
Gain value from Infor by:

- Delivering continuous access to mission-critical systems and data
- Providing uninterrupted continuity of operations
- Maintaining data integrity

The cost of downtime

There's no doubt about it: system downtime costs money. When your organization has to resort to manual processes during a system outage, you lose the efficiencies you've worked to create. Patient care takes longer, patient charges are less accurate, labor tracking is less precise. And when your system is back on line, there's a duplication of efforts required to synch the backlog of manual data back into your automated systems.

No organization can afford downtime. The exact costs are difficult to pinpoint, and vary by type and size of organization. However, a recent Aberdeen Group research report seeking to quantify the costs of downtime estimated the average cost per hour at \$138,000 for 2012. This amount was up 41% from the findings of a similar study in 2010.*

In healthcare, the costs may be even greater. How do you put a price on patient safety and well-being? How do you quantify the cost when a system outage disrupts the reporting of a lab result or delays a clinician's notes from showing up in a patient's electronic health record?

In healthcare, system failures cost more than money. They can cost lives.

As hospitals and healthcare delivery networks strive to become Accountable Care Organizations, ensuring that patient information is flowing to and through the designated systems is critical. There's zero tolerance for downtime and unplanned failover of systems—especially during peak hours for patient care.

* Aberdeen Group, "Datacenter Downtime: How Much Does It Really Cost?" March, 2012, p. 2, available from Aberdeen Group, accessed Oct. 19, 2012.

There's no predicting the time involved to remedy a server issue, and a failure during peak hours can have devastating consequences. When the flow of data is disrupted, the effect is viral and impacts patient health and safety, internal processes, and revenue.

Downtime: the ripple effect

System downtime affects more than just data availability. When your systems are down, your organization suffers in three main areas:

Patient safety

Downtime for your mission critical applications—such as the ones that run your cath lab or radiology—can result in incorrect diagnoses and compromise patient care and patient safety. Patient safety issues, in turn, can compromise your organization's accreditation standards, hospital ratings, patient satisfaction, and revenue.

Productivity and Labor

When your mission-critical apps go down, by necessity your process reverts to a manual or paper charting of patient data, tests, and procedures. As a result, your medical staff takes longer to provide patient care and the entire process becomes slower and less efficient—not to mention the extra labor involved in entering manual data back into your systems when service is restored. The total effect is a higher cost for you to manage your patients.

Revenue

If your admission, discharge, and transfer (ADT) or supply chain system is down, think of the patient charges that may be lost—and the impact on revenue. Ultimately, downtime of nearly any application or system within your healthcare organization will impact your bottom line.

High availability is no longer a nice-to-have—it's a must-have.

Industry regulations have created a healthcare industry driven by complex IT systems that support and ensure day-to-day patient care and business activities, and downtime greatly impacts all areas of a healthcare organization's operations. There's no predicting the time involved to remedy a server issue, and a failure during peak hours can have devastating consequences. When the flow of data is disrupted, the effect is viral and impacts patient health and safety, internal processes, and revenue.

Standard servers using RAID (redundant arrays of independent disks) storage devices typically provide uptime of around 99% (two nines) and do not come with standard high availability features. A failure on a standard server will very likely result in a loss of data if no high-availability option is added.

Infor Cloverleaf High Availability provides a safety net that allows your team more time to pinpoint, evaluate, and address server issues without data loss—and without impacting your critical systems for business operations and patient care.

Demystifying the 9s

Across all industries, high availability is measured in nines. “One nine” refers to 90% system uptime; “five nines,” a standard reference point, refers to 99.999% uptime.

When speaking the language of nines, it’s easy to lose sight of the real meaning behind these percentages. The table below translates these percentages into the actual costs to an organization, in terms of lost time and dollars.

Of course, even these specifics don’t tell the whole story. Downtime that occurs during peak patient care hours will have a bigger impact to your organization than downtime occurring in off-peak hours.

Number of nines	Uptime percentage	Time lost	Dollar cost*
One nine	90%	36.5 days	\$120,888,000
Two nines	99%	3.65 days	\$12,088,800
Three nines	99.9%	8.76 hours	\$1,208,880
Four nines	99.99%	52.56 minutes	\$131,100
Five nines	99.999%	5.26 minutes	\$12,098
Six nines	99.9999%	31.5 seconds	\$1,208
Seven nines	99.99999%	3.15 seconds	\$121

**Based on Aberdeen Research’s average downtime cost of \$138,000 per hour. Aberdeen Group, “Datacenter Downtime: How Much Does It Really Cost?” March, 2012, p. 2, available from Aberdeen Group, accessed Oct. 19, 2012.*

The Infor philosophy: Downtime is not an option

At Infor™ Healthcare, we believe that information should always be accessible at the point of decision-making. This means having data flowing and accessible, 24/7.

The Infor Cloverleaf Integration Engine is designed to connect the myriad of specialized software solutions used by healthcare organizations today and to ensure that data flows freely and securely among this complex network of disparate systems. Cloverleaf is an industry workhorse in the US and around the world, well-known for its resilience and its ability to integrate complex systems. In fact, it was named Top Complex Data and IT Integrator in a 2012 Black Book study, and was rated number one in security and workflow productivity. (See infor.com/blackbook to download the complete report.)

We offer Infor Cloverleaf in a High Availability configuration to further ensure there are no interruptions in the flow of data and that patient and administrative functions are being performed according to standards. We believe a high availability solution for healthcare should:

- Ensure continual access to your most critical business and clinical applications.
- Provide 99.999% uptime.
- Be an extension of your integration engine or platform.
- Allow for scheduled system maintenance and upgrades without disruption in service.
- Mask planned and unplanned outages by providing continual service.
- Employ fault-tolerance and automated failure detection.
- Reduce the burden on your internal IT resources.
- Work within your budget to achieve your desired level of uptime.

The Infor Cloverleaf High Availability solution in detail

Infor Cloverleaf High Availability works through server clustering. Two servers, or nodes, are linked together and continually communicate with each other via clustering software. The clustering software does a pulse check of the other node, looking for a system failure. If a failure is detected, the applications on the failing node “failover” to the second node and automatically resume operations seamlessly, without any manual intervention. A SAN (storage area network) is linked to both servers and provides access to the shared data required to enable failover.

High-availability clustering is available in a variety of configurations:

- **Active-Active Clustering:** Each node is running production services, while also serving as the backup for the other node, in a configuration known as mutual takeover. In case of failover, the running node becomes double-active.
- **Active-Passive Clustering:** The active node is actively handling Infor Cloverleaf requests and running applications. The passive node is idle and on standby, waiting for the other node to fail.
- **Active-Test Clustering:** The primary node is hosting the production instance of Infor Cloverleaf, and the secondary node is hosting a test instance of Cloverleaf.
- **Geo-Clustering for Disaster Recovery:** In this scenario, the cluster is configured across two data centers, providing a greater level of availability and protection by having a second copy of the data at a second physical location. The second location can accept an automated or manual push of data from a production instance of Infor Cloverleaf from the other location.

Infor Cloverleaf High Availability can also be configured in a virtual mode, allowing your organization to run multiple virtual machines on a single physical machine, sharing the resources of that single computer across multiple environments.

What does an Infor Cloverleaf High Availability solution include?

- Hardware, software, and network infrastructure recommendations
- SAN replication software
- Clustering software
- A standby instance of Infor Cloverleaf

Reduce your organization’s risk

Prevent downtime and data loss, and gain peace of mind, with Infor Cloverleaf High Availability.

Infor Cloverleaf High Availability:

- Allows for scheduled system maintenance and upgrades without disruption in service.
- Eliminates single points of failure by providing the application in a redundant, automated clustered environment.
- Automatically detects hardware failure and routes transactions to the backup system with “near real-time” failover.
- Provides data protection by tracking transactions in a recovery database—allowing the system to pick up where it left off in the event of a system failure.
- Provides the desired level of uptime for your active Cloverleaf environment in keeping with your budget.
- Responds seamlessly to unplanned system outages.
- Scales easily to meet your volume requirements.

Infor Cloverleaf High Availability in action

A 452-bed, non-profit healthcare organization in the southern United States uses Infor Cloverleaf to connect 95 percent of their systems, including hospital EHRs.

They implemented Cloverleaf in a High Availability, active-passive configuration, where one server is hosting Cloverleaf and a second server serves as the standby node, waiting to take over in case of a failover.

When a system failover occurred at 2:30 a.m. local time, here's what happened:

- Application data files automatically moved from the SAN to the standby server
- Application IP addresses automatically moved to the standby server
- Application automatically restarted on the standby server
- Users may have seen a brief interruption in service requiring a refresh or reconnect
- System automatically generated an email to the integration team, alerting them to the failover
- Integration team called Infor support to validate that the system resumed normal operation following the failover

Zero tolerance for downtime

A private, 506-bed, not-for-profit hospital on the East Coast implemented Infor Cloverleaf more than ten years ago.

Today, Cloverleaf connects more than 70 percent of their systems. The organization upgraded to High Availability in an active-active configuration more than five years ago to accommodate the organization's zero tolerance for down time.

The organization's prior failover process and stand-alone, AIX platform offered some level of redundancy—but hardware issues could create downtimes of up to four hours while waiting for vendor support. Even planned system maintenance, with scheduled downtime of as little as 10 minutes, was met with resistance from hospital employees, since users would have to revert to manual processes.

Since implementing Infor Cloverleaf High Availability, downtime is nonexistent, and patient information, such as lab test results, flows uninterrupted among the hospital's various systems.

Summary:

The healthcare information availability revolution is here. And it requires a shift in attitude.

Healthcare organizations that have previously regarded downtime as inevitable, or as a necessary evil, must adapt to increasing availability requirements. As the volume of information exchanged across your health system grows—so will the amount of revenue lost per second of downtime.

Infor Cloverleaf's High Availability solution is a direct reflection of the demands of today's healthcare IT environment and the zero-downtime requirement of today's health IT leaders.

To learn more, visit: www.infor.com/cloverleaf.



641 Avenue of the Americas
New York, NY 10011
800-260-2640
infor.com

About Infor

Infor is the world's third-largest supplier of enterprise applications and services, helping more than 70,000 large and mid-size companies improve operations and drive growth across numerous industry sectors. To learn more about Infor, please visit www.infor.com.

Copyright© 2012 Infor. All rights reserved. The word and design marks set forth herein are trademarks and/or registered trademarks of Infor and/or related affiliates and subsidiaries. All other trademarks listed herein are the property of their respective owners. This document is provided for informational purposes only and does not constitute a commitment to you in any way. The information, products and services described herein are subject to change at any time without notice. www.infor.com.