

Using Entire System Management for Disaster Recovery

By Detlaff Ross

Sales Consultant

Software AG Deutschland GmbH

September 2007

Disaster recovery may be much-lauded as a concept, but there are still potential threats that remain ignored.

There will always be disasters that incur losses beyond expectations.

Dear reader, allow us to recapitulate a few of the significant disasters of the past 10 years and discuss the threats presented or even manifested by these. In this article, I shall refrain from addressing topics such as ROI, damaged hard drives, or power outages. Instead, I wish to bring awareness to the potential for complete destruction of data centers, outline the consequences of such an event and, finally, propose a solution. Over the years, I have broached this issue with many colleagues and, most of all, end users. In the course of these discussions, I have consistently detected a small, seemingly harmless, yet extremely critical flaw in disaster planning: Namely, the lack of real-time backup capabilities for operating data.

Just a few years ago, the exceptional Elbe flood threatened two data centers near the Oder and Elbe rivers in Germany. Should the local levees have failed, these centers would have “gone down” in the truest sense of the phrase. Just recently, we heard that a nuclear power plant in Japan was heavily damaged by an earthquake — data centers, on the other hand, seem less likely to make headlines. Since 9/11 at the very latest, we have become aware of the threat posed by terrorist attacks and their potential consequences for state and society beyond the tragic fate of those immediately affected. Of course, data centers are one of the potential targets of such terrorist attacks. Whether by water, fire, terror, or even military means, these threats often incur the largest possible damages.

In the face of such potential total losses, companies require IT solutions that fully ensure that their data is backed up securely. And they often come to us because we have been involved in the IT sector for longer than most. The degree to which an organization is able to absorb losses and gaps along the IT processing chain — before their very existence is threatened — is a decision to be made by corporate directors. In cases of disaster, two factors play a critical role in the interim period until IT operations are back to normal: firstly, the reliability of the emergency recovery procedures and, secondly, how much is known about the key company data at the time of disaster. The higher the quality of the operational data at “failover,” the more secure the disaster recovery becomes and the better the chances of success for an automatic procedure. Ever larger degrees of automation in disaster recovery procedures require commensurate levels of precision in a company’s operating data at time X. As failover procedures become increasingly automated, companies are able to return to “business as usual” quicker than ever.

As a necessary — although not fully adequate — prerequisite for recovering from disaster scenarios such as those outlined above, a company must have backup computing facilities located sufficiently far away from the main operating site. After all, what good is a backup data center situated at the earthquake’s epicenter or in the path of a shock wave that destroyed the

central computing center? Another requirement is to have trained personnel and precise rulebook documentation of failover procedures so that once key data has been recovered, company operations can resume. Here the critical question is: how much documentation is enough? And furthermore, how many man-hours will be required to read and comprehend the documentation and perform the various necessary tests and everything else? In these times of service-oriented architectures, comprehensive interoperability of applications in real time, and integration of systems and applications, it no longer suffices to simply “go get the last backup from the vault” and restore the operating data back to a functioning computer.

Today, the term “real time” has become as integral to IT business processes as it has been for many years in control engineering applications. Integration and interoperability — the conjoining of batch and online processing, or “Batch Application Integration” as formulated by the Gartner Group — require more extensive methods to withstand the threats and hazards posed by the above scenarios.

What is operating data? Data centers utilize various systems for executing job processes according to specific priorities, for controlling started tasks, starting and stopping services and more — and all of these generate operating data. According to the classic definition, these are job scheduling systems specially designed to handle special and/or technical events and map the diverse relationships between processing and services. A second major application for these systems involves event controlling and rule-based post-processing of events — in a nutshell, output management systems. In today's world, the sheer range of potential data sources (computers, applications, individual creation on workplace computers) and output destinations (not only printers but also email, electronic files and dossiers, automatic emails, and so on) requires systems that are much more complex than the data tapes and preprint processing centers of yesteryear. Modern-day systems transmit messages with internal statuses to consoles, log files, and even SNMP ports. Now even though experience has shown that the information content of 99% of these messages is low, there are still some relevant messages that require automatic or manual responses.

The data pool for these controlling and monitoring systems consists largely of process and production rules together with log data of the current production statuses. With most of the systems available on today's market, this data can be found in VSAM clusters, PDS libraries, sequential file systems and, on rare occasions, in actual databases. It is important to note that, in this analysis, I define “databases” to be those systems that contain secure transaction mechanisms with automatic recovery/restart features.

As you can well imagine, all of the above systems are designed to restore all of a company's essential operating data following a catastrophic incident. The backup data center will most certainly have a copy of the rulebook from the previous day or perhaps even a current copy. But it is nearly absolutely certain that the last few minutes of production from the affected data center will not be present at the backup data center. In all likelihood, the time discrepancy between the two will be larger than just a few minutes. For the most part, the distributed operating data is not grouped by transaction and certainly not replicated in real time. Of course, we are talking about extreme disasters here, that is, situations where even the personnel will have to be “backup personnel.”

Readers of TECHniques will be familiar with the real-time data replication features implemented in the Adabas Event Replicator. However, most readers are probably unaware of the fact that Entire Systems Management has been available from Software AG for more than 20 years as a job scheduling, output management, and event management solution. Entire Systems Management has been implemented in Adabas with Natural and EntireX. Because all data is stored in Adabas, all state changes are documented in transactions and persistently replicated to the backup copy.

It is important to note that, contrary to common opinion, the capabilities offered by Entire Systems Management for fully automated production management are by no means limited to Adabas/Natural applications. Several of the major ESM end users in Germany use only ESM and do not — unfortunately — make any use of the Adabas/Natural applications. After all, ESM is not at all restricted to just mainframes, as Adabas, Natural, and EntireX are also available for other platforms.

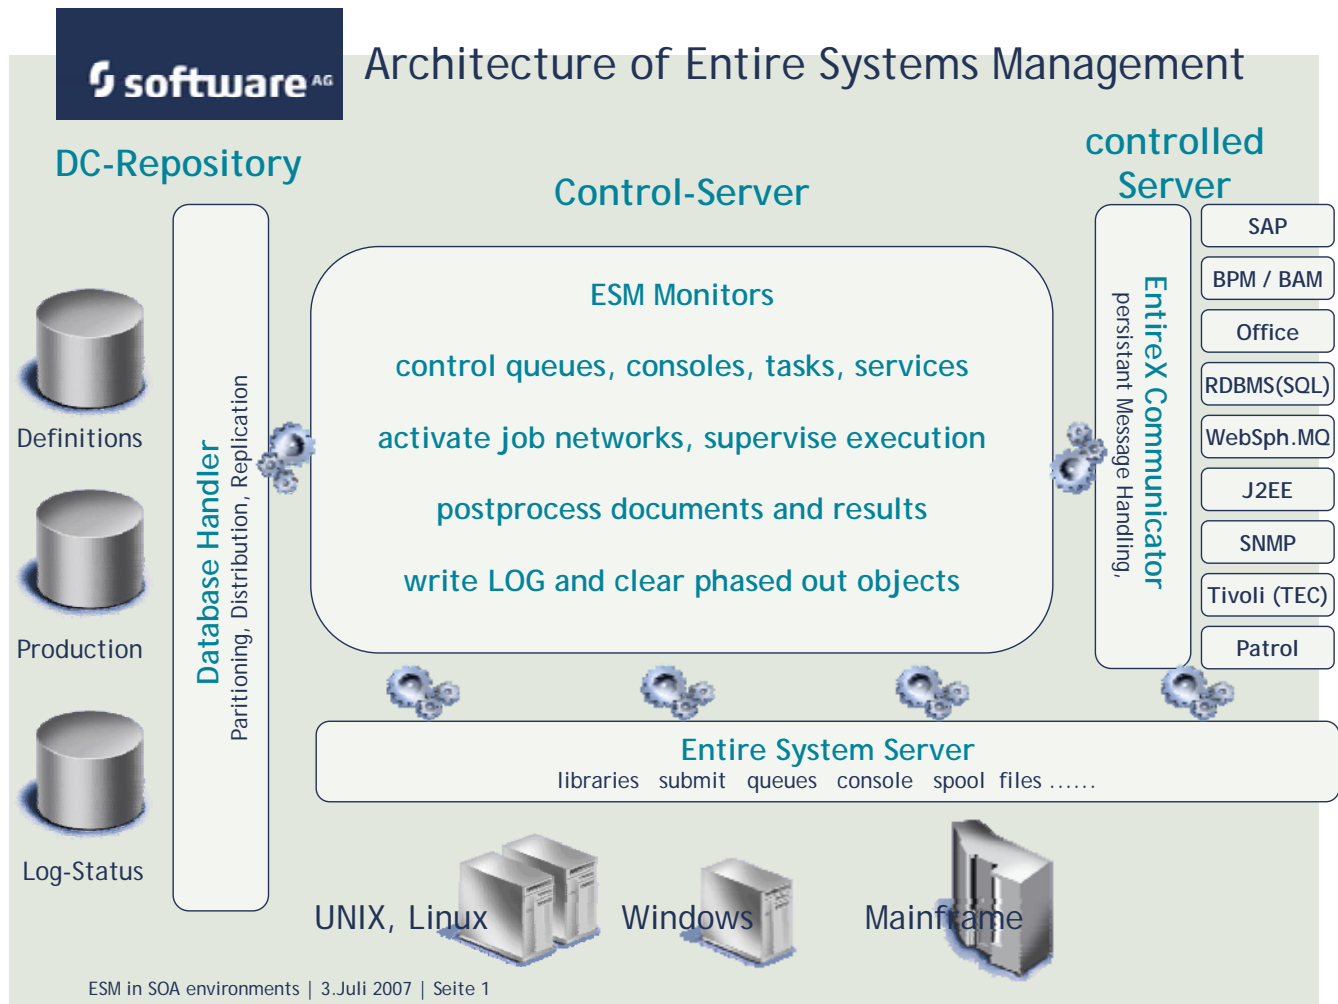


Figure 1 - ESM architecture overview

Of course, many may well be asking how ESM and the Adabas Event Replicator together can restore operating data in real time following a disaster. In order to answer this question, a brief excursion into the theoretical background of these systems is necessary (see also Figure 1). All of the following rulebooks are managed within Adabas with particular emphasis on JCL and scripts:

- Job networks, jobs, dependencies, schedules
- JCL/scripts and configuration settings
- Post processing rules for lists (separation, bundling, formatting, printing, distributing, archiving)
- Messages, message filters, event reactions

Together, these are referred to as the definition database. Whenever rules (job networks, separating/bundling, event trees, etc.) take effect according to schedules and/or events, a copy of the rules is stored together with the current values in the production database. This database documents the following:

- Current statuses of job networks, jobs, and current values for their interdependencies

- JCL to be executed by computer and/or current versions of scripts with evaluated parameters
- Status data from post processing to spool containers
- Relevant messages and events filtered from the stream, together with corresponding reactions or planned reactions
- Event trees and dependency statuses in job networks
- Planned processes and/or expected events for the immediate future

The LOG database contains the history of all changes made to production statuses, together with time stamps, causes, computer nodes, messages, and so forth. The entire system has been implemented as a finite state machine, and the automation rules utilize the log data as a source of information for both the past and the future. Of course, these can also be accessed by any (authorized) auditor (key words: SOX, ITIL, etc.). I have met auditors who have been very pleased that ESM reports can be generated for IT production processes using the correct format and terminology. Please excuse this brief digression, this was only intended as an aside.

The ESM data in Adabas is effectively the operating data repository. As described above, this data is complete, traceable, and transaction-based up to the very last moment before the total failure. By replicating this data with Adabas Event Replicator via EntireX Node-to-Node to a backup data center, the backup data center will then contain a complete repository of the operating statuses for the lost data center.

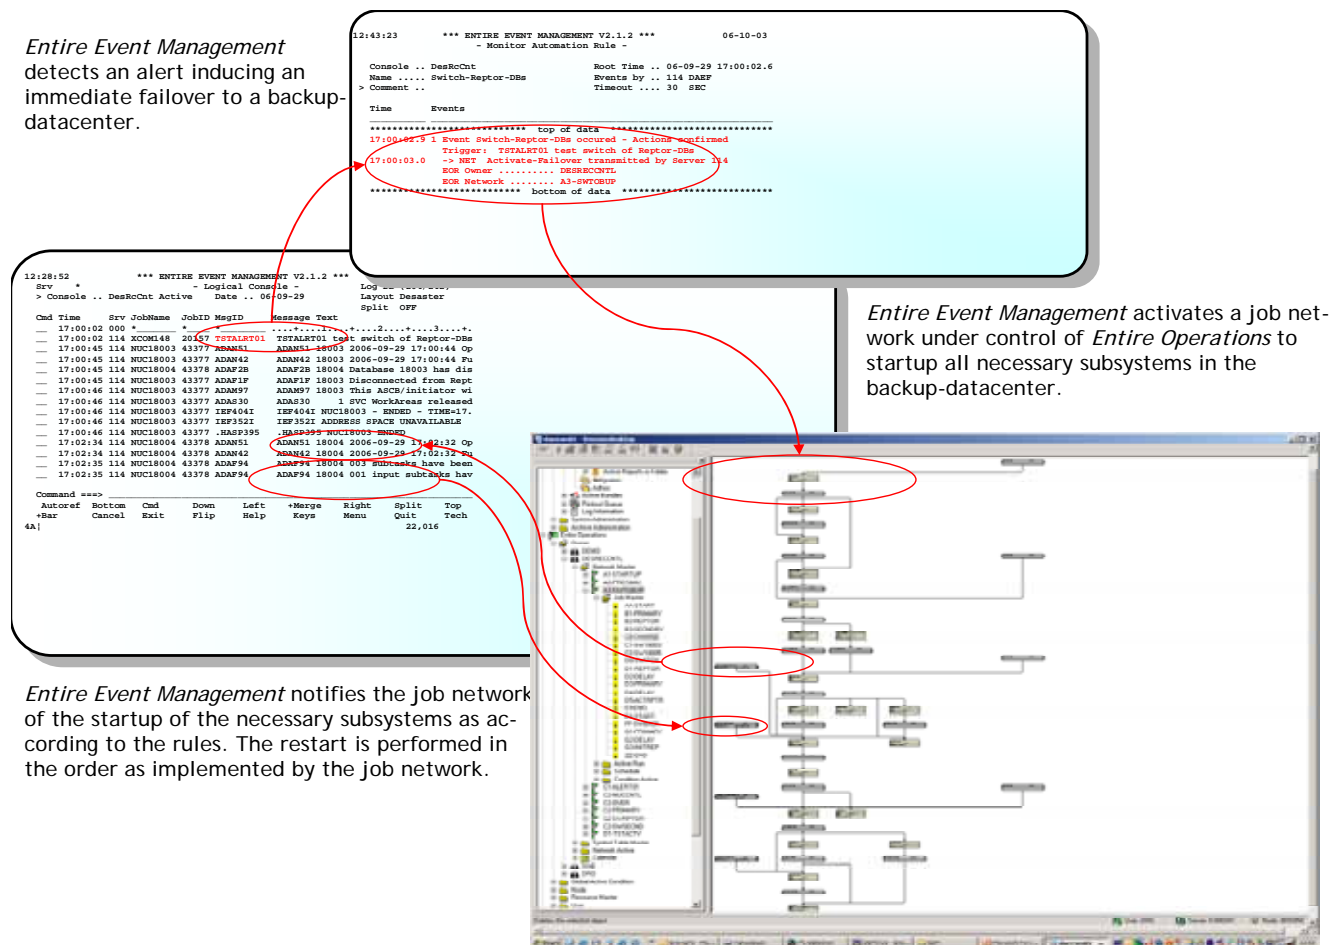


Figure 2 - ESM recovery flow

After monitoring proceedings, the "Entire Event Manager" of the ESM will initiate the restart/recovery process at the backup data center and schedule the required jobs using Entire Operations. In such a scenario, Entire Output Management is very helpful for evaluating job protocols from the lost data center, for they have been replicated. Refer to Figure 2 for a largely simplified scenario.

For attentive readers, there may be one question that lingers: How does the ESM actually control the statuses in operating systems? Long-standing customers of Software AG are probably familiar with the Entire System Server (formerly known as Natural Process). This subsystem makes it possible to control object statuses in operating systems from the "outside" but still from within Natural. It is thus possible to monitor jobs, tasks, and more without having to actually change them.‡

In summary:

1. Entire Systems Management together with Adabas Event Replicator make it possible to implement completely automated disaster recovery processes for operating data centers.
2. Entire Systems Management is compatible with all forms of applications and computer platforms; it is not restricted to Adabas/Natural applications or mainframes.
3. By organizing data centers using Entire Systems Management and securing them against "catastrophes" using Adabas Event Replicator, these centers will also be safe against large and small disruptions in daily production cycles.

*) Many vendors, especially for job scheduling systems, use the concept of "mini-steps." The jobs executed for this purpose generate program calls that transmit control data to the controlling system. However, considering that audit-proof and unalterable management processes are usually required and manual intervention is not desirable, this is a highly questionable methodology.

About the author

Detlaff Ross is a mathematician and sales consultant. For the past 20 years, he has represented the Entire Systems Management products and their customers within Germany and throughout the world. In addition to sales, his activities have included consulting work for customer projects as well as development work in ESM concepts. Please send any comments, thoughts, and opinions concerning this article to [Detlaff Ross](#).