



EASYSTREET
ADVISOR SERIES

The Reliability Roadmap:

A guide for evaluating your IT processes,
system infrastructure and facilities



THE NEED FOR RELIABILITY

The Reliability Roadmap:

A guide for evaluating your IT processes,
system infrastructure and facilities

<i>INTRODUCTION</i>	3
<i>PROCESSES</i>	6
<i>Good Processes — Repeatable</i>	6
<i>Better Processes — Standardized</i>	6
<i>Best Processes — Certified</i>	7
<i>SYSTEM INFRASTRUCTURE</i>	9
<i>Good System Infrastructure — Fortified</i>	9
<i>Better System Infrastructure — Virtualized</i>	10
<i>Best System Infrastructure — Optimized</i>	10
<i>FACILITIES</i>	12
<i>Good Facilities — Redundant</i>	13
<i>Better Facilities — Multiplied</i>	13
<i>Best Facilities — Distributed</i>	14
<i>CONCLUSION</i>	15

Introduction

When executives consider what quality they need most from their IT infrastructure, they usually conclude: Reliability.

They realize that, all factors considered, if their IT services are interruptible and undependable, their entire business operations could falter and the consequences could be disastrous.

Without reliability, nothing else really matters.

But there is no big stamp in the sky certifying an IT infrastructure as “reliable.” Whether an infrastructure is reliable is related to its specific design, its data requirements and its complex safeguards. And an infrastructure reliable for one business is not necessarily reliable for the business next door.

Yet when business requirements for IT infrastructures are examined on a case-by-case basis, a useful grid can be drawn

enabling you to pinpoint your location on a technology continuum and then to evaluate whether sufficient reliability exists for protecting your business information and critical processes.

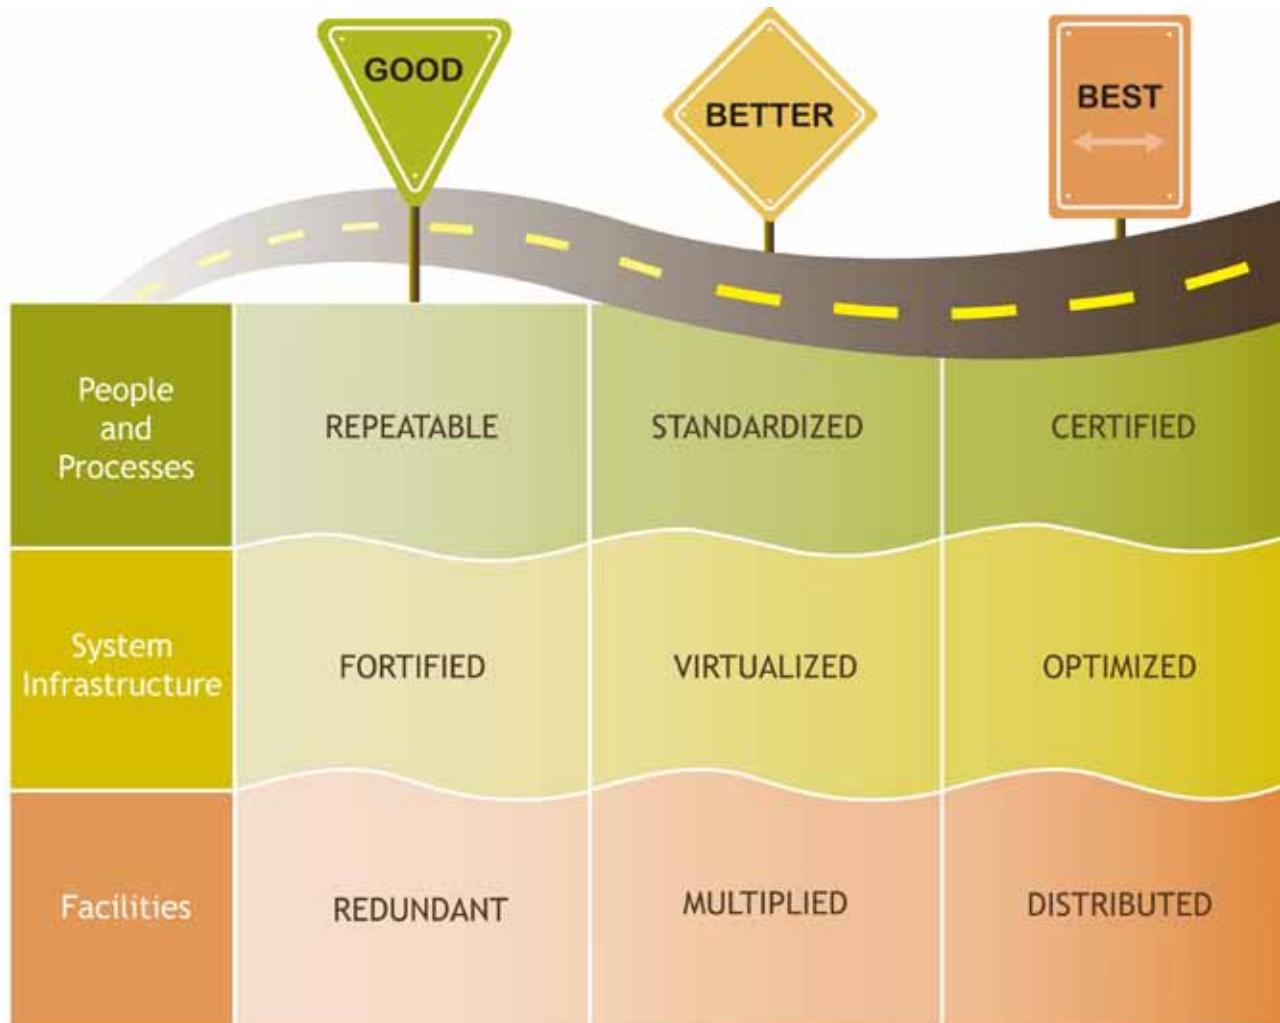
We call this grid the Reliability Roadmap.

The road to reliability

To keep the discussion manageable, we use three generic categories comprising the overall IT infrastructure for the roadmap:

- **Processes:** The rules, procedures and behaviors governing the people who manage your overall infrastructure.
- **Systems Infrastructure:** The network, hardware, systems software and related virtualization technologies.
- **Facilities:** The physical data center, including its power, heating, cooling, monitoring and security systems.

The Reliability Roadmap



Then we apply common comparative terms — good, better, and best — to lay out the map of IT infrastructure attributes and corresponding reliability levels. Keep in mind, too, that these qualitative categories tend to be fluid. One size doesn't fit all. For your specific business needs, meeting "good" category requirements might be sufficient.

- "Best" for the three categories is not reserved for large companies with sizeable IT budgets. Small and mid-size (SMB) companies can achieve "better" and "best" levels through astute alignment of their IT, business and financial objectives.
- A business's IT infrastructure can have a "better" facility, but merely "good" processes and still be considered reliable for that business. In other words, trying to achieve a "best" classification across the board may not make sense for your particular business.

The Reliability Roadmap takes into account several respected evaluation tools pertinent to IT infrastructures. Among them are the Capability Maturity Model (CMM) from the Software Engineering Institute, the Control Objectives for Information and Related Technology (COBIT) from the Information Systems Audit and Control Association and IT Governance Institute, the TIA-924 Telecommunications Infrastructure Standards for Data Centers from the Telecommunications Industry Association, the Statement of Auditing Standards Number 70: Service Organizations (SAS 70) from the American Institute of Certified Public Accountants, and the Information Technology Infrastructure Library (ITIL).

Pressures to move and improve

For some companies, a "good" facility might make the best sense in relation to its current IT infrastructure. But for another company, some movement — perhaps from "good" to "better"

The Reliability Roadmap: Pressures to Improve



— might occur due to factors beyond the company’s control. In today’s business climate, companies of all sizes confront serious IT pressures causing them to rethink fundamental practices. Some examples are:

- **Efficiencies:** The need to do more with existing resources. Competent IT administrators are becoming scarcer, and a deficit in IT talent is driving many companies to seek IT services elsewhere, offshore or locally outsourced.
- **Certifications:** Practices mandated by Sarbanes/Oxley and HIPAA, as examples, demand new levels of IT process maturity and can come with deadlines carrying adverse consequences.
- **Growth:** IT scalability cannot be effectively accomplished manually or ad hoc. Companies in that mode — throwing more people at the problem and spending more time trying to communicate and coordinate instead of actually getting things done — are an IT ‘house of cards’ waiting to fall.

Companies usually find these pressures driving them to higher performance points on the Reliability Roadmap. As one IT manager in a mid-size health-care

company recently put it: “If you don’t care that your email system is reliable only 75 percent of the time, that’s great. But if you need it 90 percent, or if it’s mission-critical, then you need a more solid foundation — as in better network, better storage, better server environment and better talent managing the system.”

From the standpoint of investment, moving from good to better, or better to best, is not trivial. If a 90-percent-available email system is required for your business, you need to make the investment. Similarly, much of your investment related to Sarbanes/Oxley certifications is required, with no choice of whether you will ever see those dollars coming back.

Seeking outside help

At a time when such pressures are forcing companies to move their IT infrastructures farther along the Reliability Roadmap, many are viewing outsourcing as a suitable alternative to in-house investments. A proven managed services provider (MSP) can offer immediate and cost-effective benefits, such as:

- **Core competency:** Few companies rank IT as a core

competency, instead treating it as a commodity. Most can achieve a decent infrastructure, but when pressures to move higher on the roadmap and costs begin escalating, turning to an MSP that already has the higher-level, proven infrastructure can make sense.

- **Economies of scale:** A good MSP offers the sufficient hardware, network connections, redundancies and security to support the IT needs of several companies. As a result, it has the IT experience and talent — the complete set of resources, several times over — to ensure reliability for its customers, whether they require good, better or best.
- **Pace of change:** Unless IT is a company’s core competency, keeping up with the pace of technological change is

daunting. It is especially difficult if you’re already behind the curve or if your ability to spend is limited. Placing their IT infrastructures with MSPs who routinely deal with changing technology is inviting to a growing number of companies, small and large.

Our hope is that this paper and the Reliability Roadmap will prove useful in evaluating your IT infrastructure requirements, help increase returns on your investment

in this critical aspect of your business, and point out the steps necessary to attain the level of reliability best suited to your requirements.

Processes

Dr. William Edwards Deming, the well-known quality guru, once observed: “Quality isn’t about people doing their best. It’s about having the right process in place.”

Processes have significant impact on reliability as IT infrastructures graduate from “good” to “better” and from “better” to “best.” Yet, it is the subtlest of the three categories examined. This is because processes are put in place and administered with no weaknesses becoming apparent until one of the other two — Facilities or System Infrastructure — either outgrows the processes, or something breaks due to an unforeseen flaw or omission in a process.



Capability Maturity Model

One of the most insightful descriptions of process implementation in the IT environment is developed by the Software Engineering Institute in its Capability Maturity Model, or CMM. It guides businesses toward a culture of software engineering and management excellence in developing and supporting their information systems.

Of the five levels CMM discusses, the first could be considered “adequate” and not relevant to the roadmap’s good-better-best matrix. It is called the “initial” level and is characterized by a handful of formal processes implemented on an ad hoc basis. They may be defined and written down — meaning that they exist, which is not always the case in many companies —



but they still are ad hoc and limited in scope. As a result, the processes are insufficiently controlled, and their operational success depends on the dedicated effort of a few individuals rather than the entire organization.

Good Processes — Repeatable

For an IT infrastructure to attain the “good” level, processes can be repeated successfully. This means they are established, and measures are applied so that repeated failures are avoided.

In CMM terms, “good” comprises the second, or “repeatable” level, as well as the third, “defined” level of process implementation. This level of more complete process documentation means new methods and tools can be added to the infrastructure and new staff trained more efficiently due to clearer understanding and implementation of the processes.

Better Processes — Standardized

The next stage of Process reliability builds upon “good” by improving IT infrastructure management through instituting

measured, standardized process to improve the quality of performance. At the “better” stage, a certified systems administrator fluent in current technologies and in implementing methodical, transparent processes is in the management role.

This level is characterized by more emphasis on quantitatively measuring the services delivered by each process. Detailed measurements are collected to identify and correct issues related to process performance. Likewise, as part of what CMM refers to as the “managed” level, the quality of all processes and their supporting activities are measured. As new sets of tools or processes are added — or adjustments made to existing processes — the ongoing measurement data enables a business to prevent recurrences of earlier process defects or errors.

An example of an institutionalized, quantified process is the change-management procedure EasyStreet uses to ensure that all changes to client data are methodical, well planned and transparent. Following change-management standards from the Information Technology Infrastructure Library (ITIL), the EasyStreet process calls for:

A request for change to be made via the Request for Change form to initiate the review process. The request spells out the desired outcome, potential negative effects, impacts on other systems and a rollback-procedure review.

EasyStreet’s Change Approval Board (CAB), representing multiple areas of the business, discusses all changes filed by midnight of the preceding day. Once the request is reviewed and approved, the change is scheduled, taking into account any conflicting changes related to other accounts.

The requester implements the change and documents the results for CAB review.

The CAB reviews all changes implemented since the previous day’s meeting to ensure the desired results were accomplished. If the implementation was incomplete or had unforeseen results, the board refers these back to the requester, who submits a modified request for change to re-initiate the process.

While the Change Approval Board is an example of a good process, a business that is growing rapidly or changing its IT direction could experience the need for an even greater process maturity to achieve a standards-based level of responsiveness and auditability that would be considered “best.”

Best Processes — Certified

For this discussion, the process level labeled “best” utilizes widely recognized process certifications, such as ITIL, COBIT and

SAS 70. At this level, the IT infrastructure's operational processes and procedures are marked by an administrative maturity not found in lower levels. There is a clear emphasis on best practices and continuous improvement.

For example, when looking at the highest level of CMM, called "optimizing," organizations proactively identify strengths and weaknesses in their processes, with the aim of preventing defects or errors. Continuous improvement is institutionalized in the development of processes, and instead of merely correcting defects as they are found, the aim is to stall future defects by eliminating their root causes through advanced planning.

ITIL

An IT infrastructure functioning at the "best" process level employs a framework of best practices, such as those found in the Information Technology Infrastructure Library (ITIL), an extensive set of management procedures to achieve high financial quality and value in IT operations. Examples are:

Incident management:
Restoring normal service operation as quickly as possible and minimizing the adverse effect on business operations, thus ensuring that the best possible levels of service quality and availability are maintained.

Problem management: Resolving root causes of incidents stemming from errors within the IT infrastructure and preventing recurrence of these errors. Under this system, a "problem" is an underlying cause of one or more incidents and a "known error" is a diagnosed problem for which a successful work-around has been identified.

Security management: Guaranteeing the safety of information, based on ISO/IEC 17799. ITIL makes it clear: "Security is the means to be safe against risks. When protecting information, it is the value of the information that has to be protected. These values are stipulated by confidentiality, integrity and availability."

ITIL's process library covers the key areas of operational concern for infrastructures striving to adopt best practices. It also includes detailed processes in areas including capacity management, service-level management, configuration management and change management.

COBIT

Another set of important IT best practices found at the roadmap's "best" level is the Control Objectives for Information and related Technology, or COBIT. Created by the Information Systems Audit and Control Association (ISACA) and the IT Governance Institute (ITGI), COBIT is a set of measures, indicators and processes for achieving maximum benefit from IT through governance and control.

COBIT's mission is to "research, develop, publicize and promote an authoritative, up-to-date, international set of generally accepted information-technology control objectives for day-to-day use by business managers and auditors." It provides 34 high-level processes covering 318 control objectives in the areas of planning and organization, acquisition and implementation, delivery and

support, and monitoring.

COBIT claims it improves the effectiveness of decision making because it:

- Defines the strategic IT plan
- Defines the information architecture
- Specifies the necessary hardware and software to execute the IT strategy
- Ensures continuous service
- Monitors the performance of the IT system

Its implications reach beyond

the IT department. Public companies subject to the Sarbanes-Oxley Act of 2002 are encouraged to adopt COBIT and the Internal Control Integrated Framework from the Committee of Sponsoring Organizations of the Treadway Commission (COSO) to achieve required levels of corporate governance.

SAS 70 (and SSAE 16)

With the passage of Sarbanes-Oxley, another set of standards gained importance related to IT infrastructure auditing. The Statement of Auditing Standards Number 70: Service Organizations — known as SAS 70 — developed by the American Institute of Certified Public Accountants defines the standards for service auditors in assessing the internal controls of entities that provide outsourced services, such as external data centers and managed services providers (MSPs).

SAS 70 allows for a single internal control review on service organizations that had previously been required to have multiple



audits under SAS 55 requirements. SAS 70 will be replaced by the new SSAE 16 audit in the fall of 2011. Under the new AICPA reporting standards, an audit that is conducted under SSAE 16 will result in a Service Organization Control (SOC) 1 report. SOC 1 reports will be available as Type 1 or Type 2 reports, very similar to the current SAS 70 reporting options.

It is evident that as an IT infrastructure attains the “best” level of process, there is considerable institutional support to identify best practices and mechanisms for continuous improvement. Again, it must be stressed that an infrastructure does not have to be huge to accommodate this level of process implementation. Being a “best” level process organization is the result of a quality IT infrastructure management team and where their priorities are placed.

System Infrastructure

The good-better-best progression for a system infrastructure reflects a steadily increasing degree of technological sophistication. A “good” system infrastructure enables IT to perform in a manner consistent with its processes and facility at their “good” levels. But at the “best” level, the system infrastructure reflects the high end of current technological equipment and strategies.

For purposes of this discussion, a system infrastructure consists of four key components:

- Infrastructure services
- Operating system
- Hardware
- Network

The operational trends and technologies affecting these four components determine where the system infrastructure is on the roadmap’s comparative spectrum, and none today is more

influential than virtualization for improving system utilization and efficiencies.

Virtualization

In its simplest terms, virtualization is a technique for greatly expanding a computing resource’s utilization and efficiency. Made possible by robust management software and native on-chip support at the processor level, virtualization on enterprise servers means that diverse operating systems and applications that previously would have required several servers now can be consolidated onto a single server.

- Virtualization greatly enhances scalability by deploying applications on virtual servers and moving away from the one-application-per-server model.
- Environments employing virtualization are less vulnerable to machine failure and witness a marked improvement in system stability.
- Virtualization lets administrators relocate workloads so maintenance does not impact uptime or service levels.

In fact, a recent Ziff Davis Media survey of more than 200 IT executives found that their perceived top three reasons for adopting virtualization were “lower hardware cost” at 43 percent, “reduce maintenance cost” at 32 percent and “server utilization rates too low” at 30 percent.

Safeguards

Network protection is another condition that increases as a system infrastructure moves higher on the reliability scale.

Examples are:

- **Firewalls:** Prevent network intrusion to the trusted internal network.
- **Network Operations Center (NOC):** Monitors a network for conditions requiring special attention — helping avoid power failures, bit errors, framing errors and the like — and enabling proactive analysis and remedy prior to performance being affected.



Virtualization technologies can deliver 60 to 80 percent server utilization.

- **Anti-virus Software:** Identifies, blocks and eliminates computer viruses and other malicious software (malware). Typically it either searches for viruses matching those in a virus dictionary or it identifies suspicious program behavior that might indicate infection.
- **Intrusion Detection Software:** Detects malicious network traffic and computer usage that a conventional firewall cannot detect. Examples are network attacks on vulnerable services, data-driven attacks on applications, host-based attacks on privileges and malware such as viruses, Trojan horses and worms.

Good System Infrastructure — Fortified

To rank higher than adequate, a system infrastructure should include sufficient name-brand hardware and spare hardware to handle processing loads in the event of equipment trouble. At the “good” level, the network is fortified with a firewall and protected by anti-virus software. A NOC monitors fundamental network performance, though its hours of surveillance may be limited to an 8 a.m. to 5 p.m. business day.

Better System Infrastructure — Virtualized

Moving to this position on the roadmap involves a sizeable technology increase along with added redundancy and system safeguards. Upgraded or expanded versions of “good” systems include:

- Implementing higher levels of anti-virus protection, both in terms of software sophistication and number of installations to more deeply protect the network.
- Installing firewalls on each server to provide maximum protection for trusted networks.
- Installing intrusion-detection systems to guard against Trojan horses, worms and other destructive hacks.
- Expanding network monitoring to 24/7, with a well-equipped NOC and clearly defined escalation procedures in the event of an alarm.

While some level of virtualization may have been present at the “good” level of system infrastructure, it is essential that virtualization be sufficiently implemented at the “better” level to effectively balance server loads for more efficient processing and to enable system maintenance without downtime.

Other attributes at this level are:

- Having servers with protocols to enable “hot swapping,” the ability to remove and replace components while the computer is operating.
- Putting virtual servers in place to boost processing capacity by managing loads as well as connections between clients

and other servers. Plus, load-balancing and high-availability server “clusters” are more evident at this level to further increase processing efficiency.

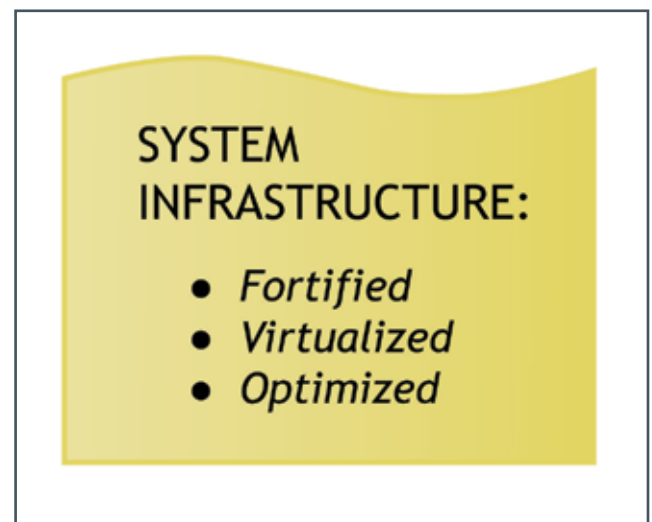
- Improving data storage and reliability through implementation of Redundant Array of Independent Drives (RAID) to divide or replicate data among multiple hard drives.
- System infrastructures of many data centers can function at the “better” level for prolonged periods, perhaps several years. Their speed, efficiency, scalability and reliability can satisfy a wide range of business sizes and types. So why would a business incur the added investment to achieve the “best” level?

Best System Infrastructure — Optimized

The highest level of system infrastructure offers state-of-the-art technology for the greatest processing efficiencies geared for optimum reliability. Just as with the “best” level of facilities — marked by multiple redundancies in physical infrastructure, including power, telecommunications, cooling and fire protection — the “best” level of system infrastructure is highly protected while remaining flexible and scalable.

Here we find blade servers designed for high-density processing. These stripped down, self-contained servers sidestep the usual rack constraints of 42 components to achieve densities of 100 or more computers per rack.

Data storage has similar exponential growth through implementations of Network Attached Storage (NAS), with more of the larger IT infrastructures now implementing Storage Area Network (SAN) architecture to allow remote storage devices — disk arrays, tape libraries and optical jukeboxes, as examples — to appear as locally attached for improved processing speeds.



The overall profile of the “best” system infrastructure makes possible the attributes of the “best” of both processes and facilities, especially:

- A system infrastructure capable of handling a distributed approach to your network, possibly utilizing multiple data centers in multiple locations.
- Sufficient redundancy and advanced technology to achieve “fault tolerant” reliability with 99.995 percent availability, with only 0.4 hours of annual downtime.
- Systems compliance with widely recognized external process certifications, such as ITIL, COBIT and SAS 70.
- Support of operational processes and procedures reflecting a high level of administrative maturity, with a clear emphasis on best practices, continuous improvement and business continuity.

In summary, scalability, adaptability and dynamism are evident across all aspects of system infrastructure at this highest level of reliability.


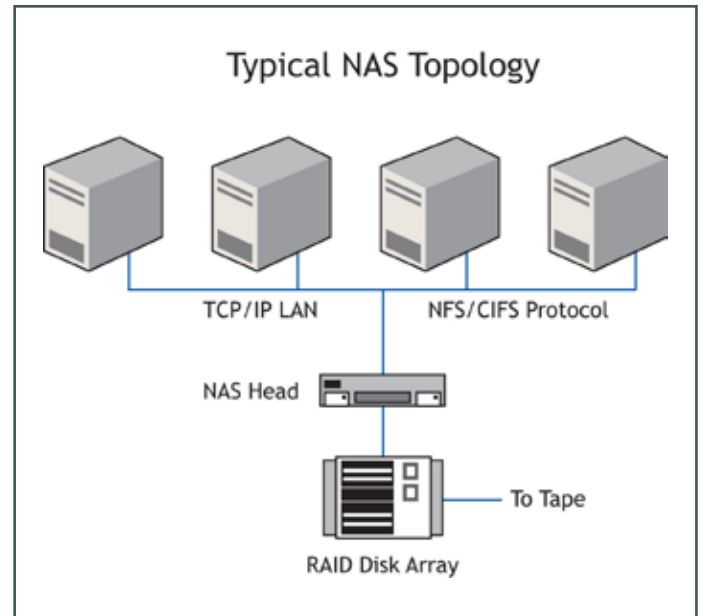
Facilities

Going from “good” to “better” to “best” in data center facilities is much more than a progression based just on physical size and data requirements. Instead, the comparative ranking is in direct relation to the degree by which a data center protects both the data and the physical structure that contains and processes it. In other words, as with Processes, a bigger IT facility doesn’t necessarily mean a “best” facility.

Before discussing the particulars of the comparative ranking, here are some fundamental elements of any data center’s facility:

Power

- A data center requires power based on the number of computers it contains. Consumption in data centers alone has jumped dramatically since 2005 with the influx of blade servers and high-speed switching. This high rate of consumption provoked Congress to order the Environmental Protection Agency (EPA) to evaluate the situation. In August of 2007, the EPA reported that data center servers accounted for 61 billion kWh in 2006, or 1.5 percent of all electrical usage in the U.S. If such trends continue unabated, usage will increase to more than 100 billion kWh by 2011.
- Data center backup power generally is provided by one or more uninterruptible power supplies (UPS) and diesel generators. To prevent single points of failure, all elements of the data center’s electrical system — including back-up



Storage methodologies like Network Attached Storage (NAS) and Storage Area Network (SAN) speed data sharing and improve data availability.

power itself — are fully duplicated, and critical servers are equipped with multiple power feeds. (See “better” and “best” discussions below.)

Cooling

- Air conditioning keeps the data center at 68-72 degrees Fahrenheit and may be used for humidity control. The goal is to keep board-level server components at the manufacturer’s specified temperature and humidity range, a critical consideration because electronic

equipment generates excess heat in a confined space and malfunctions unless it is kept cool.

- Relative humidity is maintained between 35 and 65 percent because too much causes water to condense on internal components. And if there is too little humidity, static electricity can damage computer parts.

Fire prevention

- Data centers are required to have fire detection and extinguishing systems to protect their infrastructure. Usually this involves an initial system that detects the slightest sign of particulate being given off by hot components so that a potential fire can be extinguished before flames erupt, sometimes by simply turning off the smoldering equipment.
- A second level of protection reacts to fire itself. But using conventional water sprinklers on electrical equipment can do as much damage as flames. For years the answer was Halon gas, but the Montreal Protocol now has banned Halon because of its danger to the ozone layer. Environmentally friendly alternatives include Argonite and FM-200. Some systems even control the air mixture to lower oxygen content below the level of combustion but high enough to support human life.

Physical security

- Physical access to a data center usually is restricted to selected personnel. (How much security is necessary is part of the good-better-best discussion below.)

TIA Infrastructure Standards

An important set of standards for data centers was released in April 2005. The Telecommunications Industry Association (TIA) wrote its TIA-942 Telecommunications Infrastructure Standards for Data Centers so designers could take cabling considerations into account when building new data centers.

This had not been the case until 2005 and, without these established standards, data center administrators confronted the challenge of choosing technologies and deciphering how to properly implement them in often-undersized spaces. The new standards link attributes such as space, layout and environmental considerations to reliability tiers.

Good Facilities — Redundant

To be considered “good” on the reliability scale, a data center will have moved from providing adequate services and protections to a level of redundancy that ensures a heightened rate of availability. The TIA puts data center availability at 99.741 percent — an estimated 22 hours of downtime a year — when it has redundant components.

At this “good” level, the data center has a single power backup provided by a UPS or generator, dedicated AC, single hardened telco entrance and a raised floor. It is much less susceptible to disruption from both planned and unplanned activity than the data center with a single path for power and cooling but no redundant components.

Basic fire detection and suppression systems are in place.

Security is on-site and physical during business hours, typically 8 a.m. to 5 p.m. weekdays.

Examples of pressures to move from a “good” to a “better” facility are a lack of physical space and inadequate backup systems for power or cooling.

Better Facilities — Multiplied

The most significant structural improvement that moves a data center from “good” to “better” is the implementation of multiple power and cooling distribution paths — with one path active — along with redundant components. Multiple redundancy also is in place for power backup systems and AC itself.

Another improvement is the installation of multiple hardened telecommunications with diverse paths.

Raised flooring has sufficient capacity and distribution to carry loads on one path while maintenance is performed on the other path. Fire prevention also is bolstered with an advanced fire suppression system.

Security is another area where expansion enables greater protection for longer periods. Proximity cards (“key cards”) are issued to persons with valid access. Surveillance cameras are installed and the facility is monitored 24/7.

Disaster Recovery

At the “better” level, a disaster recovery plan also is in place with related procedures to protect against disruptions — from natural disasters and terrorist attacks, to computer viruses and worker strikes. It should include:

- Explanation of procedures, such as regular off-site backups of all data, use of Storage Area Networks (SANs), down to the level of surge protectors and fire extinguishers.
- A disaster recovery framework based on Recovery Time Objective (RTO), the acceptable amount of time between the disaster and resumption of function as well as restoration of data, and the Recovery Point Objective (RPO), the acceptable data rollback that determines how current restored data must be.
- A formal, written Disaster Recovery Plan that addresses, at a minimum, the response, recovery and resumption

of services and that explains all tasks involved in these activities.

With these facilities capabilities in place, the TIA-924 puts data center availability at 99.982 percent, or only 1.6 hours of downtime in a year. Planned activities no longer affect operations, with the only disruptions coming from unplanned events.

Best Facilities — Distributed

Achieving the “best” reliability level means that a facility contains sufficient redundancy and advanced technology to achieve nearly perfect availability. It also might mean the business adopts a distributed approach to its IT, utilizing multiple data centers in multiple locations.

The TIA-942 standard calls its top tier “fault tolerant” and gives it a 99.995 percent availability, with only 0.4 hours of annual downtime. This top tier is characterized as: “Planned activity does not disrupt critical load and the data center can sustain at least one worst-case, unplanned event with no critical load impact.”

This is achieved by use of multiple active power and cooling distribution paths with redundant components, such as two UPS, each with N+1 redundancy.

In addition to the TIA-942 standard, a “best” data center has advanced fire detection and suppression systems. Security should be biometric, where attributes such as face recognition or fingerprints ensure that a proximity card is being carried by its proper owner. Security should be physically manned 24/7 and supported by multiple surveillance cameras.

Business Continuity

Another area of upgrade from “better” to “best” involves moving from Disaster Recovery planning to a full Business Continuity Plan (BCP) in the event of local incidents such as fires, regional incidents such as earthquakes, or national incidents such as pandemic illness. The key reference document is the BCP manual:

- It should contain the names, addresses and phone numbers for crisis management staff, general staff members, clients and vendors, along with the location of the offsite data backup storage media, copies of insurance contracts and other critical materials necessary for organizational survival.
- At its most complex, it could outline a secondary work site, technical requirements and readiness, regulatory reporting requirements, work recovery measures, the means to reestablish physical records, the means to establish a new supply chain, or the means to establish new production centers.

An excellent last word on the Facilities roadmap comes from the TIA’s standards for data centers:

Each of the components of the data center and its supporting systems must be planned, designed and implemented to work

together to ensure reliable access of data center resources while supporting future requirements. Neglecting any aspect of the design can render the data center vulnerable to cost failures, early obsolescence and intolerable availability.



Conclusion

With reliability as the most critical factor for IT infrastructures large and small, businesses must

weigh their level of investment against the risk of downtime. Maintaining high reliability involves a careful blend of technology — sophisticated network and computing tools to protective infrastructure redundancies — and skillful IT management.

The issue facing countless CEOs, CFOs and CTOs today is how to best leverage their IT investment in light of quickening technology changes, tightening labor markets and evolving strategic business directions. Whether the IT infrastructure resides inside the business, is distributed across several locations or is outsourced to a reputable Managed Services Provider, reliability will depend on a balance of processes, facilities and system infrastructure configured to address a business’s specific requirements.

About EasyStreet Online Services, Inc.

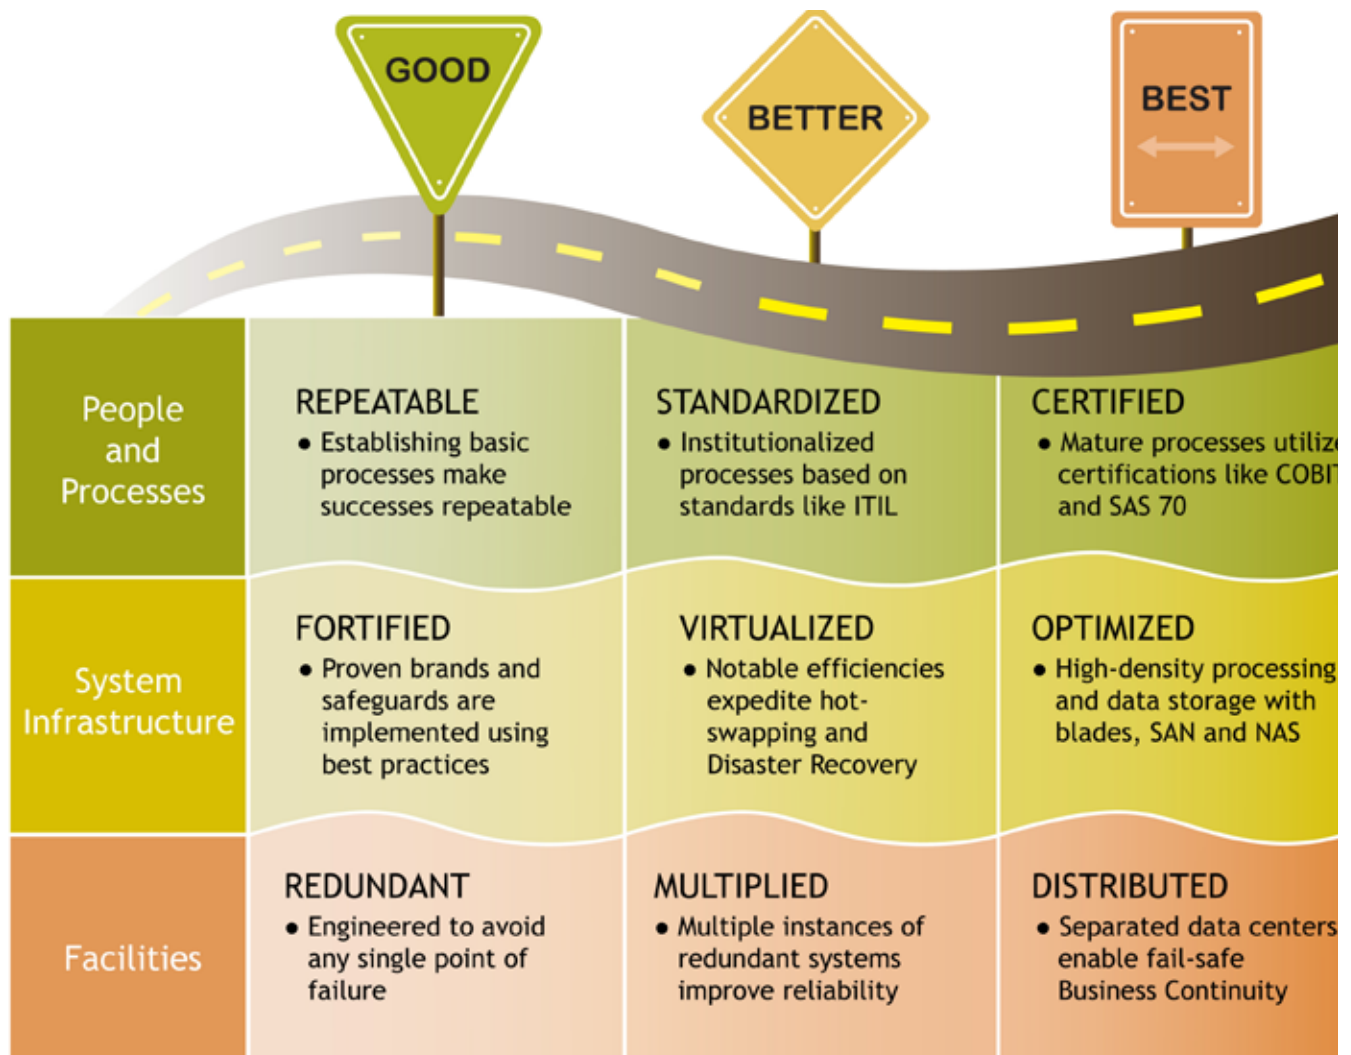
EasyStreet owns the longest-standing record of customer satisfaction of any IT-as-a-Service provider in the Pacific Northwest. The reason is its people — from its team of skilled IT technicians and expert support staff to its forward-thinking senior management.

EasyStreet was founded in 1995 to provide an array of services satisfying the IT needs of area businesses. EasyStreet has been

recognized repeatedly by the Portland Business Journal as one of Oregon's "Most Admired Technology Companies," one of the state's "100 Fastest Growing Companies," and ranks among the Deloitte & Touche "Technology Fast 50."

EasyStreet has been connected to the Internet since 1995 — continuously, without a single interruption.

The Reliability Roadmap



EasyStreet Online Services, Inc.
 9705 SW Sunshine Court
 Beaverton, OR 97005

(503) 646-8400
 (877) 567-EASY

info@easystreet.com
 www.easystreet.com