of hard drives, an industry benchmark is 350,000 MTBF (Mean Time Between Failures). This means that if large quantities of the hard drive were operated concurrently, the drives would amass 350,000 hours of operating time among them before the first drive failed. In practice this is hard to translate into an actual failure rate, but clearly failure rates are very low (when operated under the correct conditions). Information from http://www.storagereview.com

Figure 1: Current Infrastructure Diagram

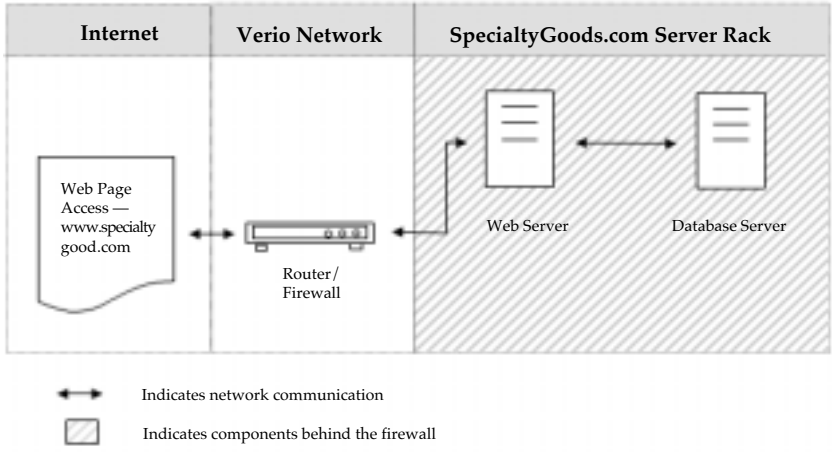

Recovery Time. Estimate: 6 hours to 3 days. Were any hardware component on the web servers to fail, we would first have to notice the problem and determine that it is a hardware problem. We would then need to call the appropriate support team (Cisco or Dell), both of whom have a 4-hour maximum response time. Once the support team arrived they would need to diagnose the problem. If it were not a hard drive problem and they had the appropriate parts with them, they would be able to fix the problem immediately. It is possible that they would have to order parts, which could take as long as two days (assuming the needed parts are not on backorder).

If both hard drives on the servers were to fail at the same time, the downed server would need to be completely rebuilt (operating system reinstalled along with all software and data from backup). This could take up to two days.

If the router were to fail, it would require two to four hours to reconfigure.

Cost in Case of Failure. Estimate: $13,500 to $168,000. During the peak holiday season, the site can produce up to $75,000 per day in sales, but most days averages $50,000. Although most of the traffic comes between 9AM and 9PM, since we cannot anticipate the time of downtime, we will

assume an even sales traffic flow and estimate the sales loss per hour at $2,100.* SpecialtyGoods would also incur other costs in handling and appeasing upset customers, holding inventory longer, and having customers order over the phone (which costs more than processing an order online). Therefore we estimate total loss per hour at $2,250.

In addition to the cost of lost sales, SpecialtyGoods would need to pay to have the servers rebuilt in case of hard drive failure. Each server would cost between $4,000 and $6,000 to rebuild. It would cost about $1,000 to reconfigure the router should that fail.

Maintenance Costs. Estimate: $24,300. SpecialtyGoods.com currently spends approximately $14,300 in annual software licensing agreements. Additional maintenance costs vary greatly by frequency and responsibility (some maintenance tasks are performed inside the company and some by outsourced firms), but we estimate the total to be around $10,000 per year.

4.0 Alternative Solutions

4.1 Solution 1: Duplicate Servers on Standby

System Details. This solution consists of adding two servers to the existing infrastructure: one web server and one database server (see diagram 4-1 below). Both servers would be fully configured, but neither would be turned on (to lower software licensing costs).

Risk of Downtime. The risk of downtime is the same as with the current solution. Although it is low, a component failure is still possible.

Recovery Time. Estimate: 6-8 hours. Recovery time is where this solution has an advantage over the current infrastructure. If a component relating to either the web or database server were to fail, the web site could be restored quickly. Once the problem was identified, the appropriate backup server could be brought up. After the data was restored from backup, the web site would be live. It should be noted that since backups are made nightly, up to a full day's worth of sales data could be lost.

Cost in Case of Failure. Estimate: $14,000 - $18,750. This would include 6-8 hours of lost sales plus the billable time to get the data restored and servers online.

*This estimate was reached by dividing $50,000 by the 24 hours in a day.

Figure 2: Solution 1—Infrastructure Diagram

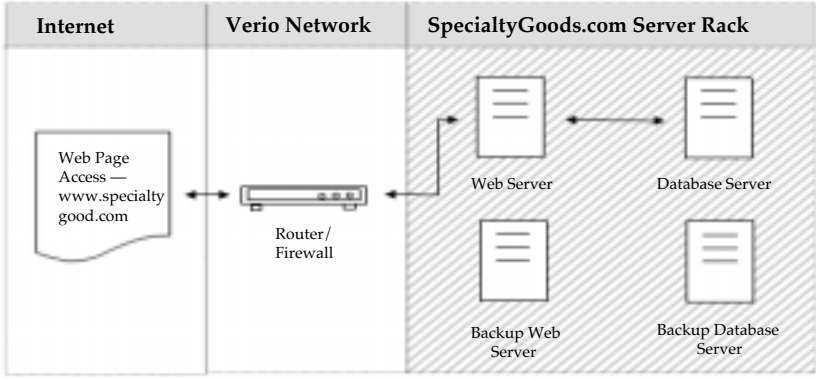

Solution Cost Estimate. The following is a table of estimated costs to implement this solution.

Table 1: Solution 1—Estimated Costs

INITIAL INVESTMENT (in dollars)

	Web Server	Database Server	Total
Hardware	5,150	4,700	9,850
Software	377	52,700	53,077
Labor*	4,000–6,000	4,000–6,000	8,000-12,000
Total	9,527–11,527	61,400– 63,400	**70,927-74,927**

ANNUAL RECURRING COSTS** (in dollars)

Software Licensing	**14,677**
Maintenance	**10,000**
Total Recurring Cost	**24,677**

* Labor costs do not include on-site network setup and ISP costs/fees

** Do not include monthly ISP fees; those will remain the same

4.2 SOLUTION 2: CLUSTERED ENVIRONMENT

System Details. This solution consists of adding a web and database server to the existing infrastructure along with an upgraded router/ firewall. Unlike Solution One, under this solution both servers would be powered on and active. The new router is actually made up of two

routers so that if one fails, the other takes over. Web site data requests would be routed to the appropriate server from the router based on current load and availability, ensuring that all servers get a similar number of requests (this is called "load balancing"). Load balancing increases performance by distributing the traffic load. It also creates a "redundant" network so that were one server to go down, the web site would remain available because the router would send all requests to the appropriate server. This redundancy makes the infrastructure extremely reliable.

Figure 3: Solution 2—Infrastructure Diagram

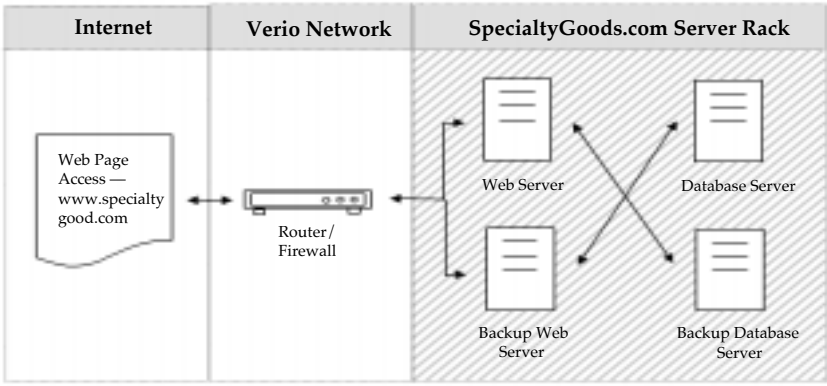

Risk of Downtime. There is no practical risk of downtime from component failure.* Every component is redundant on the servers and router / firewall. Downtime risk is limited to Internet connection failure from Verio or natural disaster.

Recovery Time. Estimate: None. Since all data is actively shared between the two database servers, if one goes down, there is no loss of data. If one web server or one database server goes down, the other would handle the entire load with no interruption of service.

Cost of Failure. Estimate: $0-$2,500. Since there would be no web site downtime, there would be no lost sales or related expenses. If a server were to fail, we would need to reconfigure it, which would be the only cost ($2,000-$2,500 as noted above). However the web site would not be down during the process.

*Downtime would only occur if two servers of the same type (database or web) failed, an extremely unlikely event.

Solution Cost Estimate. The following is a table of estimated costs to implement this solution.

Table 2: Solution 2—Estimated Costs

INITIAL INVESTMENT (in dollars)

	Web Servers	Database Servers	Router	Total
Hardware	N/A	15,000	11,500	26,500
Software	31,500	93,000	N/A	124,500
Labor*	10,000-13,000	12,000-15,000	Incl.	22,000- 28,000
Total	41,500-44,500	120,000-123,000	19,490	**173,000-179,000**

ANNUAL RECURRING COSTS** (in dollars)

Software Licensing	36,000
Maintenance	15,000
Total Recurring Cost (Annual)	**51,000**

* Labor costs do not include on-site network setup and ISP costs/fees

** Do not include monthly ISP fees; those will remain the same

4.3 CURRENT SOLUTION

A third option is to make no changes to the existing infrastructure. Please see section 3 for information on the infrastructure details and risk of downtime for this solution.

5.0 Summary of Costs

Table 3: Summary of Costs (in dollars)

	Initial Investment	Annual Recurring Costs	System Failure Costs
Solution 1: **Duplicate Standby Servers**	70,927–74,927	24,677	14,000–18,750
Solution 2 : **Clustered Environment**	173,000–179,000	51,000	0–2,500
Solution 3: **Existing System (no improvements)**	0	24,300	13,500–168,000

6.0 Analysis and Recommendation

6.1 Analysis

The current configuration brings little possibility of server failure. The dual RAID hard drives in the current systems allow for uninterrupted service in case of a single hard drive failure (hard drives being the most likely component to fail). Four-hour Dell technical support should allow for relatively quick recovery from motherboard, processor, power supply, or memory failure, the other most likely points of failure. Assuming the worst, however, recovery could take several days, and the financial effects of this downtime could be detrimental. Solutions One and Two protect against these worst-case scenarios.

We then must determine which of these two solutions is best. Under Solution One, if a server were to go down, the web site would go down with it (although the downtime would be minimal: only 6-8 hours maximum). Given this risk of downtime, the core question is whether the cost of the clustered environment is justifiable. Using the cost estimates presented in Section 3.2, the maximum cost of downtime is estimated at $2,250 per hour. Under this condition, the costs associated with Solution One are justifiable if downtime were to exceed 33 hours with the current infrastructure. Further, it is the least expensive option for bolstering redundancy. Solution Two (the clustered environment) is not justifiable based on costs, however, as there is virtually no chance that the current systems would be down long enough to cost SpecialtyGoods in sales the estimated cost of the clustered server environment. The increase in recurring costs is also substantial for Solution Two.

Depending on web site traffic growth, Solution One might be an interim solution if one web server and one database server are not able to handle the load demands. If performance does begin to degrade due to increased demand, the servers and router purchased for the Solution One upgrade could be used to implement Solution Two (the clustered environment), and SpecialtyGoods would incur only the additional costs for labor and software licensing.

6.2 Recommendation

Given the above analysis, we recommend Solution One: Duplicate Servers on Standby. It is the most cost effective option and provides

protection against the most likely points of failure, with a short recovery time in case of failure.

6.3 Considerations

Ultimately, the important decision of which solution to implement is up to SpecialtyGoods. But we hope that this document has helped clarify the current risk, options available to mitigate the risk, and some of the considerations that should be used in analyzing each solution.

Regardless of what solution is chosen, it is critical that we assemble a disaster recovery plan that includes:

- Specific steps to diagnose, fix, and restore the site in case of hardware or software failure
- Emergency contact information for all key personnel and suppliers

The estimated time to complete this plan is 6-8 hours; it can be completed in one week.

Work Cited

StorageReview.com. *StorageReview.com*. 27 Oct. 2002 <http://www.storagereview.com>.