

GREEN IT OVERVIEW

Unit Structure

- 1.0 Objective
- 1.1 Introduction
- 1.2 Environmental Concerns and Sustainable Development
- 1.3 Environmental Impacts of IT
- 1.4 Green IT
- 1.5 Holistic Approach to Greening IT
- 1.6 Greening IT
- 1.7 Applying IT for Enhancing Environmental Sustainability
- 1.8 Green IT Standards and Eco-Labeling of IT
- 1.9 Enterprise Green IT Strategy
- 1.10 Green Washing
- 1.11 Green IT: Burden or Opportunity?
- 1.12 Summary
- 1.13 Exercise
- 1.14 References

1.0 OBJECTIVE

- Explains what green IT is and examines the significance of green IT.
- Discusses environmental concerns, global warming and the principles of sustainable development.
- Examines the environmental impacts of IT.
- Describes the three key dimensions of green IT and explains green IT 1.0 and 2.0.
- Presents a holistic approach to greening IT.
- Discusses how data centres, cloud computing, storage systems, software and networks can be made greener.
- Highlights how IT could help businesses in their environmental initiatives and reduce their carbon emissions.
- Outlines enterprise green IT strategy

1.1 INTRODUCTION

Enterprises, governments and societies at large have a new important agenda: tackling environmental issues and adopting environmentally sound practices. Over the years, information technology (IT) has

fundamentally altered our work and life and improved our productivity, economy and social well-being. IT now has a new role to play – helping to create a greener, more sustainable environment whilst offering economic benefits. But IT has been contributing to environmental problems which most people do not realize.

Computers and other IT infrastructure consume significant amounts of electricity, which is increasing day by day, placing a heavy burden on our electric grids and contributing to greenhouse gas (GHG) emissions. Additionally, IT hardware poses environmental problems during both its production and its disposal.

Whilst many people consider IT to be part of the problem to environmental pollution, it can be its savior too. In other words, IT is both a solution and a problem for environmental sustainability.

Green IT benefits the environment by improving energy efficiency, lowering GHG emissions, using less harmful materials and encouraging reuse and recycling. Thus, green IT includes the dimensions of environmental sustainability, the economics of energy efficiency and the total cost of ownership, which includes the cost of disposal and recycling.

Increased awareness of the harmful effects of GHG emissions, new stringent environmental legislation, concerns about electronic waste disposal practices and corporate image concerns are driving businesses and individuals to go green.

The imminent introduction of more green taxes and regulations will trigger a major increase in demand for green IT products, solutions and services. Hence a growing number of IT vendors and users have begun to develop and offer green IT products and services. As business and governments try to balance growth with environmental risks, we will be legally, ethically and/or socially required to ‘green’ our IT products, applications, services and practices.

To foster green IT, we should understand the following issues: What are the key environmental impacts arising from IT? What are the major environmental IT issues that we must address? How can we make our IT infrastructure, products, services, operations, applications and practices environmentally sound? What are the regulations or standards with which we need to comply? How can IT assist businesses and society at large in their efforts to improve our environmental sustainability?

1.2 ENVIRONMENTAL CONCERNS AND SUSTAINABLE DEVELOPMENT

Numerous scientific studies and reports offer evidence of climate change and its potential harmful effects. Specifically, the growing accumulation of GHG is changing the world’s climate and weather patterns, creating droughts in some countries and floods in others and pushing global temperatures slowly higher, posing serious worldwide problems. Global

data show that storms, droughts and other weather-related disasters are growing more severe and frequent.

1.2.1 The Inconvenient Truth

Not everyone agrees, however, with these predictions regarding global warming and its impacts. For instance, controversies exist concerning the causes of global warming, whether this warming trend is unprecedented or within normal climatic variations, predictions of additional warming, what the consequences are and what actions should be taken. Environmental groups, numerous governmental reports and many in the media are, however, in agreement with the scientific community in support of human-caused warming. Several scientific societies and academies of science, including all major countries' national academies of science, endorse that global warming is mainly caused by human activity and will continue if GHG emissions are not reduced. Driven by the disastrous impact of recent storms, floods, droughts and excessive heat that many people have experienced around the world, various studies on global warming and its impact and major global campaigns, many people have begun to think seriously about global warming and its impacts and to do whatever they can to address this problem. Governments, enterprises and people all have roles in combating global warming and building a sustainable environment. There is now greater awareness and a growing commitment to address environmental problems. Inaction to arrest environmental degradation would significantly affect not only current but also future generations and our further

1.2.2 Sustainable Development

Sustainability is all about meeting needs and seeking a balance between people, the environment and the economy. According to the United Nations Global Commission on the Environment and Development's 1987 Brundtland Report, sustainable development is the 'development that meets the needs of the present without compromising the ability of future generations to meet their own needs'. Sustainable development comprises economic, environmental and social dimensions.

1.3 ENVIRONMENTAL IMPACTS OF IT

As mentioned in this chapter, IT affects our environment in several different ways. Each stage of a computer's life, from its production, through its use and to its disposal, presents environmental problems. Manufacturing computers and their various electronic and non-electronic components consume electricity, raw materials, chemicals and water, and generate hazardous waste. All these directly or indirectly increase carbon dioxide emissions and impact the environment.

Total electrical energy consumption by servers, computers, monitors, data communications equipment and data centre cooling systems is steadily increasing. This increase results in greater GHG emissions, as most electricity is generated by burning fossil fuel like coal, oil and gas. For

instance, each PC in use generates about a ton of carbon dioxide every year. Computer components contain toxic materials. Increasingly, consumers discard a large number of old computers, monitors and other electronic equipment 2–3 years after purchase, and most of this ends up in landfills, polluting the Earth and contaminating water.

The increased number of computers and their use, along with their frequent replacements, make IT's environmental impact a major concern. Consequently, there is increasing pressure on the IT industry, businesses and individuals to make IT environmentally friendly throughout its life cycle, from birth to death to rebirth. As many believe, it's our social and corporate responsibility to safeguard our environment.

1.4 GREEN IT

IT now has a new role to play in creating a greener, more sustainable environment, whilst offering economic benefits by becoming greener.

Green IT is an umbrella term referring to environmentally sound information technologies and systems, applications and practices.

1. the efficient and effective design, manufacture, use and disposal of computer hardware, software and communication systems with no or minimal impact on the environment;
2. the use of IT and information systems to empower – that is, support, assist and leverage – other enterprise-wide environmental initiatives and
3. the harnessing of IT to help create awareness among stakeholders and promote the green agenda and green initiatives.

Green IT is not just about creating energy-efficient IT systems (hardware, software and applications), though this is an important component, especially as the use of IT proliferates. Green IT is also about the application of IT to create energy-efficient, environmentally sustainable business processes and practices, transportation and buildings.

IT can support, assist and leverage environmental initiatives in several areas and also help create green awareness. IT contributes to only about 2–3% of GHG emissions. The vast majority of emissions come from non-IT sources. So, broader applications of IT in other areas of the economy could bring significant energy savings and improve overall environmental sustainability. According to the SMART 2020 report, IT's largest influence will be by enabling energy efficiency in other sectors, an opportunity that could deliver carbon savings five times larger than the total emissions from the entire information and computer technology (ICT) sector in 2020. IT can help organizations to minimize their environmental impacts in areas such as GHG emissions, toxic contamination and energy and water consumption.

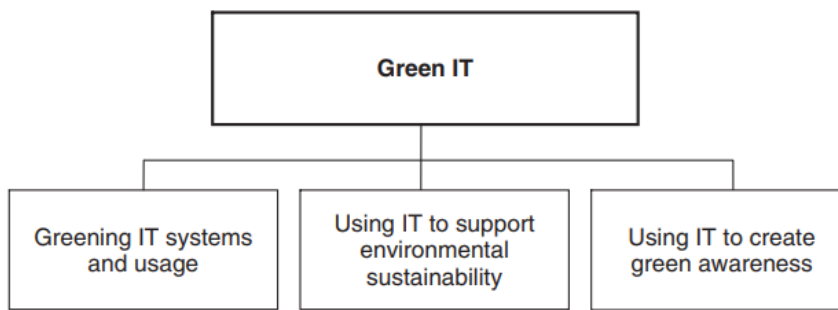


Figure – 1.1 Green IT

1.5 HOLISTIC APPROACH TO GREENING IT

To comprehensively and effectively address the environmental impacts of IT, we must adopt a holistic approach that addresses the problems along these six complementary directions :

1. **Green design:** Design energy-efficient and environmentally sound components, computers, servers and cooling equipment.
2. **Green manufacturing:** Manufacture electronic components, computers and other associated subsystems with minimal or no impact on the environment.
3. **Green use:** Reduce the energy consumption of computers and other information systems, and use them in an environmentally sound manner.
4. **Green disposal:** Refurbish and reuse old computers, and properly recycle unwanted computers and other electronic equipment.
5. **Green standards and metrics :** These are required for promoting, comparing and benchmarking sustainability initiatives, products, services and practices.
6. **Green IT strategies and policies :** These effective and actionable strategies and policies add value and focus on both short- and long-term benefits. These are aligned with business strategies and practices, and are key components of greening IT.

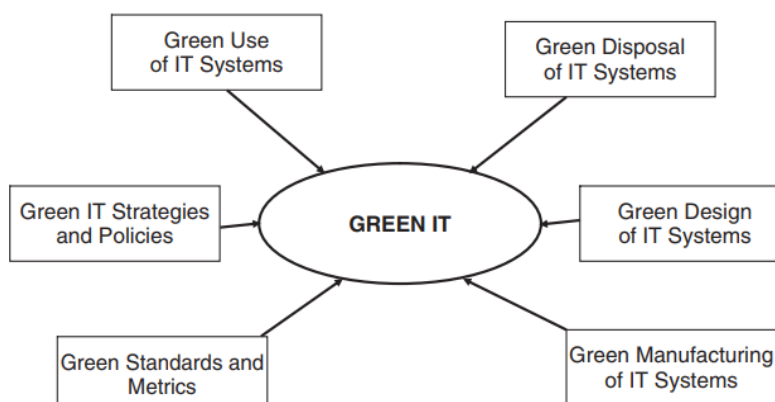


Figure 1.2 Holistic Approach to Greening IT

1.5.1 Greening Computer's Entire Life Cycle

As shown in Figure 1.3, the entire life cycle of a computer, server and storage system could be made greener, reducing their GHG emissions and carbon footprint and minimizing or eliminating toxic materials used and/or released to the environment.

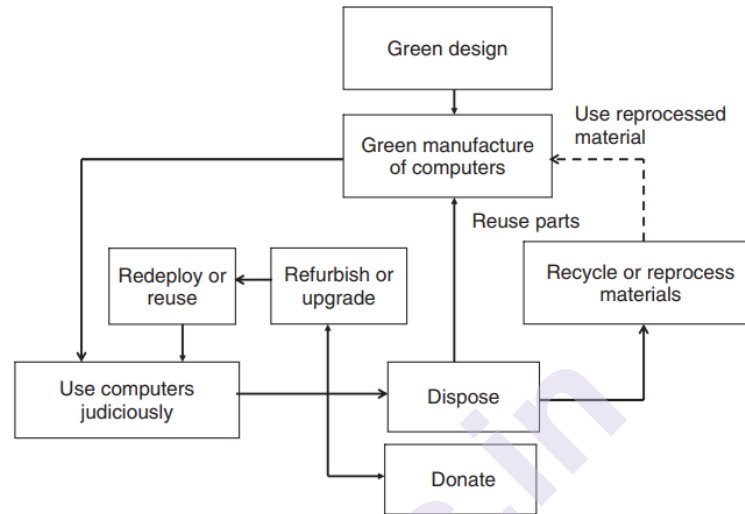


Figure 1.3 Greening Computer's Entire Life Cycle

1.5.2 The Three Rs of Green IT

Unwanted computers, monitors and other hardware should not be thrown away as rubbish, as they will then end up in landfills and cause serious environmental problems. Instead, we should refurbish and reuse them, or dispose them in environmentally sound ways.

Reuse, refurbish and recycle are the three 'Rs' of greening unwanted hardware.

- **Reuse** - Many organizations and individuals buy new computers for each project or once every 2–3 years. Instead, we should make use of an older computer if it meets our requirements. Otherwise, we should give it to someone who could use it in another project or unit. By using hardware for a longer period of time, we can reduce the total environmental footprint caused by computer manufacturing and disposal.
- **Refurbish.** We can refurbish and upgrade old computers and servers to meet our new requirements. We can make an old computer and other IT hardware almost new again by reconditioning and replacing some parts. Rather than buying a new computer to our specifications, we can also buy refurbished IT hardware in the market. More enterprises are now open to purchasing refurbished IT hardware, and the market for refurbished equipment is growing. If these options are unsuitable, we can donate the equipment to

charities, schools or someone in need, or we can trade in our computers.

- **Recycle-** When we cannot refurbish or otherwise reuse computers, we must dispose of them in environmentally friendly ways by depositing them with recognized electronic recyclers or electronic waste (e-waste) collectors. E-waste – discarded computers and electronic goods – is one of the fastest-growing waste types and poses serious environmental problems. The United Nations Environment Program estimates that 20–50 million tons of e-waste are generated worldwide each year, and this is increasing. IT hardware contains toxic materials like lead, chromium, cadmium and mercury. If we bury IT hardware in landfills, toxic materials can leach harmful chemicals into waterways and the environment. If burned, they release toxic gases into the air we breathe. So if e-waste is not discarded properly, it can harm the environment and us. Waste electrical and electronic equipment (WEEE) regulations aim to reduce the amount of e-waste going to landfills and increase recovery and recycling rates.

Disposal of e-waste, if not properly done, causes serious environmental damage and health problems particularly to those directly involved in the disposal or recycling. Despite bans on the export and import of e-waste, e-waste gets into developing countries (such as India, China and Philippines) for ‘recycling’ as the cost of recycling is lower there. Unfortunately, as environmental regulations and proper means of e-waste disposal and recycling are not enforced in practice in these countries, e-waste is handled ‘informally’ in unofficial recycling markets by manual, crude, hazardous means to extract metals and other valuables.

Computer manufacturers should take their share of responsibility and take action to reduce pollution caused by their products’ end of life. For instance, they should widely adopt a take-back option, whereby they take from consumers the computers that they no longer need, and arrange for their disposal in an environmentally friendly manner through an e-waste recycling plant. They should educate customers on what they should do with their old computers. They should also gradually eliminate or minimize the use of toxic materials in computers, which some computer manufacturers are doing.

1.6 GREENING IT

One might wonder how we should go about greening IT, and what subsystems of IT can be greened. In fact, every subsystem and peripheral of IT can be greened. The key among them are PCs, notebooks and servers, data centres and cloud computing, software (system and application software, along with the processes of software design and development), storage systems and networking and communication systems and protocols. Peripherals such as printers can be made energy efficient and environmentally friendly.

1.6.1 Green PCs, Notebooks and Servers

We can significantly reduce energy consumption by making small changes to the ways we use computers. Most desktop computers run even when they aren't being used, because users needlessly leave them on, wasting electricity

Furthermore, computers generate heat and require additional cooling, which add to the total power consumption and cost. Whilst the savings in energy costs per PC may not seem like much, the combined savings for hundreds of computers in an enterprise is considerable. We can reduce PC energy consumption by adopting several measures.

- **Enabling power management features** - Without sacrificing performance, we can program computers to automatically power down to an energy-saving state when we are not using them.
- **Turning off the system when not in use** - This is the most basic energy conservation strategy for most systems.
- **Using screensavers** - A blank screensaver conserves more power than a screensaver that displays moving images, which continually interacts with the CPU. But even that reduces the monitor's energy consumption by only a small percentage.
- **Using thin-client computers** - Users can choose to employ thin-client computers, which draw about a fifth of the power of a desktop PC.

These measures, though easily adoptable, will not become a practical reality without users' wholehearted willingness and active participation. Even simple steps by one individual or organization can make a huge difference when leveraged across the vast number of individuals and organizations across the world. Smart companies will adopt innovative environmental strategies to innovate, create value and build a competitive advantage. Chapter 2 describes the concept of green hardware including PC power management, energy-efficient power converters, the use of multicore processors, newer types of displays, the use of less toxic materials and related topics.

1.6.2 Green Data Centres

Enterprise data centres, the modern engine rooms that power the Internet and corporate computing, are growing in their number, capacity and power consumption. For instance, according to an IBM estimate, the power demanded worldwide by data centres currently stands at 100 billion kWh a year, and data centres are one of the fastest-growing users of power (Pritchard, 2007). The carbon footprint of data centres has been increasing dramatically as they consume much energy to power their IT systems and data centre cooling systems.

A study by Jonathan Koomey reveals that the total power used by servers represented 0.6% of total US electricity consumption in 2005 (Pritchard, 2007). With the power needed for cooling and other auxiliary services,

electricity use rises to 1.2%, equivalent to the power used by all the country's colour televisions. Aggregate electricity use for servers doubled between 2000 and 2005, and most of this came from businesses installing large numbers of new servers (Pritchard, 2007).

The continued rise of Internet and Web applications is driving the rapid growth of data centres and an increase in energy use. To handle more transactions in less time, to process and store more data and to automate more business processes, enterprises are installing more servers or expanding their capacity, all of which demands more computing power.

As energy prices increase worldwide, data centres' operational costs also increases. Energy costs now account for nearly 30% of a data centre's operating costs. As a result, IT is increasingly coming under scrutiny, and data centre efficiency is a major issue facing IT departments.

The number of server computers in data centres has increased sixfold, to 30 million, in the last decade, and each server draws far more electricity than earlier models. Besides the cost, the availability of electrical power is becoming a critical issue for many companies whose data centres have expanded steadily. Energy suppliers need time to design, build and supply the huge amounts of additional electrical power (a few megawatts) demanded by data centres. These social, financial and practical constraints force businesses and IT departments to consider how to reduce, or at least limit, energy consumption by data centres.

So data centres must become greener as well. A green data centre is one in which IT system, air-conditioning systems, electrical and mechanical systems and the buildings that house the data centre are designed and operated for maximum energy efficiency, low carbon footprint and minimum environmental impacts. The data centre uses advanced cooling, heating and IT systems to tailor power consumption to processing and operational needs. Ways to save data centre energy consumption include server, storage and network virtualization, the use of blade servers, server clustering and consolidation and the use of energy-efficient power supplies. Chapter 5 details key sustainability challenges facing data centres and discusses strategies to minimize energy consumption and reduce the carbon footprint of IT systems and data centre facilities.

The EU Code of Conduct on Data Centres' Energy Efficiency is a voluntary initiative aimed at reducing the environmental, economic and energy supply security impact of data centres. The scope of the Code of Conduct encompasses both the equipment and system levels, focussing on two primary areas: the IT load (i.e. the IT capacity available for the power consumed) and facilities load (equipment and systems that support the IT load, such as cooling systems, power distribution units (PDUs) and uninterruptable power supply(UPS)). The flexible and adaptable Code of Conduct encourages data centre operators and owners to undertake a cost-effective adoption of energy-efficient practices without hampering their data centres' mission-critical functions.

1.6.3 Green Cloud Computing

Cloud computing represents a paradigm shift. It is a transition from computing-as-a product to computing-as-a-service, which is shared and scalable on demand. Driven by the benefits that cloud computing offers, businesses, educational institutions, governments and individuals in both developed and emerging markets have begun to use it for several applications. Cloud offerings and use are growing, and hence these created huge demands on data centres that house the clouds. To cater to the growing demands of cloud-computing services, vendors use large-scale data centres which consolidate thousands of servers with other infrastructure such as cooling, storage and communication networks. With the growth of the cloud, therefore, comes increasing energy consumption by data centres. Then how can clouds be greener? Cloud computing is a green solution as cloud infrastructure embraces two critical elements of a green IT: resource efficiency and energy efficiency.

1.6.4 Green Data Storage

Data and information storage requirements keep growing drastically. Storage systems in data centres consume significant amounts of power and cooling. For instance, in a data centre, storage systems consume anywhere between 24% and 40% of total IT power usage and are the centre's biggest power hogs. So, besides making servers energy efficient, the focus is on greening data storage. Several approaches including MAID (Massive Array of Idle Disks), disk spin down, tiered storage and solid-state drives (SSDs) are used for improving energy efficiency and cutting the overall costs of storing persistent data. System-level approaches such as storage virtualization, thin provisioning (TP) and data de-duplication helps reduce required storage space and make effective utilization of available space. For in-depth coverage of green data storage, refer to Chapter 6. Green storage has to be part of a bigger 'greener ICT' strategy.

Data de-duplication is the elimination of coarse-grained redundant data, typically to improve storage utilization. De-duplication may occur in line, as data are flowing, or post process after data have been written. It reduces the required storage capacity since only the unique data are stored. Depending on the type of information stored, de-duplication of data can yield a compression ratio from 3:1 to 10:1. Data de-duplication also reduces the data that must be sent across a network for remote backups, replication and disaster recovery.

TP is a method of storage resource management and virtualization that lets IT administrators limit the allocation of actual physical storage to what applications immediately need. TP operates by allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time. NetApp, EMC, Compellent, Xiotech, Dell, 3PAR and several others offer tiering and TP systems.

Tiering, in addition to filtering out unnecessary data and files, allocates business data and files to the most efficient layer of storage available: for example, Tier 1 (on-demand data), Tier 2 (data not critical but still timely)

and Tier 3 (archival data). Tiering also provides immediate access to timely business data so they can be used for internal corporate analytic if needed. Chapter 6 discusses the power consumption characteristics of different storage solutions and media, and highlights different energy management techniques for hard disks and system-level green measures.

1.6.5 Green Software

Does software impact the environment? Yes. Software plays an important role in

determining overall energy consumption and computational efficiency. For instance, a single ill-behaving, computationally inefficient or power-unfriendly software component on a system can thwart all of the power management benefits built into the hardware.

Software is a key element in improving environmental sustainability. Green software is environmentally friendly software that helps improve the environment. Green software can be classified into four broad categories:

- Software that is greener – consumes less energy to run;
- Embedded software that assists other things in going green (smart operations);
- Sustainability reporting software, or carbon management software (CMS);
- Software for understanding climate change, assessing its implications and forming suitable policy responses.

The manner in which software is developed and the quality attributes of software impact the environment. Sustainable software development refers to creating software addressing environmental requirements and perspectives. Development-related attributes such as modifiability, re-usability and portability and performance attributes such as computational time and efficiency, usability and dependability influence software's environmental impact. For further discussion on green software and on software design considerations and software methodologies to improve software energy efficiency, refer to Chapters 3 and 4.

Using open source methodologies for application development is expected to result in energy savings as collaborative development processes tend to be more efficient than traditional processes.

Organizations are now required to account and manage their carbon footprint, the amount of GHG

(in CO₂ equivalent) that they produce. This requires IT systems that can measure, analyse and manage carbon emissions in a cost-effective and efficient manner, called a carbon management system. A number of carbon management systems such as Carbon view (www.carbon-view.com), Carbon Planet (www.carbonplanet.com), Greenstone

(www.greenstonecarbon.com/software.php) and Emissions Logic (www.cemsus.com/cems/emissions-logic) are available in the market (see also the CMS Directory at www.carbonmanagementsoftware.com). CMS is also integrated with ERP software from vendors like SAP, Oracle and Microsoft.

1.6.6 Green Networking and Communications

Networks and communications play more significant roles than ever before, and facilitate data transfer and sharing. They enable us to communicate and share information, shop, learn and socialize online, make our work environment smarter and do many other things. The demands on communication networks, wired and wireless, have been constantly increasing, and as a result the energy consumption of communication systems has increased considerably. Traditional networking systems and communication protocols are not particularly designed for energy efficiency. Thus, they have some negative impact on the environment

1.7 APPLYING IT FOR ENHANCING ENVIRONMENTAL SUSTAINABILITY

IT can be a key driver in greening several industries and activities, and a positive force towards environmental sustainability initiatives. We should make IT a positive force in environmental change. As mentioned in this chapter, several studies reveal that IT contributes only about 3% of GHG emissions; thus the vast majority of emissions come from non-IT sources, almost all of which can realize enhanced energy efficiency and minimize their environmental pollution through the smarter use of IT.

Besides IT becoming green, it can also be a very helpful enabler and aid to create a better environment. Some of the opportunities for this are as follows:

- Software tools for analyzing, modelling and simulating environmental impacts and for environmental risk management;
- Platforms for Eco-management, emission trading and ethical investing;
- Tools for auditing and reporting energy consumption and savings and for monitoring GHG emissions;
- Environmental knowledge management systems, meaning the acquisition and transfer of environmental knowledge, decision support systems and collaborative environments; environmental ontologies;
- Environmental information systems engineering, including geographic information systems and environmental (meta-)data standards;
- Urban environment planning tools and systems;

- Technologies and standards for interoperable environmental monitoring networks; smart in situ sensor networks;
- Integration and optimization of existing environmental monitoring networks, easy plugin new sensors, sensor cooperation and networks;
- Tools and systems for optimizing organizational work-flows.

To reduce their carbon footprint, organization can dematerialise some of their products and activities embracing IT. Dematerialization refers to the transformation of physical goods to information goods represented in digital form – “turning atoms to bits.” Organisations can adopt electronic billing (e-billing) instead of paper-based billing, offer music and videos for online download rather than on CDs, use e-books and electronic documents rather than printed documents and do video conferencing rather holding face-to-face meetings.

1.8 GREEN IT STANDARDS AND ECO-LABELLING OF IT

To promote and adopt standardization, a number of green IT standards and directives have emerged. Key among them are EPEAT (Electronic Product Environmental Assessment Tool), RoHS (Restriction of Hazardous Substances Directive), WEEE, Energy Star, LEED (Leadership in Energy and Environmental Design), the ISO14001 core set of standards for designing and implementing an effective environmental management system and the EN 16001 Energy Management System. EPEAT is a popular, easy-to-use assessment tool to help organizations compare computer desktops, laptops and monitors based on their environmental attributes. EPEAT-registered products are classified as bronze, silver or gold (www.epeat.net) and they have reduced levels of cadmium, lead and mercury to better protect human health. They are more energy efficient and easier to upgrade and recycle.

1.9 ENTERPRISE GREEN IT STRATEGY

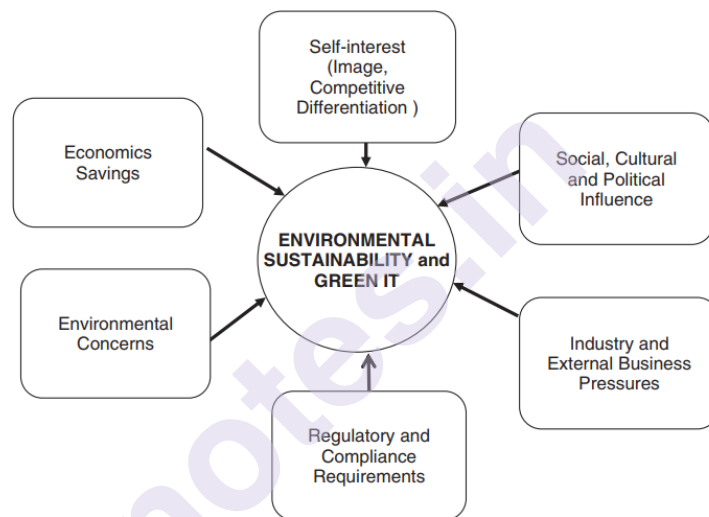
Green IT and green initiatives are becoming a key agenda for enterprises and governments. They are driven by the benefits they offer and by several on-going developments such as concerns about climate change, government regulations and peer pressure and influence, as shown in Figure 1.5.

Each enterprise must develop a holistic, comprehensive green IT strategy, which should be a component of, and aligned with, an enterprise-wide green strategy. It should then develop a green IT policy outlining aims, objectives, goals, plans of action and schedules.

Large enterprises should also appoint an environmental sustainability officer to implement their green policy and to monitor their progress and

achievements. To green their IT, enterprises can take any one or a combination of the following three approaches:

1. **Tactical incremental approach** - In this approach, an enterprise preserves the existing IT infrastructure and policies and incorporates simple measures to achieve moderate green goals such as reducing energy consumption. These measures include adopting policies and practices such as power management, switching off computers when not in use, using compact energy-efficient light bulbs and maintaining an optimal room temperature. These measures are generally easy to implement without much cost. However, enterprises should work towards these measures only as short-term, ad hoc solutions.



2. **Strategic approach** - In this approach, an enterprise conducts an audit of its IT infrastructure and its use from an environmental perspective, develops a comprehensive plan addressing broader aspects of greening its IT and implements distinctive new initiatives. For example, an enterprise may deploy new energy-efficient, environmentally friendly computing systems, or it may develop and implement new policies on procuring, operating and/or disposing of computing resources. Whilst the primary rationale is still cost efficiency and a reduced carbon footprint, this approach also considers other factors such as branding, image creation and marketing.
3. **Deep green approach** : This approach expands upon the measures highlighted in the strategic approach, wherein an enterprise adopts additional measures such as implementing a carbon offset policy to neutralize GHG emissions – including planting trees, buying carbon credits from one of many carbon exchanges or using green power generated from solar or wind energy.

An Accenture Report suggests that about 96% of CEO are in favour of integrating sustainability issues with organizational strategies (Lacy et al.,

2010). Chapter 8 discusses crucial steps and considerations in developing green IT strategies, and challenges in implementing green strategies and policies. Chapters 9 and 10 deal with ‘Sustainable Information Systems and Green Metrics’ and ‘Enterprise Green IT Readiness’, respectively. Chapter 14 outlines how to manage IT with a focus on environmental sustainability.

1.10 GREEN WASHING

Green washing refers to the practice of organizations exaggerating their green credentials and environmental sustainability attributes, and making false claims. Green washing is an amalgam of the terms green and whitewash. It is an unjustified appropriation of environmental virtue. This socially irresponsible and unethical practice misleads customers and the public regarding the company’s environmental practices or the environmental benefits of its product or services. It is a marketing ploy to establish an eco-friendly image to consumers, investors, businesses and regulators.

1.11 GREEN IT: BURDEN OR OPPORTUNITY?

The green philosophy in general, and the ‘go green’ movement and green demands on corporate IT in particular, do not excessively or unduly burden IT systems, corporate IT departments or functional units. In fact, these initiatives provide an opportunity to revisit and examine our IT systems and their operations in terms of energy efficiency and resource utilization, and thereby enable us to go lean on IT, minimize IT’s energy consumption and save on energy bills. Until recently, IT functions and activities primarily focused on meeting their functional and performance requirements. Very little attention was paid to aspects such as energy consumption, effective utilization of IT resources, IT’s operational costs or IT’s negative impact on environments at the stages of design, manufacturing, use, reuse and disposal. There is a pressing need to address these neglected or overlooked aspects as they are now important for safeguarding our environment. IT is required to go green. It is good for IT, businesses and the entire planet. Though initially some might view going green as a burden, a closer examination of green philosophy reveals that it includes improving energy efficiency, improving resource utilization, reducing waste, promoting reuse and recycling and more such benefits. This will give the necessary impetus and motivation to turn IT green and use IT in innovative new ways to green all other corporate functions.

We also need to look at green requirements from another viewpoint. The implications of not going green might cost a lot in the context of emerging stricter environmental regulations, stakeholder demands, competitiveness, brand or corporate image and social responsibility. A holistic and objective view would reveal green IT – greening of IT and greening by IT – to soon become a necessity, not an option. Even if one feels overburdened with ‘go green’ initiatives and demands, it is better to adopt them in the interest of self and our planet.

Green IT will be a top priority for several years to come, as it is both an economic and environmental imperative. Several case studies on greening efforts reveal that businesses Harnessing Green IT that reduce their environmental (carbon) footprint can also reduce costs and improve their public image. IT professionals, CIOs and IT support staff are thus being called upon to deliver environmentally sustainable IT solutions (Wilbanks, 2008). Even simple steps that one individual or organization takes can make a huge difference when leveraged across the vast number of individuals and organizations across the world.

However, there is a disparity in the level of green IT understanding across companies, IT professionals, students and IT users. Many do not know how or where to begin or are unwilling to implement green IT. Although green initiatives are catching the attention of the corporate world, some IT professionals, executives and IT departments feel excessively burdened by the green philosophy. However, upon closer examination, they will find that going green is a sound strategy.

Green initiatives let us revisit and examine our IT systems and their operations in terms of energy efficiency and resource utilization, and thus can reduce energy bills. Until recently, very little attention was given to IT's energy consumption, effective use of resources, operational costs and negative environmental impacts during manufacturing, use and disposal. Now, however, a spotlight has been turned on IT, and there is a pressing need to address these overlooked aspects, which are important in safeguarding the environment for future generations.

Businesses also need to look at green requirements from another viewpoint – that is the implications of not going green in the context of stricter environmental regulations, stakeholder demands, competitiveness, branding and corporate image and social responsibility. Smart companies will adopt an environmental strategy to innovate, create value and build a competitive advantage. They will benefit by viewing these challenges as strategic opportunities.

Again, the greening of and by IT will soon be necessities – not options. To help create a more sustainable environment, IT professionals must understand green IT and its potential. Publications that describe new advances, outline current trends and present solid case studies demonstrating green IT's benefits will help provide this understanding and the motivation to 'green' IT.

1.12 CONCLUSION

As the climate debate heats up, IT finds itself part of the problem – and part of the solution. Environmentalism and economic growth can go hand in hand in the battle against global warming.

A vigorous green IT plan is an economic – as well as an environmental – imperative.

Companies can out compete their peers by tackling sustainability head on, engaging stakeholders, developing partnerships and adding environmental stewardship to their corporate culture. Every business, big or small, faces environmental risks and opportunities. Companies have benefited from taking these challenges as strategic opportunities (Esty and Winston, 2006).

Businesses must develop a positive attitude towards addressing environmental concerns and adopt forward-looking, green-friendly policies and practices. The challenges are immense; however, recent developments indicate that the IT industry has the will and conviction to tackle these environmental issues head on.

1.13 EXERCISE

1. Briefly describe climate change, global warming, greenhouse gases and the greenhouse effect.
2. What is meant by green IT? Why is it gaining greater relevance and importance now?
3. What are the different dimensions or directions of green IT?
4. How can software impact the environment and the energy consumption of computing systems?
5. What are key subsystems of IT that could be made greener? Briefly explain.
6. Describe the 3Rs of green IT.
7. Why there is growing demand, and need, for greening data centres?
8. What is meant by 'green washing'? Explain with examples.

1.14 REFERENCES

- Harnessing Green IT: Principles and Practices, San Murugesan, G. R. Ganadharan, Wiley & IEEE.
- Green IT, Deepak Shikarpur, Vishwkarma Publications, 2014
- Green Communications: Principles, Concepts and Practice-Samdanis et al, J. Wiley
- Green IT for Sustainable Business Practice: An ISEB Foundation Guide, Mark G. O'Neill, The Chartered Institute for IT, 2010

GREEN DEVICES AND HARDWARE

Unit Structure

- 2.0 Objective
- 2.1 Introduction
- 2.2 Life Cycle of a Device or Hardware
- 2.3 Reuse, Recycle and Dispose
- 2.4 Summary
- 2.5 Exercise
- 2.6 References

2.0 OBJECTIVE

- Introduces the concept of green devices and hardware.
- Sensitizes the reader to environmental issues arising from electronic devices.
- Describes the life cycle of electronic devices.
- Examines electronic devices' impact on the environment during each phase of their life cycle, and possible causes for this impact.
- Discusses a range of solutions to address this problem and highlights best practices in each life cycle phase

2.1 INTRODUCTION

Electronic devices have become ubiquitous and are an intrinsic part of our lives. Whilst these devices provide the convenience of faster and better access to people, information and services, the downside is their negative impact on available resources and our environment. With the threat of global warming looming large, prudent use of these devices becomes urgent. Awareness of our usage patterns' impact and ways to minimize the impact comprises the first step towards a more sustainable practice.

The number of computers and other electronic devices in use has been increasing exponentially, and newer more powerful devices continue to replace older versions. This has led to a very short useful lifetime of devices, leaving behind a trail of obsolete devices. During the early stages of the electronic revolution, also called the digital revolution, the manufacture, distribution, use and disposal of devices had only a small impact on our environment. For instance, power consumed by these devices during use was not a major constraint. Development efforts were focussed primarily on improving processing speed, increasing device

density (number of transistors per integrated circuit) and reducing the cost of production. Until recently, little attention was paid to minimizing or optimizing devices' power consumption. With the advent of battery-powered devices such as notebook computers and mobile phones, power consumption in these devices started to become an important consideration. However, environmental impact, especially in the disposal stage of the devices, was largely ignored.

These features coupled with users' behavioral changes can go a long way in optimizing devices' energy consumption. The process of manufacturing the devices also has a significant environmental impact. It is important to examine manufacturing processes to ensure minimal negative impact. Yet another dimension to the phenomenon of unprecedented growth in the number of devices is the accumulation of electronic waste (e-waste) as these devices become obsolete. Given that many devices, especially ones built during the early phase of the electronic revolution, contain toxic materials that are potentially harmful, disposing them safely poses a major problem. This makes it imperative to closely examine the materials used in manufacturing to ensure minimal impact on the environment even during the disposal stage. The problem of e-waste disposal is further aggravated by the proliferation of newer, better devices.

With these backdrops, we explore 'green IT' from different perspectives – that of the producer and the user. Following the green IT definition of Murugesan (2008), we use the term green devices and hardware to explore the aspects of 'greenness' in hardware. In this chapter, we explore the environmental impact of electronic devices, primarily hardware devices, at every stage of their life cycle. We examine the manufacturing process in an attempt to identify best practices. We also explore resource usage patterns in bringing products to market. Further, we look at usage patterns and the adoption of environmentally friendly practices, and make the reader sensitive to these issues.

2.2 LIFE CYCLE OF A DEVICE OR HARDWARE

A green device cannot be built by just having an additional step in the life cycle of the device. There needs to be a concerted effort at every stage of the device life cycle – from the moment the device is conceived, to its development, to the time when it is used and recycled or disposed (i.e. from cradle to grave). Each stage of the cycle has varying levels of impact on the environment. In this section, we will look into each stage that a computer device goes through and discuss green considerations for each stage. The typical life cycle of a device, shown in Figure 2.1, consists of five stages:

1. Design.
2. Manufacture and facilities.
3. Packaging and transportation.
4. Usage.

5. Reuse or disposal.




Figure 2.1 : Life Cycle of a Device or Hardware

2.2.1 Design

In the design stage, the idea is conceptualized and the device is designed, prototype and tested. Though this stage does not have direct impact on the environment, design stage decisions such as those regarding architecture, constituent components or materials and layout have a huge effect on the environmental impact of other stages. Hence, it is important to design devices so as to keep their environmental impact as low as possible during subsequent stages whilst meeting performance and other requirements. We will look at a possible process improvement that will help design a green device.

Figure 2.2 shows the typical steps in the design process. When a device gets into the design stage, the first step should be to set environment targets for it (Samsung, 2011). As indicated in Figure 2.2, setting environmental targets should be done in parallel to conceptualization. These targets can be derived from the device's environmental objectives, and also based on environmental impact assessment reports of similar devices that are in the market. These form a minimum set of targets that the device has to meet when released to the market. These targets should act as a benchmark and should help with decision making whilst designing the device.

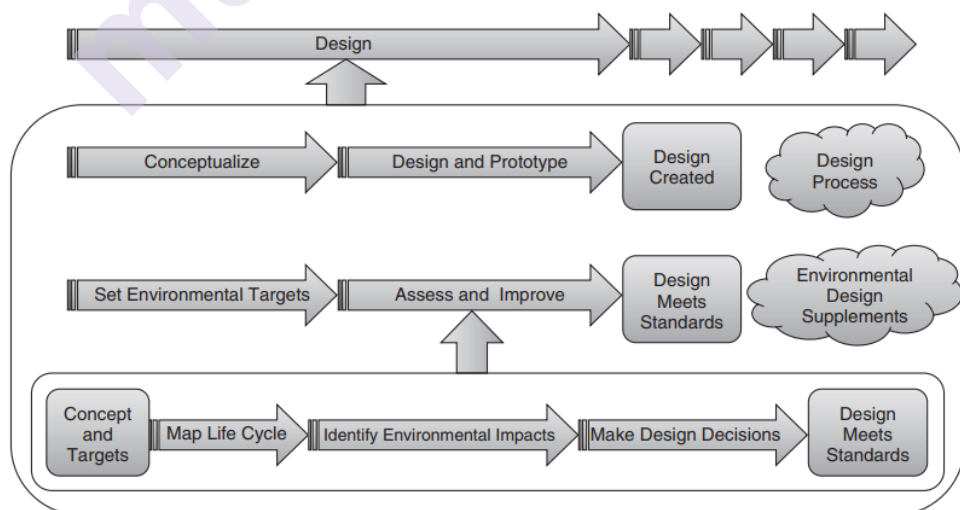


Figure 2.2 Environmental Targets

design and improve it until the targets are achieved. To assess and improve the prototype, we need to take a holistic approach to the device's life cycle (Sustainability Victoria).

As shown in Figure 2.2, we first need to map the device's entire life cycle to understand the places where we perceive and estimate its environmental impact. This assessment should include each stage of the life cycle from manufacturing and transportation to usage and disposal. Effort should be made to quantify the impact at every stage of the life cycle. This will help us identify main sources of environmental impact. Design alternatives to reduce the environmental impact need to be considered. Design changes have to be implemented and the prototype has to be assessed again for environmental impact. This process should go on until the device meets the environmental targets set for the device.

2.2.2 Manufacturing

The manufacturing process is one of the main sources of environmental impact in the life cycle of a device. Manufacturing processes are resource intensive and consume a lot of raw materials, water and energy; they create many different categories of waste, some which are toxic. It is very important to minimize the use of environmentally sensitive substances and reduce the amount of harmful waste.

The EPEAT (Electronic Product Environmental Assessment Tool, n.d.) ratings system serves as a good guideline for understanding how green a device is. EPEAT grades electronic devices and computer systems into three grades (bronze, silver and gold), helps customers identify how green a device is and helps them to make informed choices.

The rating system takes into account the following aspects:

- Reduction or elimination of environmentally sensitive materials;
- Material selection;
- Design for end of life;
- Device longevity or life extension;
- Energy conservation;
- End-of-life management;
- Corporate performance;
- Packaging.

Since it covers most aspects of a manufacturing process, EPEAT serves as a good way to assess the greenness of a device with respect to its manufacturing. Electronic devices may contain hazardous and environmentally sensitive materials like lead (Pb), cadmium (Cd), mercury (Hg), hexavalent chromium (Cr6+), polybrominated biphenyls (PBBs), polybrominated diphenyl ethers (PBDEs), arsenic and polyvinyl chloride (PVC). It is important to reduce, if not eliminate, these materials both in the device and in its manufacturing process. The Restriction of Hazardous

Substances (RoHS, n.d.) directive restricts the hazardous substances commonly used in electronics and electronic equipment.

The impact of some of those key materials is as follows:

- Lead is generally used in printed circuit boards and in glass monitors. The pathological effects of lead are most prominent in the nervous system, hematopoietic system and kidneys. Excessive exposure to lead (Pb) results in clinical toxicity. There has been extensive research to find a replacement for lead (Norwegian University of Science and Technology, 2010). There are also on-going efforts to recover lead from the e-waste (Duffy, 2008).
- Cadmium is generally used in low-temperature soldering, plating for corrosion protection, colorants in plastics and contact buttons in relays (Penica and Hilty, 2004). Cadmium when ingested affects the liver and kidneys. Research is on-going to produce alternatives for its various uses.
- Mercury is mainly used in monitors and batteries. Mercury is a well-known pollutant (US Geological Survey, 2000), and it affects the immune system, alters genetic and enzyme systems and damages the nervous system. New light-emitting diode (LED) backlit technologies have been developed which make monitors mercury free. Mercury use in batteries has also been reduced considerably over the past few years.
- PBDEs and PBBs are used as flame retardants (Wisconsin Department of Natural Resources, 2011). These materials are ubiquitous environmental pollutants; they are extremely toxic and are banned in many countries. Efforts are on-going to remove these materials from the manufacture of computers.
- Arsenic is used in glasses and in manufacturing semiconductors. Arsenic interferes with cellular longevity by allosteric inhibition of an essential metabolic enzyme. Many arsenic prevention and recovery methodologies have been devised, and these need to be actively deployed to reduce arsenic usage in the manufacturing process.
- PVC is one of the most commonly used plastics worldwide. It has been used extensively in the manufacture of computer parts. Vinyl chloride, one of the main ingredients of PVC, is a carcinogen and also affects the human reproductive system (US Environmental Protection Agency (EPA), 2011). Of late, many companies have eliminated the use of PVC in their devices.

The EPA Web site gives further information on these and other toxic materials and how each of these chemicals affects the environment. It is important that electronic device manufacturers avoid using these chemicals in their devices and use environmentally friendly alternatives. Further, their facilities' carbon footprint can be reduced by employing energy conservation techniques and by using energy from renewable sources. Energy conservation techniques like sensors for switching off

lights automatically, and shutting down servers during weekends (where possible), not only help reduce the carbon footprint but also result in lower energy bills. In addition, environmentally sensitive materials like chlorofluorocarbons (CFC) should be avoided in air conditioning and cooling facilities. Such measures also help companies in projecting a good public image.

2.2.3 Packaging and Transportation

Packaging and transportation also contribute to the carbon footprint in a device's life cycle. The two main contributors in this segment are the materials used for packaging and the carbon footprint of the vehicles used in transportation.

Since the materials used in packaging a device have an effect on the device's carbon footprint, the amount of material used for packaging needs to be kept at a minimum. The packaging should be done in such a way that it minimizes the amount of material used whilst ensuring the integrity and security of the device.

The size of the device has an effect on the amount of packaging required, so even during the design stage, effort must be made to make the device as compact as possible. Eco-friendly materials like recycled paper, potato starch and recycled board can be used as packaging materials, and soy ink can be used for printing. The following are general recommendations for designing suitable packaging (US Environmental Protection Agency, 2011):

1. Packaging materials should be recyclable.
2. The amount of packaging materials should be kept at a minimum.
3. Various materials used in packaging should be easily separable to ease the recycling process.
4. Adhesive use should be reduced by using folds and tabs instead.
5. All the additives, coatings and inks that get added to the package should be eco-friendly.
6. Printed documentation can be avoided wherever possible. Instead, documentation could be provided through the device's Web page.

In addition, freight transportation leads to a lot of carbon emission. Along with making the transport system efficient, the design of the device should facilitate smarter transportation. For example, the smaller the device, the more of them can be transported in a given space, thus reducing the number of vehicles required to transport them from manufacturing facilities to distribution centres.

2.2.4 Use

Quite a significant amount of energy is consumed by devices when they are being used (powered). This increases a device's carbon footprint, thus resulting in a profound impact on the environment. The Energy Star rating system helps customers choose the most energy efficient devices and, thereby, reduce their energy consumption and cost.

Often people tend to apply a usage pattern without being aware of a device's characteristic. For example, the commonly held belief that, to ensure longer life of a battery, the battery needs to be drained completely before recharging is not necessarily true for all battery types. In case of lithium-ion batteries, commonly known as Li-ion batteries, the maximum number of charge cycles decreases with increase in depth of discharge (Buchmann, 2011). Since the best usage patterns are tightly tied with a device's characteristics, it is important that manufacturers provide guidelines on its best usage practices from environmental and energy perspectives.

Devices like 'Kill a Watt' (P3 International, 2011) help measure the amount of energy consumed by a device over time. When 'Kill a Watt' is plugged between the device and the power socket, it measures the amount of energy that the device consumes. This can be used to measure the amount of energy consumed over time and with different configurations of the device, and thus helps one find an optimum configuration that consumes the least energy. It also helps one figure out whether older devices are still energy efficient.

The amount of energy consumed by different kinds of electronic devices varies widely, and ways to reduce their energy use vary widely too, hence we will look into different device categories separately, and look at ways to reduce energy usage in notebook computers, desktop computers, servers, mobile devices and other special devices.

2.2.4.1 Notebook Computers

Though notebook computers are designed to be energy economical compared to desktops, the energy consumption of a laptop depends on the usage pattern. To reduce energy consumption, it is necessary to have a usage pattern that results in optimal energy consumption. Let's examine the power consumption of a laptop's different components and discuss how each component's power consumption could be reduced or optimized.

Notebook computers are generally powered by a rechargeable battery, which when in good condition powers the device for about 3–5 hours. The battery is recharged with the help of a charger, which can charge the battery irrespective of whether the device is running or is in a powered-down state. The lifetime of a battery is generally short: They last for only about 300–600 charge cycles. This results in the need to replace the battery a few times during the life of a laptop, thereby contributing to e-

waste. A few companies are also coming out with notebook computers which run on solar power.

Notebook computers are also characterized by their liquid crystal display (LCD) or LED monitors, which are integrated into the device. The monitors and CPU consume the largest share of power that is utilized by the laptop. The larger the monitors, the greater the laptop's power consumption and the more intensive the operations performed by the processor, the higher its power consumption.

Energy can be conserved in a laptop in several ways. According to various studies, quite a lot of the energy wasted in a laptop happens when the laptop is being charged. The chargers convert the AC to DC and steps down the voltage. This happens as long as the charger is connected to the power socket, irrespective of whether the laptop is connected to the charger or not, thus resulting in wasted energy. Hence it is important to switch off the power supply and unplug the charger from the power socket when it is not in use. This will reduce power wastage. Of late there have also been green chargers (such as the Go Green Charger) available which can detect whether a charger is connected to a notebook computer or any other device, and reduce the power consumption when a charger is not connected to a device.

Monitors consume about 20–30% of the total energy used by a laptop; hence it is important to reduce their power consumption. Strategies that help in reducing their power consumption include the following:

1. Reduce the brightness of the monitor to an appropriate level. A brighter screen consumes more energy.
2. When some background task is running on the computer and there is no need to use the monitor during this time, switch off the monitor instead of using screen savers as screen savers also consume some energy.
3. Most computer operating systems provide power-saving profiles which, when enabled, reduce the amount of energy consumed by the computer. For example, when the laptop is starting up or shutting down, these applications reduce screen brightness to a minimum.

Laptop processors also consume a lot of power particularly when carrying out computationally intensive tasks such as encryption, analytic, computer games and image and video processing. Many aspects of applications that run on a system impact power consumption. The following guidelines will help in optimizing energy consumption:

1. The background processes and other applications which are not being used are kept running to keep the processor active, thus resulting in energy wastage. So when an application is not in use, close the application and also stop the background processes that are not being used.

2. Multitasking is a trade-off between CPU time spent executing tasks and inefficiencies brought in by context switching. When a task starts running slower than it would normally run, one of the likely causes is due to inefficiencies due to context switching. It is advisable to reduce the number of tasks when such slowing down is noticed.
3. The more processes that are set to start at start-up, the longer amount of time the laptop needs to boot up, thus resulting in more unproductive time for the processor and the monitor. So keep the number of processes at start-up to a minimum. The required applications can be started when needed.
4. When playing games or other multimedia applications, the greater the level of detail, the greater the processing power consumed. So it is better to keep the level of detail to the required level; most games support this feature.

The higher the speed at which a processor runs, the more energy is consumed, so advanced users could run the processor at a lower frequency (i.e. 'under-clocking' it) to save power. This can be done by going to the basic input-output system (BIOS) of the laptop and setting the appropriate processor speed. In most computers going to the BIOS is easy, though it depends on the computer configuration. In contrast, if a computer is perceived to be slow and you are contemplating replacing the processor, a better option

would be to try over-clocking the processor to see whether the increased speed meets the need. Over-clocking a processor can harm the computer, so practice caution and have a computer expert do it. In addition multicore processors consume relatively less energy for the same amount of processing done by multiple single-core processors. This is primarily because of the decreased power required to drive signals external to the chip.

The other major component of the laptop that consumes a lot of power is the hard disk. Since hard disks are physical devices, they consume relatively more energy than many other components. Hence it is important to keep hard disk access to a minimum and use it only when necessary. The spinning of the hard disk is one of the important parts that results in higher energy usage. Lessening the spin of the hard disk results in lower energy consumption. The spin of the disk is greater when the files are fragmented and scattered all over the hard disk. Defragmenting the hard disk reduces the spin of the hard disk and will result in less energy usage. Solid-state drives (SSDs) are known to consume less energy than comparable hard disk drives (HDDs). In addition, many power-saving options are available on computers that switch off the hard disk when it is not in use. For detailed coverage of energy-efficient, environmentally friendly computer storage, refer to Chapter 6.

Peripheral devices connected to a computer also consume energy, even though they might not be in active use. For example, devices connected to

USB ports in a computer draw power even when the device is not in use. So when these external devices are not in use, it is better to unplug them.

Most notebook computers can operate in at least two idle power modes – standby and hibernate. These modes help keep the laptop's energy consumption to a minimum whilst still retaining the state of the computer. In standby mode, the laptop's internal devices and optical drives are powered off, but the power is still maintained to the RAM, where the state of the device is preserved to facilitate instant resumption. Though this state requires considerably less power compared to when the computer is running, it still requires some power, which needs to be supplied by the battery or an external power source, or else the system will shut down immediately, thus losing the state. Retaining the state is very helpful when we are away from the computer for a short period of time.

The hibernate mode makes the laptop completely shut down whilst still retaining the powered-on state, thus reducing the start-up time. When a laptop goes to hibernate mode, the current state of the machine is stored in the hard disk, and all the devices including the RAM are powered off. When the computer is resumed from this state, the state information from the hard disk is reloaded into RAM. It takes more time to resume from hibernate mode than from standby mode, but it is more energy efficient. This mode can be used when the state of the computer needs to be retained for a longer period of time. Most laptops today are provided with power management features and software. Power management software can help in regulating the use of the battery and electric power. It monitors the load on the hard disk, the activities on the laptop and the ambient brightness, and changes various computer settings to make optimal use of power. This software also allows users to set timings when the computers can automatically go to standby or hibernate mode. For example, notebook computers can be set to go to standby mode after 20 minutes of inactivity or when the laptop lid is closed. This helps in automatically managing the power consumption of the laptops.

2.2.4.2 Desktop Computer

Desktops are personal computers that are meant for use from a single location and are not portable. Though these devices are being replaced by laptops, there are certain applications where these devices are still used, such as intensive processing applications (but not in a scale where servers are required) like image processing.

The architecture and physical configuration of desktops support upgrading and replacing individual defective sub-units and parts. Since desktops are quite rapidly being discarded in favour of laptops, the biggest problem facing the industry is disposing of or recycling these devices. In addition, quite a few desktop users are home users who are not fully aware of disposal mechanisms unlike the enterprises that are obliged to safely dispose of their assets.

One of the areas where a desktop varies from laptops is in its mode of usage. Desktops are powered from external power, so, in many cases, they

are kept always on so that they can be remotely accessed from a laptop or some other desktop. In most cases desktops are left on 24/7 even during holidays. Hence most of the time desktops remain idle and waste a lot of energy. To avoid having desktops switched on all the time, remote wake-up methodologies have been devised. In the 1990s, Advanced Micro Devices (AMD) and several other companies came together to develop the magic packet technology to accomplish this. It is commonly referred to as Wake On LAN (WOL). This relies on a special kind of packet called the magic packet to wake up the computer. When the computer is in sleep state, the Ethernet Network Interface Controller (NIC) is left powered on. This network interface controller can detect the magic packet and trigger an interrupt which wakes up the desktop. Using this technology, desktops can be awakened at any time from remote locations, thus removing the need to have the desktops switched on all the time.

Though most operating systems have group policies which allow certain configurations to be enforced on desktops, they are not as effective as management tools which allow granular control over the desktops. Tools like Night Watchman help enterprises have more control over the desktops. These tools allow enterprises to securely, remotely and centrally power down desktops. They also allow enterprises to apply power schemes at different times in multiple locations, globally from a single console, maximizing power savings without impacting users. Though there might be resistance from employees to the implementation of such systems, it is important to make employees understand the benefits and advantages of the system and then implement solutions.

2.2.4.3 Servers

Servers are computers designed to serve the needs of other computers. In general, server computers run one or more services that will be used by other computers in the network. They have powerful CPUs and a large amount of memory (RAM). They are generally placed in racks. Examples of servers are mail servers, database servers, file servers and others. They are on most of the time, and in most cases to provide increased reliability redundant servers are used. These servers are generally placed in data centres. Even though the growth of energy usage of data centres was less than predicted, it still accounts for between 1.1% and 1.5% of total electricity use.

2.2.4.4 Mobile Devices

Mobile devices comprise mobile phones, personal digital assistants (PDAs) and other smart devices that people carry around. These devices are small in size and are powered by a small rechargeable battery. These devices are produced in large numbers and are designed for operation with low power consumption. These are used by people all over the world. www.it-ebooks.info 34 Harnessing Green IT the world from both developed countries and developing countries. The International Telecommunication Union (ITU) estimates that by October 2010, over 5 billion mobile phones were in use (International Telecommunication Union, 2010).

Mobile devices use external chargers for their batteries. The power efficiency of these chargers is low, and most people keep the charger powered even whilst mobile devices are disconnected after being charged. Chargers consume power even whilst not charging a battery and hence waste power. So, chargers should be switched off or unplugged when they are not charging, or disconnected from mobile devices. Awareness needs to be created among users to switch off chargers when they are not needed.

2.2.4.5 Specialized Devices

Specialized devices are designed for a specific purpose, such as set-top boxes, play stations and medical equipment like X-ray machines and computerized tomography (CT) scanners. Since these devices are designed to accomplish a specific purpose, the designer can optimize the device for the particular use and for low power consumption. Devices like set-top boxes and play stations consume power even when they are not in use (idle), and idle power could be as much as about two-thirds of the power the device consumes when it is used. For such devices, it is important to switch off the devices so that they do not consume (waste) power when not in use. Play stations in general have powerful graphics processors which consume a lot of power. It is a common practice among users of play stations to leave them in pause mode for long periods of time, which results in power wastage. It is important to create awareness among users of special devices to switch them off when they are not in use.

2.3 REUSE, RECYCLE AND DISPOSE

Mobile phones and computers, as well as some other electronic gadgets, become obsolete quickly, typically about every two years, due to continued advancements in technology and the introduction of new gadgets with enhanced features. This leads to an accumulation of much e-waste. Most of these devices end up in landfills or get exported to developing countries for recycling. In many developing countries, however, due to lack of sufficient knowledge about the toxicity of the chemicals involved and lack of regulations enforcement relating to e-waste recycling, unhealthy recycle methodologies are used, leading to health hazards and pollution.

One way to reduce waste is to increase the lifespan of the devices. Increasing the lifespan of devices saves life cycle energy¹ and resources and reduces the amount of hazardous materials ending up in landfills. Reusing an old computer is a great way to increase its lifespan. Reusing a computer is environmentally friendlier than recycling.

It has been shown that reusing a computer is 20 times more effective at saving life cycle energy than recycling (Computer Aid International, 2010). As there is a lot of environmental cost involved in a device's production, it becomes important to use a computer to its fullest life before discarding it.

1 Life cycle energy is a term used to refer to all energy inputs to a product, not only direct energy inputs during its manufacture, but also all energy inputs needed to produce components, materials and services needed for the manufacturing process.

Once we have decided to replace a device, the first option to consider is to send the computer for reuse to some other place like an educational institution. Both the EPA in the United States and WEEE in the European Union provide guidelines on enabling computer reuse. Before sending the device for reuse, adhere to the following guidelines (U.S. Environmental Protection Agency, 2011):

1. Make sure that the receiving organization does need the device being sent and can reuse it. A better alternative would be to send the device to a local refurbisher who will make sure the device is in a working state and can send it to an appropriate organization that is in need of it.
2. Copy the information that has been stored in the device to a backup device or another device.
3. Wipe off personal or sensitive information in the device using tools specific for this purpose like Symantec's WipeInfo.

Another way to increase the lifespan of a device is by enabling easy upgrade of its various parts. This will encourage people and organizations to use their device longer by upgrading certain parts of their device that have become obsolete. Some design considerations that enable the easy upgrade of a computer are as follows:

1. The device should be modular.
2. Each part of the device should be easy to replace.
3. The device should be easy to disassemble and reassemble. Techniques like screw-less casing make it easy.

Wherever reuse is not a viable option, the next best option is to recycle. In the process of recycling, most of the original device's materials are reused as raw materials for a new device, thus resulting in less waste. To enable recycling, it is important to collect back the old devices from customers. Many manufacturers have their own recycling programs.

Take-back, mail-in and trade-in programs are some of the popular programs that encourage customers to return devices to the manufacturers. In addition, companies also support local organizations to collect the devices. In order to increase awareness, collection events and local recycling events are conducted. These programs help individual users easily recycle their computers.

An electronic device consists of various materials which take quite a lot of energy to extract from ores. During recycling these materials can be recovered and some of the recovered materials can be used in making new

devices, thus reducing the amount of environmental emissions compared to making these devices from virgin materials.

Materials like gold, silver, copper and plastic can be recovered and reused. Some companies use a measurement called recycling rate, which is the weight of the materials the company recycles each year as compared to the total weight of the devices the company sold a few years earlier.

As discussed in this chapter, one of the problems regarding recycling of e-waste is the export of e-waste to developing countries for recycling. Developed countries export their e-waste to developing countries since the cost of recycling is cheaper there. For example, recycling a computer in the United States costs about US\$20.00, whereas recycling a computer in developing countries like India or China costs only about US\$2.00.

But e-waste is not properly recycled in developing countries. In many cases, the labourers involved in recycling simply burn off the e-waste, follow crude unhealthy practices to remove components and metals from e-waste or dump the waste into rivers, resulting in huge environmental impact. They are unaware of the environmental and health hazards posed by the chemicals and poisonous gases released in the crude recycling process they follow. They are not required to wear protective clothing, and thus face severe health problems. Efforts have been made by organizations such as Greenpeace to regulate the export of e-waste to developing countries and to adopt healthy e-waste recycling processes.

Whereas refurbishing and recycling are effective techniques for the end-of-life management of electronic devices, often they are not practical. The only option left at this stage would be disposal. There are a few known techniques of waste disposal that are commonly used in e-waste disposal as well. These are incineration, chemical decomposition and landfill. There is no single silver bullet that works in all cases for all materials or all kinds of e-waste. Incineration may result in the release of toxic gases and can pose a hazard. On the positive side, the heat generated may be used in generating electricity.

Due to the non biodegradability of materials used in electronic devices, landfill may not be a viable option either. Chemical decomposition may be an option in some cases and will have to be executed in a managed manner. In short, safe disposal is a major problem and the solution depends on the specific material used in manufacturing the device.

2.4 SUMMARY

The last couple of decades have seen an explosion in the use of electronic devices. We now face the problem of e-waste emerging as a significant challenge. This, coupled with increasing concerns on global warming, has made the need to manage and contain the environmental problems caused by electronic devices and IT significant and very urgent.

In this chapter, we explored ‘green’ aspects of electronic devices and computers. We examined the environmental problems created by

electronic devices and hardware and explored how the industry is addressing these issues. We outlined best ‘green’ practices from the perspective of manufacturers, and discussed practical green measures that users should adopt. We also outlined the importance of designing a device with environmental considerations in mind so that the entire life cycle of the device could be better managed.

We highlighted some of the major problems we encounter in recycling e-waste. Many governments and other agencies have realized that the cost of inaction is enormous and the consequences of inaction – that is, the resulting environmental damage – may be irreversible. This has led governments, the United Nations and other agencies to come up with guidelines and regulations to address the environmental impacts of electronic devices throughout their life cycle – from design and manufacturing to their use and disposal followed by recycling.

2.5 EXERCISE

1. Briefly outline different stages in the life cycle of an electronic device. Order these stages in decreasing order based on their environmental impact from your perspective, and give your rationale.
2. How would you design the manufacturing process of a product (say, a mobile phone) while also considering its environmental impact during manufacturing?
3. What environmentally sensitive materials are typically used in the manufacturing of a product (e.g. a notebook computer or a smart mobile phone), and what are their impacts? Suggest safer alternatives for each of those materials.
4. What are the tools that can be used to monitor and manage the power consumption, or usage patterns, of an organization’s computers? Briefly describe the features of three of those tools, and compare them.
5. What are the various e-waste disposal techniques, and which is the most effective among them and why?

2.6 REFERENCES

- Harnessing Green IT: Principles and Practices, San Murugesan, G. R. Ganadharan, Wiley & IEEE.
- Green IT, Deepak Shikarpur, Vishwkarma Publications, 2014
- Green Communications: Principles, Concepts and Practice- Samdanis et al, J. Wiley
- Green IT for Sustainable Business Practice: An ISEB Foundation Guide, Mark G. O'Neill, The Chartered Institute for IT, 2010



GREEN SOFTWARE

Unit Structure

- 3.0 Objective
- 3.1 Introduction
- 3.2 Energy-Saving Software Techniques
- 3.3 Evaluating and Measuring Software Impact to Platform Power
- 3.4 Summary
- 3.5 Exercise
- 3.6 References

3.0 OBJECTIVE

- Describes software methodologies, designs and software development tools that help to improve the energy efficiency of application software.
- Discusses methods to reduce energy consumption during software development.
- Recommends a set of tools for evaluating and measuring the energy consumption of software.

3.1 INTRODUCTION

A green IT infrastructure is incomplete without green software. Software plays an important role in overall platform energy efficiency. Whilst platform hardware designers work diligently to reduce components' power consumption, a single ill-behaving or power-unfriendly software component on the system can thwart all of the power management benefits built into the hardware. This chapter describes the characteristics of green software and the software design considerations and methodologies to improve software energy efficiency.

A computing 'platform' is a combination of hardware, software and other technologies that allow software to run. The techniques we describe in this chapter apply to all types of platforms including embedded platforms, smart phones, tablets, netbooks, notebooks, desktops, data centre servers and supercomputers. Every one of these devices runs software, and every one of them requires energy to do it. Most agree that a platform's energy efficiency is extremely important. Nowadays less than a full day's use of one's smart phone is intolerable to most. Notebook computer manufacturers are still striving for the elusive all-day PC, though some of the thin and light models are getting close. A big concern for data centres

is the cost of electricity to run the servers and air-conditioning units to keep the room and servers cool. To make progress in improving energy efficiency, it is important to understand how energy is consumed. Figure 3.1 shows a typical power breakdown for a notebook PC during various operating conditions (provided by Si Software's Sandra benchmark suite). There are many components to target for improvement, and we believe that a focus on software will lead to significant impact.

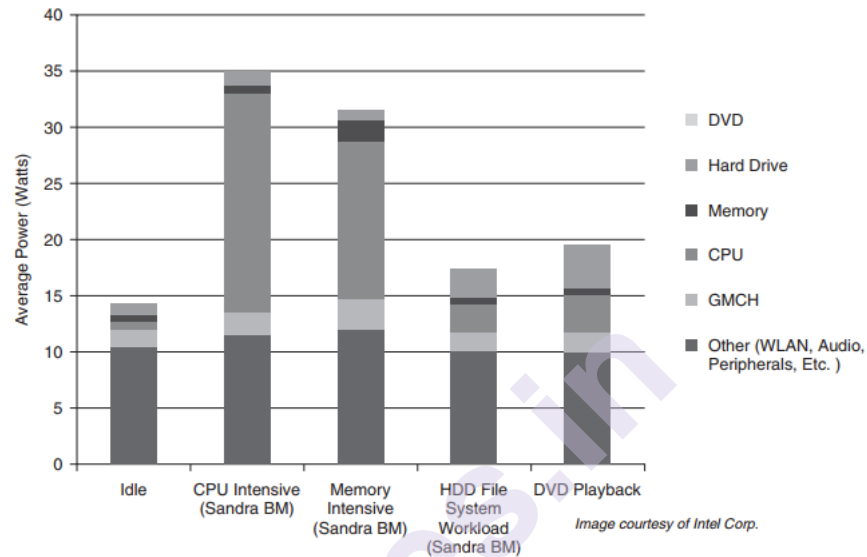


Figure 3.1 : typical power breakdown

3.2 ENERGY-SAVING SOFTWARE TECHNIQUES

Many might wonder why a software developer needs to know about the component energy efficiency and states of the CPU. Whilst most of the energy-saving features in the platform are transparent to the software developer and applications have very little direct control, the behaviour of the software has significant influence on whether the energy-saving features built into the platform are effective. Well-behaved software allows the energy saving features to work. Poorly behaved software inhibits the energy-saving features and leads to lower battery life and higher energy costs. Consider, for example, DVD playback applications; some are more energy efficient than others. A number of techniques are available to the developer to implement a more energy-efficient player. For instance, you could implement a read-ahead buffering methodology and allow the optical drive to take breaks rather than continuously spinning the disk. The application could be developed to use a hardware accelerator for video decode instead of using the CPU. That also saves energy. A third mechanism might be to use an alternate audio decoder when the system is running on battery power, one that may sacrifice a small bit of quality for longer battery life. Software developers have choices and design decisions to make that can influence the energy efficiency of applications.

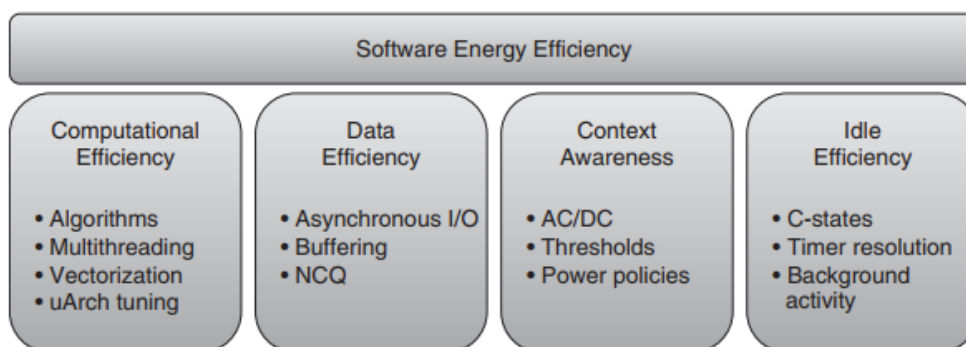


Figure 3.2 Software energy efficiency

Before examining software energy-saving techniques, it is important to understand the distinction between active and idle software. Active software is software that is fulfilling its intended purpose such as by computing a spreadsheet, playing music or a movie, uploading photos to a Web site or browsing the Internet. In all of these cases, there is a workload that the CPU or graphics processing unit (GPU) is busy working on. Idle software is software that is essentially running but waiting for an event to make it active.

Examples of idle software are a browser that is started but has not been pointed to a Web site, an open word processor programme in the background or whose window is minimized, or an instant messaging programme that is running but not sending or receiving a message. The DVD playback scenario provides a good example for both states. It is idle when it is started and not playing a movie, and it is under a workload when it is playing a movie.

Software energy efficiency can be improved by improving computational efficiency, data efficiency and idle efficiency and by enhancing context awareness. Figure 3.2 provides a snapshot of the techniques that can be applied in each of these areas.

3.2.1 Computational Efficiency

Simply put, computational efficiency means getting the workload done quickly, with minimal energy consumption. The amount of effort that computer scientists have put into software performance saves not only time but also energy. We call this the race to idle. The faster we can complete the workload and get the computer back to idle, the more energy we can save. For faster processing, approaches such as efficient algorithms, multithreading, vectorization and uArch tuning are used.

3.2.1.1 Efficient Algorithms

Considerable effort has gone into research to find more efficient means to solve problems and to investigate and document the corresponding time and space trade-offs. It is clear from computer science theory that the choice of algorithms and data structures can make a vast difference in an application's performance. All other things being equal, an algorithm that

computes a solution in $O(n \log n)$ time is going to perform better than one that does the job in $O(n^2)$ time. For a particular problem, a stack may be better than a queue and a B-tree may be better than a binary tree or a hash function. The best algorithm or data structure to use depends on many factors, which indicates that a study of the problem and a careful consideration of the architecture, design, algorithms and data structures can lead to an application that delivers better performance.

Figure 3.3 illustrates how better algorithm performance leads to energy savings. The two diagrams in Figure 3.3 show the simplified, hypothetical power usage over time of two algorithms operating on the same workload. The power at idle for this computer is about 8 W. The areas within the dotted lines represent the total energy consumed over 12 sec. In the first case the total energy was 220 J, and in the second case the total energy was 112 J. The better algorithm saves time and energy. Of course, the big hole in our logic here is the assumption that both algorithms push the system to consume an identical 20 W. But what if the more time-efficient algorithm has a complexity that drives the system to consume more power than the slower algorithm?

Could the total energy for the more time-efficient algorithm be worse? The answer is probably yes, but we believe that in the vast majority of cases, the race to idle yields favourable results, and in the rest of this chapter, we'll show a number of case studies that demonstrate why we believe that.

3.2.1.2 Multithreading

Multithreading the software delivers better performance and as well as better energy efficiency. To help you take advantage of the multiple cores available in modern compute platforms, many threading methodologies and libraries such as OpenMP, OpenCL and Thread Building Blocks are available. To illustrate how well-developed, balanced multithreading saves energy, consider a favourite performance benchmarking application – Cinebench 11.5 from Maxon Computer GmbH – and the measured energy consumed on a four-core, eight-thread Intel processor. The hypothesis: Completing the Cinebench workload with the full four cores and eight threads available will save energy versus completing the workload with a single thread. As you can see, completing the workload with a single thread takes considerably longer and uses more energy than any of the multithreaded runs. In fact, the eight-thread case (normalized) uses about 25% less energy than the single-threaded case. A well-balanced multithreaded workload is more energy efficient than running the same workload with a single thread. The added advantage is that the eight-thread run was completed in about one-fourth the time and the processor is available for other computation.

3.2.1.3 Vectorization

Another method used to achieve better computational efficiency is by vectorizing the code instead of using scalar C-code by using advanced instructions such as single-instruction multiple data (SIMD) for

instruction-level data parallelism. If you can vectorize your solution, you will get better performance and a corresponding power benefit. To test this, we took two different audio decode algorithms and optimized them using Intel's Advanced Vector Extensions (AVX) instruction set. Using a method we have to turn the AVX feature on and off, we measured both the performance and the average power of an audio decode workload. When AVX is enabled, the workloads run considerably faster – by 1.65X and 1.34X – with a huge savings in power. Table 3.1 summarizes the performance and power impact of using AVX.

3.2.2 Data Efficiency

Data efficiency reduces energy costs by minimizing data movement and delivers performance benefits. Data efficiency can be achieved (Arnout, 2005) by designing:

- Software algorithms that minimize data movement;
- Memory hierarchies that keep data close to processing elements;
- Application software that efficiently uses cache memories.

Methods for achieving data efficiency include managing disk input–output (I/O), using block reads, native command queuing, file fragmentation, disk I/O with multithreaded code as well as pre-fetching and caching.

3.2.2.1 Managing Disk I/O

This section summarizes the analysis of a disk's power characteristics during sequential and random reads and native command queuing, and provides an analysis on file fragmentation and disk thrashing. This section also provides guidelines on optimizing the power during disk I/O in various usage models, along with the power impact. For additional details as well as sample code, see Krishnan and De Vega (2009).

The analyses are based on the typical performance characteristics of hard disk drives (HDD) which are affected by rotational speed, seek time, rotational latency and the sustainable transfer rate. Furthermore, the actual throughput of the system will also depend on the physical location of the data on the drive. Since the angular velocity of the disk is constant, more data can be read from the outermost perimeter of the disk than the inner perimeter in a single rotation. When a read request is placed by an application, the disk may have to be spun up first, the read-write head must be positioned at the appropriate sector, data are then read and optionally placed in an operating system (OS) file system cache and then they are copied to the application buffer. Table 3.2 shows the relative time (in milliseconds) involved in these operations based on the theoretical specification of the serial advanced technology attachment (SATA) drive used. Whilst the actual numbers will vary between different drives, this gives a relative idea of the times taken. Figure 3.5 shows the average power consumed during idle, read-write and spin-up.

It is clear from the data that disk spin-up takes the most time and consumes the most power. Applications performing disk I/O should take this power profile into consideration and optimize for the power and performance of the hard disk. To help better understand the energy usage of HDD using various I/O methodologies, we present below the setup, results and recommendations of five separate experiments:

- Impact of block size on sequential reads;
- Effect of buffering during multimedia playback;
- Impact of file fragmentation;
- Impact of native command queuing on random reads;
- Disk I/O in multithreaded code.

3.2.2.2 Pre-Fetching and Caching

We briefly outline outcomes of a study to determine if pre-fetching and caching can save energy during DVD playback (for details, see Chabukswar, 2009). The study analysed the power consumption of three different DVD playback software applications (DVD App 1, 2 and 3) with the multiple out-of-box configurations available whilst taking power measurements – primarily maximum power-saving mode versus no power-saving mode.

From the results of the studies performed, three guidelines emerge that can help save energy during DVD playback:

- **Buffering:** Apply buffering techniques, which can reduce DVD power consumption by 70% and overall platform power consumption by about 10%, compared to other techniques.
- **Minimize DVD drive use:** Reduce DVD spin-up, spin-downs and read accesses in order to save energy.
- **Let the OS manage the CPU frequency:** Allow the OS to set the appropriate P-state, adjusting the CPU frequency as needed. (We do not recommend changing the CPU power scheme to run the processor at the highest available frequency.)

3.2.3 Context Awareness

Humans naturally use context to understand our world, makes decisions and adapt to the environment. Context awareness in computers means that they can sense the environment in which they are operating and software can be designed to react to changes in the environment. Context awareness was first introduced by Schilit, Adams, and Want (1994, p. 89). The objective is to create applications that can respond or adapt to changes in the environment.

For the physical environment, this requires sensors and the ability to generate events or state changes to which the applications can react. Examples of context-aware behaviour are as follows:

- A PC or smartphone warns you when your battery has reached a low-energy state.
- A notebook PC responds to a change from AC to DC power by automatically dimming the display.
- A tablet PC or smartphone responds to ambient light level and adjusts display brightness.
- A notebook PC quickly parks the hard drive heads when sensors detect that the device is falling – to avoid a head crash.
- A handheld device writes cached data to flash memory when the battery is getting critically low.

Embedded systems are particularly context aware because, in many cases, they are designed specifically to monitor environmental conditions from sensor data and react to them. The use of sensors is growing rapidly in smartphones and tablets and includes light sensors, gyros, accelerometers, Global Positioning System (GPS) receivers, near-field communications and others.

3.2.3.1 Awareness of Power Source

Does it benefit an application to know the source of power – for example, if a notebook PC is plugged into an AC power source or operating on battery – and accordingly tailor its operational modes or application behaviour? Consider the following scenarios:

- Should your virus checker start a full-system scan when you are running on battery?
- Should your BluRay movie player application provide you with the option to deliver longer playback when operating on battery, possibly sacrificing some quality?
- Are you willing to give up some special effects in games to get more game play on a single battery charge?

If you answer yes to any of these questions, then it is important for the application to be aware of whether the platform is plugged into a continuous power source (AC) or operating on battery (DC). In Windows, you can achieve this by querying a unique GUID called GUID_ACDC_POWER_SOURCE (Microsoft Corporation, 2007). Armed with this information, you can adapt the application's behaviour and possibly deliver extended battery life for the use.

The logical next step is to receive a notification when the power status changes. To get this notification, your application must register for the event. Important event notifications include the switch from AC to DC power and reaching a battery threshold. Below is some sample code to register for the AC–DC change event. More details can be found at Microsoft Corporation (2010a).

```
HPOWERNOTIFY ACDCNotificationHandle;
ACDCNotificationHandle = RegisterPowerSettingNotification(
GUID_ACDC_POWER_SOURCE,
DEVICE_NOTIFY_WINDOW_HANDLE);
```

3.2.3.2 Platform Power Policies

Microsoft Windows provides built-in power policies – ‘High performance’, ‘Balanced’ and ‘Power saver’. They give the system user the option to choose between better performance and better battery life. Application software can use power policies in the following ways:

1. Adjust application behaviour based on the user’s current power policy. Windows provides a method to query for the current power policy. Similar to the query for the AC–DC setting, the application would query a GUID called GUID_POWERSCHEME_PERSONALITY.
2. Change application behaviour in response to a change in power policy. This requires that the application register for an event notification. To register for that event, your application would use a mechanism that would look similar to this:

```
HPOWERNOTIFY ACDCNotificationHandle;
ACDCNotificationHandle = RegisterPowerSettingNotification(
GUID_ACDC_POWER_SOURCE,
DEVICE_NOTIFY_WINDOW_HANDLE);
```

3. Change the power policy to suit the application behaviour. The application must first determine if the system will allow a policy change and then make the change if necessary. When the application is finished, it is best practice to return the power policy back to its original state. To change the power policy, you must call the Win32 application programming interface (API) PowerSettingAccessCheck to determine if the user has access rights to override the power settings. This API call queries for a group policy override for specific power settings.

```
DWORD WINAPI PowerSettingAccessCheck (
__in POWER_DATA_ACCESSOR AccessFlags,
__in_opt const GUID *PowerGuid);
```

The API will return the active power settings identified by the GUID, allowable values and default values for AC and DC. For more information on how to use this API, see Microsoft Corporation (2011).

3.2.3.3 Other Context-Aware Behaviours

Developers may also consider the status of other components on the platform and use that information for intelligent application behaviour and

energy savings. It might be useful to know information about the status of components such as network cards, Bluetooth, Wi-Fi, universal serial bus (USB) devices and monitors. For example, should an application continue to render video if the monitor has turned off due to a timeout or if the application window is minimized?

Another set of devices to consider are local area network (LAN) cards and radios. Networking increases energy consumption not just because the LAN card uses energy for transmitting and receiving data, but also because it will remain in an active state for a long period of time on the chance that some useful network communication may occur.

PCs and set-top boxes are the most notable examples of this. The Microsoft System Event Notification Service (SENS) can help alleviate this and other mobile application issues. The SENS API, distributed with the Platform Software Development Kit (SDK), provides a simple function call for checking if the network connection is alive and another that will ping a specified address for you. In addition to these functions, you can also register with the service to receive events when a connection is made or lost, to ping a destination or even as an alternative method to detect when the system changes power states (battery on, AC ON or battery low).

3.2.4 Idle Efficiency

Idle Power for mobile platforms is defined as the power consumed when the system is running in ACPI S0 state with software applications and services running but not actively executing workloads. In this state, there should be minimal background activity. Figure 3.6 shows energy measurements taken whilst a computer was running an office productivity suite. The idle floor is about 8 W. By examining the C-state data from this run, we determined that even with application activity, the platform was in C0 only 5–10% of the time. In other words, the processor was in an idle state, something lower than C0, 90–95% of the time. The challenge is to lower the idle floor by improving application idle efficiency which will lead to a significant increase in battery life. This also benefits average power scenarios and helps all but the most demanding (thermal design power (TDP)–like) workloads.

3.2.4.1 Deep C-State Residency

One of the key requirements for achieving idle efficiency is to keep the platform in deeper C-states for as long a duration as possible. For a platform in idle state, ideally the residency in the deepest C-state (C6/C7) should be more than 90%. Software should aim to keep the number of C-state transitions coming and going out of deep C-state as low as possible. Frequent C-state transitions from C6/C7 to C0 are not energy efficient. Activity should be coalesced whenever possible to allow for higher residencies in deep C-states. This concept was verified by measuring captured power data from the idle platform and a singlethreaded application with two processes that communicate frequently with each other, each waiting for the other to complete. Each process ran for a very

small duration ($\sim 50 \mu\text{s}$). In a multicore system, when these two processes are scheduled on two different cores, the communication between these two processes generates an inter-processor interrupt (IPI). Whilst one process is waiting for the other to complete, the core with the waiting process goes into a lower power C-state. This sort of frequent C-state transition impacts power consumption in two ways:

- The energy requirements to enter and exit deep C-states are nontrivial. When the C0 (active) duration is very small, the latency to transition in and out of the C-states in comparison is appreciable and may result in net energy loss.
- The hardware policy may demote the C-state to a lower state based on heuristics. Even if the frequent C-state transition behaviour occurs for only 2–3 ms in a 15.6 ms window, hardware policies may either demote the core C-state or reopen the package-level cache, and this will impact power for the remaining $\sim 12\text{--}13$ ms of idle period.

To reduce C-state transitions in applications and services, it is recommended not to split a task between processes and threads unless parallel execution can occur. If it is necessary that a task be split between processes, then schedule the work so that the number of C-state transitions can be reduced. Also, applications and services should coalesce activity whenever possible to increase idle period residency.

3.2.4.2 OS Timer Resolution

The default system-wide timer resolution in Windows is 15.6 ms, which means that every 15.6 ms the OS receives a clock interrupt from the system timer hardware. When the clock interrupt fires, Windows performs two main actions: It updates the timer tick count if a full tick has elapsed, and it checks whether a scheduled timer object has expired. A timer tick is an abstract notion of elapsed time that Windows uses to consistently track the time of day and thread quantum times. By default, the clock interrupt and timer tick are the same, but Windows or an application can change the clock interrupt period.

Many applications call `time Begin Period` with a value of 1 to increase the timer resolution to the maximum of 1 ms to support graphical animations, audio playback or video playback. This not only increases the timer resolution for the application to 1 ms, but also affects the global system timer resolution, because Windows uses at least the highest resolution (i.e. the lowest interval) that any application requests. Therefore, if only one application requests a timer resolution of 1 ms, the system timer sets the interval (also called the system timer tick) to at least 1 ms. Generally, setting the timer to a value less than 10 ms can negatively affect battery life. Modern processors and chipsets, particularly in portable platforms, use the idle time between system timer intervals to reduce system power consumption. Various processor and chipset components are placed into low-power idle states between timer intervals. However, these low-power idle states are often ineffective at lowering system power consumption when the system timer interval is less than 10 ms.

If the system timer interval is decreased to less than 10 ms – including when an application calls `time Begin Period` with a resolution of 1 ms – the low-power idle states are ineffective at reducing system power consumption and system battery life suffers. System battery life might be reduced as much as 25%, depending on the hardware platform. The reason is that transitions to and from low-power states incur an energy cost. Therefore, entering and exiting low-power states without spending a minimum amount of time in the low-power states may be more costly than if the system simply remained in the high-power state.

Windows 7 comes with a command line utility called `PowerCfg` (Microsoft Corporation, 2010b). Using the `/energy` option, `PowerCfg` can be used to determine whether an application has increased the platform timer resolution. Run the `PowerCfg` utility when the application is running and examine the resulting energy report to see if the application changed the platform timer resolution. `PowerCfg` will also show the entire call stack for the request. The report lists all of the instances of increased platform timer resolution and indicates whether the process hosting the application increased the timer resolution.

It is important that application developers understand the power impact of using higher resolution timers and set them to the lowest resolution that meets the performance requirements of the application for the specific platform. Applications and services should use the lowest timer resolution possible that meets the performance requirements of the application. If the application requires a high-resolution periodic timer, increase the timer resolution only when the application is active and then return the timer to its default state after the application exits. Software applications should minimize the use of APIs that shorten the timer period. In Windows this includes, but is not limited to, the `time Begin Period Multimedia Timer` API and `NtSetTimer Resolution` low-level API. Instead, applications should use the newer timer-coalescing API supported by the OS such as `SetWaitableTimerEx` (API) and `KeSetCoalescableTimer` (a device driver interface, or DDI) in Windows 7. Some other recommendations are as follows:

- If your application must use a high-resolution periodic timer, enable the periodic timer only whilst the required functionality is active. For example, if the high-resolution periodic timer is required for animation, disable the periodic timer when the animation is complete.
- If your application must use a high-resolution periodic timer, consider disabling use of the periodic timer and associated functionality when a power-saving power plan is active or when the system is running on battery power.

3.2.4.3 Background Activity

Frequent periodic background activity increases overall system power consumption. It impacts both the processor and chipset power. Long-running, infrequent events also prevent the system from idling to sleep.

Background activity on the macro scale (e.g. minutes or hours) such as disk defragmentation, antivirus scans and the like are also important for power. Windows 7 has introduced a unified background process manager (UBPM) to minimize the power impact from background activities. The UBPM drives scheduling of services and scheduled tasks and is transparent to users, IT pros and existing APIs. It enables trigger-start services. For instance, many background services are configured to auto start and wait for rare events. UBPM enables the trigger-start services based on environmental changes. Some of the environmental change activities include device arrival or removal, IP address change, domain join and so on. An example of a trigger-start service is starting Bluetooth services only when the Bluetooth radio is currently attached.

Some of the other Windows 7 improvements to minimize frequent idle activity are as follows:

- Elimination of transmission control protocol (TCP) distributed program call (DPC) timer on every system timer interrupt;
- Reduction in frequency of USB driver maintenance timers;
- Intelligent timer tick distribution;
- Timer coalescing.

Table 3.4 summarizes the methodologies for energy conservation presented in this section and highlights their potential benefits. Note that the energy saved and other benefits will vary significantly and depend heavily on the specific application and workloads.

3.3 EVALUATING AND MEASURING SOFTWARE IMPACT TO PLATFORM POWER

To improve the energy efficiency of application software and to optimize power in mobile platforms, it is important to understand the impact on power consumption of the different hardware functional units. Various tools are available to provide a high-level estimate of the power consumed by a mobile platform. Whilst these tools may be sufficient to understand the high-level power consumption of a particular platform, they do not provide fine-grained details of specific components. The more accurate and invasive method to measure power is to use data acquisition (DAQ) tools where specific hardware components are instrumented and more granular power measurement can be logged. DAQ tools require an instrumented platform for actual power measurement.

3.3.1 Fluke NetDAQ® (Networked Data Acquisition Unit)

The Fluke NetDAQ is one of a class of DAQ tools that can be used to measure platform power consumption whilst running different applications on the system. A NetDAQ has many channels (10–100) that may be programmed individually to perform virtually any signal conditioning function the user requires. NetDAQ logger software for

Windows is a configuration and data management programme for the NetDAQ Systems. The details of a NetDAQ setup are as follows:

- The target PC has a special motherboard with built-in sense resistors. For each target component (i.e. CPU), all sense resistors are wired and soldered at both ends before being connected to a module attached to the NetDAQ unit.
- The NetDAQ has modules that are attached with individual wires to the target PC and measures the current and voltage drop across the sense resistors. The NetDAQ is connected to the host PC via a cross-over network cable.
- The host PC can be any IA32 system with NetDAQ logger software installed on it. The logger collects the measured current and voltages and calculates the average power (W). The sampling interval can be changed on the basis of individual analysis.

3.3.2 Software Tools

Tools are an important asset for tuning software for performance and energy efficiency. The better your tools, the easier it will be for your developers to create powerful applications on top of any platform. The power tools discussed here help diagnose energy problems with the applications and system and provide a starting point towards creating energy-efficient apps and platforms.

3.3.2.1 Windows 7 PowerCfg

As mentioned in this chapter, PowerCfg is a command line tool that lets users control their system's power management settings. Users can use PowerCfg to view and modify power settings, detect common energy efficiency problems when the system is idle. Used with the /energy option, PowerCfg provides an HTML report or a checklist of problems based on a snapshot of the system's energy consumption over a 60 s period. Users can detect power policies, sleep states that are supported in the system, USB devices that are not suspending and applications that have increased the platform timer resolution. Using PowerCfg, typical diagnostics are generated and the following issues are reported:

- Processor utilization: overall utilization and per-process utilization;
- Power policy settings: idle timeouts, processor power management (PPM) configuration and wireless power save mode;
- Platform timer resolution requests;
- Outstanding power requests: display, sleep and away;
- Platform capabilities: sleep state availability, display dimming capability, firmware validation problems and Peripheral Component Interconnect Express (PCI Express) active
- state power management (ASPM) status;

- Battery capacity: battery static data and last full-charge capacity or design capacity;
- USB device selective suspend issues.

To see the available options for PowerCfg, on a Windows Vista or Windows 7 system, run the command line tool as administrator and type 'powercfg /?'. For further details on PowerCfg, refer to Microsoft Corporation (2010b).

3.3.2.2 PowerInformer

PowerInformer (Intel Corporation, 2008) is a tool developed by Intel to provide basic power-relevant statistics to a developer. Developers can use these statistics to optimize their application against battery life constraints whilst meeting performance requirements.

PowerInformer features include the following (see Figure 3.7):

- Battery and power status indicators;
- Processor (physical, core and logical) detection;
- Percentage time of C1, C2 and C3 states of the system and of all logical processors;
- Average residency of C1, C2 and C3 states of the system and of all logical processors;
- Percentage time of the system and of all logical processors;
- Percentage time of the P stages of the system;
- System calls per second;
- Interrupt rate of the system;
- Disk I/O operation (access, read and write) rates;
- File I/O operation (control, read and write) rates.

3.3.2.3 Energy Checker

The Intel® Energy Checker SDK was developed to help software developers measure software energy efficiency and write energy-aware software. Energy-aware software should be capable of measuring and reporting dynamically its energy efficiency metrics. The SDK was designed to simplify this task, so the software vendor can focus on developing the heuristics required to save energy. The Energy Checker API provides the functions required for exporting and importing counters from an application. A counter stores the number of times a particular event or process has occurred, much like the way an odometer records the distance a car has travelled. Other applications can read these counters and take actions based on current counter values or trends derived from reading those counters over time. The core API consists of five functions to open, reopen, read, write and close a counter. In practice, using the API exposes metrics of 'useful work' done by an application through easy software instrumentation. For example, the amount of useful work done by a payroll application is different from the amount of useful work

performed by a video-serving application, a database application or a mail server application. All too often, activity is measured by how busy a server is whilst running an application rather than by how much work that application completes. The Energy Checker SDK provides a way for the software developer to determine what measures of ‘useful work’ are important for that application and expose those metrics through a simple API. Figure 3.8 shows a sample screenshot of the Energy Checker. More information on the SDK can be found at Intel Corporation (2010).

3.4 SUMMARY

Energy efficiency will be crucial for the computing industry in the future both to increase battery life for mobile platforms and to reduce energy expenses for desktops, server platforms and data centres. The demand for higher performance from mobile devices and new usage models will also continue to grow. The world is moving towards green technologies, and consumer demand for longer battery life in mobile devices will continue to increase.

As discussed in this chapter, software behaviour can have a significant effect on platform power consumption and battery life. Modern processors and platforms have many energy saving features, particularly for gaining performance improvements and the ability to enter low-power states when idle. But to benefit from these features, software should harness and work in harmony with these features. Software developers should do the following:

- Take advantage of performance features by emphasizing computational efficiency.
- Be frugal with data movement to improve data efficiency.
- Implement intelligent application behaviour by exploiting context awareness.
- Seriously consider the impact of software at idle to improve idle efficiency.

There are many free tools from Intel and others to help you get started. Even small improvements when amplified across millions of systems can make a dramatic difference.

3.5 EXERCISE

- How do processor C-states save energy?
- What are some programming methods used to achieve computational efficiency?
- Explain how data buffering can save energy.
- Give a few examples to illustrate how context awareness leads to ‘smarter’ devices.

- How much more battery life can be obtained from a 60 W h battery if we reduce the idle floor of the system from 8 to 6 W?

3.6 REFERENCES

- Harnessing Green IT: Principles and Practices, San Murugesan, G. R. Ganadharan, Wiley & IEEE.
- Green IT, Deepak Shikarpur, Vishwkarma Publications, 2014
- Green Communications: Principles, Concepts and Practice-Samdanis et al, J. Wiley
- Green IT for Sustainable Business Practice: An ISEB Foundation Guide, Mark G. O'Neill, The Chartered Institute for IT, 2010

munotes.in

SUSTAINABLE SOFTWARE DEVELOPMENT

Unit Structure

- 4.0 Objective
- 4.1 Introduction
- 4.2 Current Practices
- 4.3 Sustainable Software
- 4.4 Software Sustainability Attributes
- 4.5 Software Sustainability Metrics
- 4.6 Sustainable Software Methodology
- 4.7 Defining Actions
- 4.8 Summary
- 4.9 Exercise
- 4.10 References

4.0 OBJECTIVE

- Discusses how software (a virtual asset) can negatively impact the environment.
- Illustrates the importance of considering software's social and economic aspects.
- Examines current approaches to address environmental issues related to software.
- Introduces the notion of sustainable software development.
- Presents a methodology to measure the sustainability performance of software projects.
- Illustrates the application of the methodology with a real-world case study

4.1 INTRODUCTION

In recent years, green IT has gained attention across the IT industry, especially in organizations that maintain energy-intensive resources such as data centres. Although most of the focus has been on minimizing hardware's energy consumption and improving resource utilization, very little attention is paid to the environmental impact of software. Software is commonly regarded as neutral due to its virtual nature, but its indirect

impact is in fact real. The most obvious impact is the higher energy consumption caused by inefficient software, but other issues, such as increasing e-waste caused by existing computers becoming obsolete due to software demands, are greatly ignored by software professionals.

Software also creates social issues besides its environment impact. For example, lack of accessibility support (i.e. support for disabled persons) effectively excludes people from the workforce. In addition, ill-designed and hard-to-use software demands more training and user support which might increase software's carbon footprint, particularly if it involves printing and distributing hard copy of training materials and additional travel. E-waste also creates a social issue, since most of the dumped and highly toxic components are shipped illegally to developing countries, where they are handled by untrained workers without protective gears, effectively risking their lives and putting additional burden on the community.

The terms green and sustainability are often used interchangeably, but there is a subtle difference: Whilst green normally refers to environmental aspects, sustainability encompasses environmental, economic and social dimensions (Figure 4.1).

These three dimensions are integrated and interrelated:

- Economic growth without addressing its environmental impact diminishes quality of life.
- Environmental protection without considering local culture and traditions creates social injustice.
- Social disruptions like riots and war jeopardize the environment and the economy.
- A healthy economy enables a better environment and social welfare.

In advocating sustainable software, we include all three sustainability factors – environment, society and the economy – and align them to the project's business goals. It is imperative to adopt a holistic approach, considering not only the environmental aspects but also social and economic issues.




Figure 4.1 Three Dimensions

4.2 CURRENT PRACTICES

Unlike other engineering disciplines such as environmental and civil engineering, currently there are no well-established standards or practices to effectively manage sustainability in software projects. In the context of software, most of the effort on software sustainability focuses on the maintenance of legacy systems, and they do not address software sustainability in terms of environmental, social and economic aspects.

Since accurate tracking was complex, managers were then instructed to estimate resources used by the team members, such as paper used, travel and meetings. The estimated impact was then converted into cash amounts and reinvested in environmental projects. Whilst the practice was criticized for not being consistently quantifiable, it helped to raise awareness about the environmental impact in the context of software projects

As Figure 4.2 shows, there are several methodologies and standards that cover the life cycle of computer hardware (design, manufacturing, usage and disposal), where the main concerns are lower energy and water consumption in data centres, lower carbon emissions, use of less toxic materials and environmentally friendly disposal of unused hardware. For example, Energy Star and Restriction of Hazardous Substances (RoHS) are now industry standards, and IBM Green Sigma is a generic framework that applies Lean Six Sigma strategies to reduce water and energy usage, and can also be used on various business operations, including IT operations.

However, these methodologies and standards do not address environmental concerns related to software, such as system longevity (minimize e-waste), algorithmic efficiency (minimize energy consumption), project's carbon footprint (minimize carbon emissions) and support for disabled persons (address societal and economic issues, such as digital inclusion and disability inclusion)




Figure 4.2 Methodologies and Standards

4.3 SUSTAINABLE SOFTWARE

Sustainable software is software that is developed considering its positive and negative impact on the environment, society and the economy, whilst still delivering the value for which the software is intended.

Having discussed the negative impacts associated with software, it is not realistic to simply expect companies to remove product features or ask programmers to simply take care of the problem without a systematic approach. Just like quality and security, sustainability should be integrated as an orthogonal aspect of the software development process.

For instance, ‘quality’ is built into software or other products through a series of actions adopted across the development cycle, supported by metrics, that eventually increases the overall quality of the product. Sustainability should be integrated with a similar systematic approach, with incremental steps driven by trade-off decisions, supported by a pragmatic analysis of its properties, across the entire development process.

It also embraces the continuous improvement ethos like IBM Green Sigma and Six Sigma. This methodology is integrated at the end of the existing software life cycle when the product is released (see Figure 4.3), and it includes three steps:

1. Metrics are collected at the end of development cycle.
2. Metrics are analysed based on their environmental, social and economic benefits.
3. Corrective or refined actions are formulated for the next development cycle and release of the product.

The goal of this approach is not to label whether a system is ‘sustainable’ or ‘not sustainable.’ Instead, it aims to measure and improve the overall sustainability performance across several product releases. This practice allows us to create an environment of ‘pro sustainability’ across the organization, by adopting small but meaningful improvements where processes are constantly measured and evaluated.

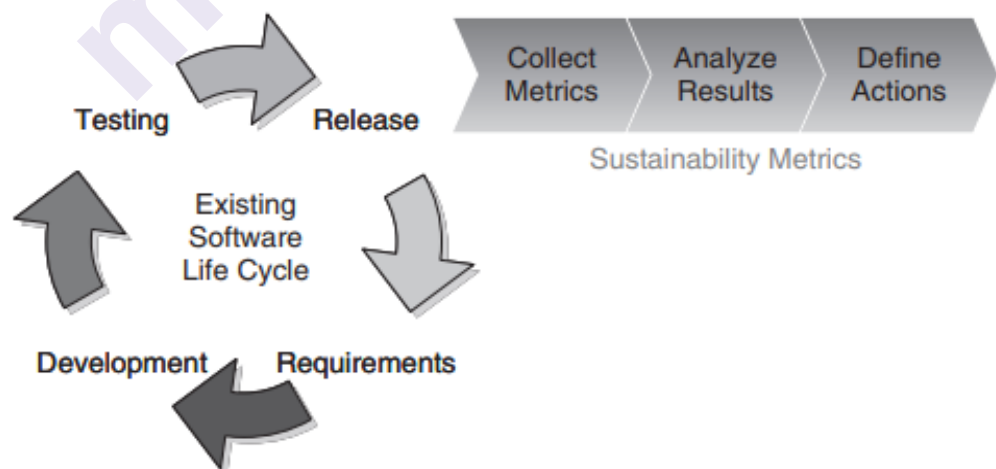


Figure 4.3 Life Cycle

To evaluate the sustainability performance of a given software system, we must first outline what attributes of the system we want to improve, and examine their impact and benefits. In other words, we need to define what a sustainable software system means in the first place.

As previously discussed, sustainability is an orthogonal property-like quality, and as such we will borrow a software-engineering concept called quality attributes to define sustainability. Quality attributes (also known as nonfunctional requirements) are overall system attributes that impact the system's operation as opposed to its functions or features, and are normally used to analyse architectural trade-offs. They also provide a foundation for defining metrics and understanding their benefits.

The following attributes are most relevant to sustainability performance:

- **Development-related attributes:** These are the attributes that impact the development process, and key among them are the following:
- **Modifiability:** The ability to make changes quickly and cost-effectively.
- **Reusability:** The degree to which system components can be reused in other systems.
- **Portability:** The ability of the system to run under different computing environments.
- **Supportability:** The system's ability to be easily configured and maintained after deployment.
- **Usage-related attributes:** These are the attributes that impact the usage at runtime.

They are as follows:

- **Performance:** The time the system takes to respond to user requests.
- **Dependability:** The ability of a system to function as expected at any given time.
- **Usability:** Features that make a system easy to use.
- **Accessibility:** The system's ability to allow users to access the system regardless of the user's location, or the type of computer or access device used.
- **Process-related attributes:** These are the attributes that impact project management.

They are as follows:

- **Predictability:** The ability to accurately estimate the effort and cost involved in developing software upfront.
- **Efficiency:** The effort towards deliverables that add direct value to the customer, as opposed to the effort towards deliverables necessary to run the project.
- **Project carbon footprint:** The net carbon emissions arising from the development of the software.

Quality attributes were originally created to minimize the cost or effort required during system development, maintenance and use, whilst maximizing the system's value.

However, these attributes can also be used to analyse the (indirect) environmental and social impacts of a given system. For example, good performance improves the user's productivity (thus minimizing effort and cost), but it also indirectly minimizes e-waste because the system can be deployed on older computers.

Though many may not recognize it, these software attributes bring environmental, social as well as economic benefits.

- **Modifiability and reusability:** Less effort in introducing changes minimizes the resources used by the development team as well as the amount of waste, and hence reduces the carbon footprint.
- **Portability:** Software that can run on different environments increases flexibility and minimizes e-waste because it can run on existing hardware, for example by installing a less demanding OS like Linux to run the application. Portability also reduces the costs of adoption, and therefore the application is accessible to more users.
- **Supportability:** Less effort in configuring and operating the system minimizes the use of physical resources such as books and training required to support the system.
- **Performance:** Increased performance with even less resources minimizes e-waste by allowing the system to be deployed on older computers.
- **Dependability:** A system that functions reliably and as expected minimizes the need for additional resources because users do not have to deal with recurring unplanned tasks or emergency maintenance.
- **Usability:** By reducing the usability barrier, the system is available to a broader cross section of users. Also, better usability of a system reduces the amount of training required, thus minimizing the use of books, travelling and other resources.
- **Accessibility:** It allows the system to be used by minorities, the elderly, people with disabilities, non-English-speaking communities and the illiterate population.

- **Predictability:** A team that develops software in a predictable manner has to deal with fewer emergencies and unplanned activities, therefore minimizing its usage resources (office resources, computer equipment, electricity, etc.).
- **Efficiency:** A more efficient software team optimizes the use of physical resources.
- **Project carbon footprint:** Decreasing the software team's carbon footprint directly reduces the amount of greenhouse gas emitted.

4.5 SOFTWARE SUSTAINABILITY METRICS

To assess the greenness and improvements in environmental performance, we need to quantify each one of these attributes. These metrics are commonly used in software engineering, and most of these measurements can be done or collected automatically (which is discussed in this chapter).

4.5.1 Modifiability and Reusability

Of all the attributes, modifiability and reusability are perhaps the most common in software projects, because they directly affect the programmer's productivity. For instance, in object-oriented systems it is a common practice to measure and minimize the degree of dependency among classes, as highly interdependent classes are difficult and risky to change.

To find the distance from main sequence, a prime measure of modifiability, we first need calculate instability and abstractness – both are measures of reusability. Instability measures the potential impact of changes in a given package. A package is considered stable when it has no dependencies on outside classes. However, it is undesirable to have a system that is completely stable, since that would make it unchangeable. In view of this, we also calculate abstractness, which measures how much a package can withstand change through the use of abstract classes. We then calculate the distance from main sequence, which represents the balance between instability and abstractness, by measuring how far a package is from the idealized balance.

$$\text{Instability} = \text{Efferent Couplings} / (\text{Afferent Couplings} + \text{Efferent Couplings})$$

where Afferent Couplings is the number of classes outside a package that depend upon classes within the package, and Efferent Couplings is the number of classes inside a package that depend upon classes outside the package.

Instability ranges from 0 to 1, where 0 means the package is maximally stable and 1 means the package is maximally unstable.

Abstractness is the number of abstract classes in a given package divided by the number of concrete classes in a given package.

It ranges from 0 to 1, where 0 means the package is completely concrete and 1 means it is completely abstract.

Distance from Main Sequence = Abstractness + Instability – 1

It ranges from 0 to 1, where 0 means the package has the ideal balance and 1 means the package requires redesign and refactoring.

These metrics help us to identify whether the package needs to be optimized in the next release.

4.5.2 Portability

An objective of sustainable software development is to minimize e-waste by limiting the implementation of software features that would render reasonably recent hardware obsolete. Since this is a subjective goal, we try to quantify it by finding the estimated system lifetime, which is an estimation of the year when the hardware required to run the system reached the market. The intent is to make a reasonable judgment that the new system will not force the target desktops to be unnecessarily upgraded.

4.5.3 Supportability

Another objective is to minimize the amount of post-deployment support required by the system, which would in turn reduce the amount of resources needed. One could use metrics already gathered by the system support team, such as number of calls and average time to resolve tickets. For new systems one can use a metric devised from the usability study, called the support rate. For further information on the usability study, refer to Section 4.6.1.

Support Rate = User's Questions That Required Assistance/Number of Minutes System Was Used in a Usability Study Session.

4.5.4 Performance

A measure of performance is relative to the objective of the system. Since our goal is to improve user productivity, we use the relative response time metric captured during the usability study.

Relative Response Time = Number of Tasks with an Unacceptable Response Time/Total Number of Tasks Tested.

4.5.5 Dependability

The goal is essentially to minimize defects, which are costly from both business and environmental standpoints, because it causes more resources to be wasted. We use three common quality assurance metrics to evaluate dependability: Defects Density (the amount of defects compared to the system's size), Testing Efficiency and Testing Effectiveness.

As the use of Defect Density alone could be misleading if the testing process is faulty, we also use the Testing Efficiency and Testing Effectiveness metrics.

- Defect Density = Known Defects (Found but Not Removed)/Lines of Code (LOC);
- Testing Efficiency = Defects Found/Days of Testing;
- Testing Effectiveness = Defects Found and Removed/Defects Found.

4.5.6 Usability

- Quality attributes are used not only to assess quality or (in our case) sustainability, but also by usability experts. For improving software ease of use, Jakob Nielsen (n.d.), a leading usability expert, recommends the following attributes:
- **Learnability:** How easy is it for users to accomplish basic tasks the first time they encounter the design?
- **Efficiency:** Once users have learned the design, how quickly can they perform tasks?
- **Memorability:** When users return to a design after a period of not using it, how easily can they re-establish proficiency?
- **Errors:** How many errors do users make, how severe are these errors and how easily can they recover from the errors?
- **Satisfaction:** How pleasant is it to use the design? A usability study is required to measure these attributes, which are defined as follows:
- **Learnability:** Minutes to Accomplish First Critical Task without Assistance/Minutes of Usability Test/Minutes System Was Used by User;
- **Efficiency:** Tasks Completed/Total Tasks Tested;
- **Error Rate:** Tasks Which Were Completed but Deviated from Normal Course of Action/Total Tasks Tested;
- **Satisfaction:** Score Based on Subjective User's Feedback.

4.5.7 Accessibility

Accessibility can be measured based on the following requirements, adapted from Ball State University's Web Accessibility Initiative (Ball State University, 2002):

- Support for motor-impaired users;
- Support for visually impaired users;
- Support for blind users;
- Support for users with language and cognitive disabilities;
- Support for illiterate users;
- Internationalization and localization support.

Each of these requirements is measured on a 4-point scale: 0 (non-existent), 1 (not adequate), 2 (acceptable) and 3 (adequate).

Most of these requirements are provided (or not provided) by the user interface (UI) platform used to develop the system, whilst some of them require the implementation of new measurements. Also, some of these items may not be relevant depending on the demographics of the target users: For instance, it does not make sense to provide support for illiterate users in a system that targets technical users who are professionally required to be literate.

4.5.8 Predictability

The goal is to measure how effectively the development team is able to estimate effort and cost. In traditional projects, we can obtain this metric by simply counting the number of days a project was delayed. However, in agile-based projects a project is never late, since agile methodologies have fixed time but variable scope. In such projects we count how many iterations were overestimated or underestimated based on an acceptable threshold

(say, $\pm 20\%$) by simply comparing the number of points allocated at the beginning of the iteration with the number of points actually used at the end.

- Estimation Quality Rate (for Traditional – Waterfall – Software Development Life Cycle Projects) = $\text{Planned Days} / \text{Actual Days}$;
- Estimation Quality Rate (for Agile Software Development Projects) = $\text{Number of Underestimated or Overestimated Iterations} / \text{Number of Total Iterations}$.

4.5.9 Efficiency

Project efficiency is measured by the effort towards deliverables that add direct value to the customer (e.g. coding and manuals) as opposed to project-related effort (e.g. infrastructure maintenance and project management). The efficiency can be represented in hours or points, depending whether your project is based on a waterfall or agile methodology.

4.5.10 Project's Carbon Footprint

To measure the project's carbon footprint, rather than quantifying the amount of all the resources used during a project, which may be hard to do, one can track the activities that can potentially cause a negative impact on the environment, such as the following:

- Work-from-home days;
- Long-haul round trips;
- Number of conference calls;
- Number of offsite meetings.

Metrics should be selected based on the nature of the project. For example, a project involving a globally dispersed team is likely to hold many conference calls, but that fact alone does not necessarily mean that the team's footprint is low. In addition, some metrics might be considered beneficial from an environmental standpoint but harmful to other areas of the project. For example, holding less physical meetings might have a positive environmental impact, but it can negatively impact the team's cohesiveness and motivation, which could ultimately jeopardize the project's sustainability.

Note that the goal of these metrics is not to answer questions such as 'How much CO₂ does my team emit?' (which is hard, if not impossible, to track), but instead to spark discussions like 'Why not visit clients by train instead of car?' The key is to select sensible and simple metrics, and stick with them across several product releases

4.6 SUSTAINABLE SOFTWARE METHODOLOGY

4.6.1 Collecting Metrics

Having defined the 'why' (benefits) and the 'what' (attributes and metrics), in this section we outline 'how' to obtain the metrics (see Table 4.3). Some of the metrics can be collected automatically, whilst others are likely to be already available in your project.

4.6.2 Code Metrics Tools

Metrics such as instability, abstractness, distance from main sequence and LOC are standard in software engineering, and most development environments have several tools to obtain such metrics automatically. The following are some code metrics tools for specific languages or development environments:

❖ Java:

- IBM Rational Software (commercial): <http://www-01.ibm.com/software/rational/>
- Eclipse Metrics (open source): <http://metrics.sourceforge.net/>

❖ NET:

- NDepend (commercial): <http://www.ndepend.com/>
- C++:
- CppDepend (commercial): <http://www.cppdepend.com/>

❖ Adobe Flex:

- FlexMetrics (open source):
<http://opensource.adobe.com/wiki/display/flexpmd/Flex>

❖ Metrics

- It Depends (open source): <http://code.google.com/p/it-depends/>
- How to get the LOC count:
<http://dougmcune.com/blog/2007/05/10/analyze-youractionsript-code-with-this-apollo-app/>

❖ Multilanguage:

- Sonar (open source): <http://www.sonarsource.org/>

4.6.3 Simplified Usability Study

Measurement of support rate, relative response time and the four metrics of the usability attribute require a usability study. Usability studies can be complex, and in large projects they are normally handled by a separate team of professionals (user experience designers or usability engineers). For small to mid-sized projects, however, a simplified usability study method is described here which can be completed in a few hours.

One or more subjects are required for the usability test. The subject's profile will depend on the target audience of the application being tested. For instance, an outsider can be recruited if the system is intended to be used by people with no previous knowledge (e.g. a Web application). However, if previous knowledge is required to use the system, then a person who is familiar with the subject matter should be recruited (the subject should not have previous knowledge of the system being tested). Recruit someone who can be classified as a new user, so the learning ramp-up required to use the system can be measured.

Then, a list of small tasks to be performed by the subject(s) must be created. The tasks in the list must follow an order that can be performed during a single test session, and they should cover the most important use cases a user is expected to accomplish the first time the subject encounters the system. The tasks must have sufficient information to allow the subject to accomplish specific tasks without assistance: for example, a task should not request something vague like 'Navigate around the map', but instead specify 'Find river X in the map'. The tests should be designed so that the entire task can be completed within 10–15 minutes.

During the actual test, the subject should be observed and the time taken to complete each task should be measured. Help should be provided only when the subject specifically requests assistance, or when he or she is clearly stuck on a task. The results of each test should be recorded: 'The task was completed without assistance', 'The task was completed with assistance' or 'The task was completed but deviated from the expected path' (i.e. the subject accomplished the task by using a different path than was normally expected).

At the end of the session, feedback about the subject's overall satisfaction about usability should be asked. Also, the system response time of each task should be recorded: good, slow or unacceptable. There are

quantifiable ways to track response time (e.g. a tool like Firebug could be used to measure page load times on Web applications), but simple subjective feedback is sufficient for the purposes of this study, based on the subject's own observation or on complaints during the test. A form like in Table 4.4 can be used to record test session data.

Finally, compile the following data that will allow you to calculate support rate, relative response time and the four usability metrics:

- Number of minutes to accomplish the first critical task without assistance;
- Total number of minutes the test session took;
- Number of tasks completed;
- Number of tasks completed but deviated from normal course of action;
- Number of tasks that required assistance;
- Total number of tasks.

4.6.4 Platform Analysis

To estimate system lifetime and the Accessibility Requirement Score, identify the minimum hardware and software requirements of the platform on which the system will run or be deployed, and the platform's accessibility compliance.

To estimate the system lifetime, we first start by listing all the major underlying software components used by the system, both client and server, such as Java, Flash Player, a database and others. We then find the software and hardware requirements based on the documentation of each component, and based on information available online (on the manufacturer's Web site or Wikipedia) we gather the year the component or required hardware was released. For example, Java 6 was released somewhat recently, but the documentation mentions that it supports Internet Explorer 6 SP1, and through Wikipedia we found that IE6 SP1 was released in 2002. The assumption here is that if IE6 SP1 was released in 2002, it is likely that a computer manufactured in 2002 will run IE6 SP1. We use the same approach for hardware requirements: For example, a software component requires an 'Intel or AMD (Advanced Micro Devices) processor running at 1 GHz', and we found in a news article that the first 1 GHz Intel chips came out in 2000.

The estimated system lifetime is the most recent year that we can find for this analysis, and we normally add two more years, assuming that this is the amount of time that it takes for a product to become broadly available. The idea is that if the year found is too recent (e.g. the system lifetime is 2010) we should try to mitigate that because it likely will force users to upgrade their existing hardware. Normally we do not consider memory and hard-disk requirements because they are not constraints per se: Normally they are standard and interchangeable (they can be reused on old and new computers). Client and server requirements should be analysed separately, and the client requirements are normally the most relevant

because they impact the largest number of physical computers, and servers tend to be more powerful and commonly reused across the organization.

Obtaining the accessibility information is usually fairly simple, search online for '[your UI platform name and version] Section 508'. Section 508 is part of the US Rehabilitation Act, and it became a standard for accessible products. You will search for the platform on which your UI relies: For example, if your application uses Java 6 on the server side and Adobe Flash Player 10 on the client side, then you would search for 'Flash Player 10 Section 508'.

4.6.5 Existing Project Statistics

Most software projects already track metrics like number of defects or effort in some capacity, either manually (using spreadsheets or documents) or automatically (using management tools, such as bug-tracking systems and support ticket systems). These metrics are normally public within the organization, because they are useful to all team members, and you can collect them from the following sources:

- Project management;
- Testing and quality assurance team;
- Customer support team.

It is a good practice to involve these teams during the sustainability analysis, as this would create a sense of buy-in and involvement for people across the entire organization.

4.7 DEFINING ACTIONS

Based on the value of the attributes measured or estimated, one can identify the attributes that need to be improved in the next development–release cycle, and select the key goals for the next release cycle, based on the following:

- Business value: the degree to which the goal supports the system's objectives.
- Effort required for implementation: the feasibility or likelihood that the goal will be actually implemented.
- Sustainability value: the degree to which the goal would be to bring environmental, social and/or economic benefits.

This criteria list helps the selection of goals that can realistically be implemented. If we evaluate the goals solely by their sustainability benefits, it is unlikely that they will be implemented because there would not be organizational buy-in. A realistic goal is one that brings tangible sustainability benefits, but is also a win–win for decision makers, business managers concerned about profitability and project managers concerned about cost (effort).

Chosen goals are to be reported and evaluated by the project team during the next product release. Goals are translated into requirements that are treated by the development team as normal feature requests.

Goals must be specifically defined based on quantifiable and realistic metrics, such as ‘Improve learnability level by 30%’. Realistic does not necessarily mean small. Although most continuous improvement actions are small but continuous steps, there are also opportunities for breakthrough improvements.

4.8 SUMMARY

The goal of this methodology is to extract quantifiable and reliable metrics, so they can be consistently managed and improved. However, it is very difficult to connect the benefit implemented in the software with the actual physical benefit. That would be possible only if there is an extensive collection of case studies that would allow us to research the actual physical impact of the implemented actions. A collection of case studies would also be helpful in defining industry benchmarks for each one of the metrics.

Although substantial effort was made to make the metrics quantifiable, some of them are still based on subjective criteria. In addition, some of the metrics are still somewhat hard to gather manually. A tool to automatically gather some of these metrics would be extremely helpful to address both issues.

This methodology is intended to be a first step, and by no means does it try to cover all possible sustainability-related indicators. Our goal is to remove some of the guesswork in implementing sustainability initiatives by offering a practical starting point. Considering that the biggest barrier for adoption is the lack of awareness regarding the physical impact of software, it is recommended to start with a subset of the metrics (specifically, the ones that can be obtained automatically). The results of such a study can then be used as a venue within the organization for discussions about the impact of software production on the environment and society, in order to get the buy-in from the different stakeholders.

4.9 EXERCISE

Review Questions

1. How can software impact the environment?
2. Distinguish between ‘green’ and ‘sustainable’.
3. What is sustainable software?
4. What are the quality attributes of software?
5. What are the attributes relevant to assessing sustainability performance?

4.10 REFERENCES

- Harnessing Green IT: Principles and Practices, San Murugesan, G. R. Ganadharan, Wiley & IEEE.
- Green IT, Deepak Shikarpur, Vishwkarma Publications, 2014
- Green Communications: Principles, Concepts and Practice-Samdanis et al, J. Wiley
- Green IT for Sustainable Business Practice: An ISEB Foundation Guide, Mark G. O'Neill, The Chartered Institute for IT, 2010

munotes.in

GREEN DATA CENTRES

Unit Structure

- 5.0 Objectives
- 5.1 Introduction
- 5.2 Data Centers and Associated Energy Challenges
- 5.3 Data Centre IT Infrastructure
 - 5.2.1 Servers
 - 5.2.2 Networking
 - 5.2.3 Storage
 - 5.2.4 IT platform innovation
- 5.4 Data Centre Facility Infrastructure: Implications for Energy Efficiency
 - 5.3.1 Power System
 - 5.3.2 Cooling
 - 5.3.3 Facilities infrastructure management
- 5.5 IT Infrastructure Management
 - 5.4.1 Server power
 - 5.4.2 Consolidation
 - 5.4.3 Virtualization
- 5.6 Green Data Centre Metrics
 - 5.5.1 PUE and DCiE
 - 5.5.2 Power Vs Energy consumption
- 5.7 Summary
- 5.8 List of References
- 5.9 Unit End Exercises

5.0 OBJECTIVES

After going through this unit, you will be able to:

- Understand the concept of Green IT and its implementation in developing the sustainable information and computing system
- Implement and developing the Green IT strategies in enterprise development
- Learn about the green network communication and data center operation

5.1 INTRODUCTION

Environmental challenges and the adoption of environmentally sound practices have become a new priority for businesses, governments, and communities as a whole. Information technology (IT) has significantly changed our work and lives throughout the years, enhancing our productivity, economics, and social well-being. IT now has a new role to play: assisting in the creation of a greener, more sustainable environment while also providing financial rewards. However, most people are unaware that IT is contributing to environmental issues. Computers and other IT infrastructure use a lot of electricity, and that quantity is growing every day, putting a strain on our electric grids and adding to greenhouse gas (GHG) emissions. Furthermore, both during creation and disposal, IT hardware has a negative impact on the environment.

While many individuals believe that information technology is a source of pollution, it may also be a source of relief. To put it another way, information technology is both a solution and a problem in terms of environmental sustainability. We can leverage the power of IT in new ways to address pressing environmental concerns and make our IT systems – and their use – more environmentally friendly. Green IT, often known as green computing, is the science and practice of efficiently and successfully designing, manufacturing, and operating computers, servers, displays, printers, storage devices, networking, and communications systems with zero or little environmental impact. Green IT also refers to the use of technology to support, assist, and leverage other environmental activities, as well as to raise green awareness. As a result, green IT refers to technology, software, tools, methods, and practices that help to improve and maintain environmental sustainability.

Green IT improves energy efficiency, reduces GHG emissions, uses less toxic materials, and encourages reuse and recycling, all of which assist the environment. As a result, green IT encompasses environmental sustainability, energy efficiency economics, and total cost of ownership, which includes disposal and recycling costs. Businesses and individuals are being driven to go green by increased awareness of the adverse impacts of GHG emissions, new strict environmental legislation, worries about electronic waste disposal techniques, and corporate image issues.

Green IT is both a financial and an environmental necessity. As many green activists would tell, it is also our social responsibility. More green taxes and regulations are on the way, resulting in a surge in demand for green IT products, solutions, and services. As a result, an increasing number of IT providers and users are developing and selling green IT products and services. As businesses and governments attempt to balance growth with environmental hazards, we will be obliged to 'green' our IT products, apps, services, and processes on a legal, ethical, and/or social level.

To promote green IT, we must first comprehend the following issues: What are the most significant environmental consequences of IT? What

are the most pressing environmental IT challenges that need to be addressed? What can we do to make our IT infrastructure, goods, services, operations, applications, and processes more eco-friendly? What are the regulations or norms that we must adhere to? What role can IT play in helping businesses and society as a whole improve our environmental sustainability?

This chapter highlights IT's environmental impact before defining green IT and presenting a holistic strategy to greening IT. It also presents a green IT strategy for businesses, highlighting how IT can help in a variety of ways to improve our environmental sustainability.

5.2 DATA CENTERS AND ASSOCIATED ENERGY CHALLENGES

The increased deployment of data center infrastructure has stemmed from the ever-increasing digitization of modern life. Data centers are complex ecosystems that interconnect elements of the information and communication technology (ICT), electrical, and mechanical fields of engineering, and they are the fastest growing contributor to the ICT sector's overall carbon footprint, according to the widely cited Global e-Sustainability Initiative Smart 2020 report (GeSI, 2008).

The Web has evolved into a worldwide platform that impacts the lives of billions around the world on a daily basis since its initial public acceptance in the early 1990s. The Web has revolutionized humanity's opportunity and outlook for now and the future through supporting commerce, education, news media, social participation, and scientific research.

The spread of digital technology, as well as humanity's natural urge to communicate, have aided the growth of the Internet. We've evolved from information consumers to digital content producers, thanks to old-school television, radio, and print media. Blogs, wikis, photo-sharing portals, social-networking sites, online auctions, and many other sites are used and contributed to by us.

According to data centre energy consumption trends compiled by Green Grid, data centre power usage in the United States is expected to account for 3.5 percent of total US power use in 2011. According to Koomey (2007), worldwide servers utilize about 120 TWh of electricity, and data centers will account for roughly 18% of the ICT sector's carbon footprint by 2020, according to GeSI (www.gesi.org/). In short, there has never been a bigger need for data-processing services from data centers. However, these technological, commercial, and societal advancements are not without their drawbacks. To mitigate the environmental impact of such rapid ICT and data centre growth, cost and provision must be made.

As a result, the impact of data centre facilities is a major concern for both businesses (in terms of meeting sustainability goals) and governments (who are scrutinizing data centre use and CO₂ emissions). The US Environmental Protection Agency's 2007 "Report to Congress on Server

and Data Centre Energy Efficiency" and the European Commission's 2008 "Code of Conduct on Data Centres Energy Efficiency" both attest to this growing focus.

incorporating energy-efficient design, high-efficiency power supply, high-efficiency cooling, and growing use of renewable energy sources into their operations (RESs). Such facilities strive for a balance that balances environmental sustainability with economic concerns. In reality, green data centres are viewed as a competitive advantage that helps to the enterprise's financial survival.

5.3 DATA CENTRE IT INFRASTRUCTURE

There are two layers of infrastructure in a modern green data centre: IT infrastructure and facilities infrastructure. The server technologies, networking systems, and storage provided within a conventional data centre are referred to as IT infrastructure.

The following major parts of a data center's IT architecture are discussed in this section:

- Server design and development to support effective data centre service delivery and a wide range of service functions.
- The importance of networking in a data centre
- The importance of storage and the various types of storage available.
- The evolution of data centre IT platforms as a result of system innovation.

5.3.1 Servers

IT servers come in a variety of shapes and sizes, and they provide a variety of services and activities, but their primary objective remains the same: to deliver a service as part of bipartite communication between a client and a server.

A server could be a software program running on the same physical machine as the client or via networking infrastructure from afar. Servers can either be software or hardware. A system created to perform a specific task: e-mail server, web server, database server and print server are all examples of servers. The IT hardware used to host a software server within a data centre varies in design, efficiency, and function. A server may be built to run a certain operating system (OS), and different OSs may be supported inside a data centre. Each server machine will be made up of physical hardware, an operating system, and a software service. This mix of options influences data centre architecture decisions and characteristics, and it must be evaluated early in the IT infrastructure development process.

A server computer built for use in a data centre will have significantly different characteristics than one designed for use at home or in the office. This is owing to the nature of its eventual application, in which enormous

numbers of servers are networked via router and switch technologies, with gateway computers coordinating massive quantities of client service requests. As a result, most data centre servers are headless, meaning they don't have a monitor or mouse attached. The machine-to-machine service interaction paradigm, which is common in data centres, necessitates human participation only when something goes wrong or a machine requires maintenance. Human engagement is necessary for administration, setup, and maintenance in this situation; hence a control terminal with a monitor, keyboard, and mouse interface is supplied. A comprehensive Network Operations Centre may be present in large-scale data centres.

The central processing unit (CPU), onboard memory, networking interface, and software onboard a data centre server are all optimized for the task at hand. A database server, for example, may have more onboard memory to allow queries to be handled in memory, reducing the time it takes to retrieve data stored on networked disc space. Data centre servers, like home computers, must be housed in an appropriate form factor to allow for efficient integration into the larger data centre infrastructure.

5.3.1.1 Rack mounted servers

Rack-mounted servers are available in two sizes: 19 in. and 23 in., with the smaller being the standard. The machine housing has predrilled ears that match precisely with the vertical posts of the rack cabinet, making it reasonably simple to fit a racked server into a standard-width cabinet. A rail-mounting system is also a basic part of a rack for heavy equipment or equipment that requires regular access to replace or maintain server hardware components. This makes accessing the area behind the faceplate as simple as opening a drawer on a piece of furniture.

5.3.1.2 Blade servers

Racked servers are self-contained devices with their own power and network connections. Blades, on the other hand, are kept in a blade system container that is mounted in a normal cabinet. All of the blades inside the system are powered and networked by the blade system. The blade cabinet system may additionally provide on board cooling and an uninterruptible power supply (UPS).

5.3.1.3 Containers

The self-contained data centre module was the next logical step in server configuration architecture for data centres. Container-based data centres, which are shipped within a standard-sized transport container similar to those used to send products overseas or by heavy-goods vehicles, give an off-the-shelf answer to data centre needs. Their energy and connectivity requirements are understood ahead of time because they are designed to be self-contained. From racks to containers, modular design allows for rapid scalability and easy replacement, with old or failed units switched out for replacements or higher-spec units swapped in as demand demands change. Modularization also simplifies power management because energy consumption scales with processing capacity.

5.3.2 Networking

A data center's gateway machine will be located at the facility's entrance. Its principal function is protocol translation in and out of the data centre, serving as a link between the data center's internal local area network (LAN) and the wide area network (WAN) outside the data centre - in most cases, the Internet Service Provider's network.

The software services hosted within the data centre must get millions of requests for Web page content (hits), e-mail, or database query results. Once within the data center's LAN, each message must be routed to a server that can respond to the request immediately. Multiple routers will be deployed to direct message traffic, routing messages to different sectors of the internal network, depending on the scope and network complexity of a data centre. Routing tables and message address standards are used by routers to control traffic within the data centre. A router will either forward a message to the next address on the path (a hop) or simply drop it if there is no next hop for the message. Because of this easy procedure, routers can be adjusted to handle the massive amounts of data passing via the data center's networking cables.

While a router works at the network layer, switches can be used to manage communications across multiple layers, albeit they are most commonly used at the data layer. A switch in a data centre network's principal role is to connect different network segments. The system designers will deploy a number of switches to optimize and control the message or data traffic between all network segments in the data centre, depending on how the center's network of servers has been setup. Servers, storage, and a gateway will all be included.

The data center's backbone exists to offer high-capacity communication between the data center's various network parts. It will typically have a much higher bandwidth capacity than the segments to which it is connected, allowing for greater traffic flow capacity between multiple locations. The backbone will frequently be required to handle the data center's whole two-way message capacity.

5.3.3 Storage

The storage of data is an important aspect of data centre design. There are several solutions available, each of which caters to the needs of other parts in the overall IT infrastructure decisions. The way the client machine – in our example, the data centre server – logically views the storage media is the primary differentiator in storage type. This will affect how the server manages the storage space as well as the access methods available for accessing the data.

5.3.4 IT platform innovation

Data centre design is more than just physical hardware computation. Software platforms and OS virtualization, like physical servers, networking, and storage design, have enabled systems to be modularized

into distinct servers (called services) at the software level. Because of the separation of concerns between hardware resources and virtualized software provision, this advancement allows infrastructure provision to be more efficient. The rest of this part takes a quick look at the evolution of IT platform infrastructures toward completely virtualized cloud-based systems.

5.3.4.1 Server Farm (Cluster)

Multiple, physically distinct machines are closely coupled to give the logical interface of a single system in cluster computing. Cluster computing, which is frequently connected with the parallelization of processing methods, necessitates the use of dedicated and highly specialized middleware to create the sophisticated message-passing architecture needed to manage the cluster's physical resources. Different requests would be split and scheduled across a huge number of individual server computers, each equipped to compute a result, using the parallel processing capacity. The outcomes of separate operations would then be combined to generate the final product. The main usage and research interest for cluster computing is to save time by parallelizing workloads.

5.3.4.2 Grid computing

As the scientific research community's demands grew, researchers looked into grid computing as a natural evolution from clusters. A new strategy was sought as data sets grew in size and the needs for processing capabilities to analyze data in a timely manner grew. Grid computing's major goal was to bring together diverse resources from various organizational areas to form virtual organizations. Grid computing, on the other hand, raises a number of difficulties, including security across many domains, the complexity of managing interaction state, and the practicalities of processing large data quantities across a network.

5.3.4.3 Service orientation

Web services are an example of service-oriented architecture (SOA), which advocates a separation of concerns between service implementation (software) and service hosting (server hardware). It's a network-based method of delivering both data and processing resources that separates service instantiation from a machine-readable service interface. This allows system designers to create network services that can be accessed and addressed using regular Internet protocols. Increases in Internet speed have overcome grid computing concerns, making service orientation more realistic in the long run.

5.3.4.4 Virtualization

Virtualization is another way to provision compute servers. A virtualized system allows a single hardware computer running a single OS to host numerous virtual machines, which may or may not be running the same OS, rather than a single hardware machine supporting a single OS that supports a single server application (Web server, mail server, etc.). As a

result, numerous virtual machines can run on a single host machine. Virtualization allows a data centre to scale service provision, make far more efficient use of hardware resources, save operating costs, and minimize energy consumption. As a result, it has proven to be a game-changing technology in data centre IT design infrastructure with larger reduction in power consumption.

5.3.4.5 Cloud computing

The combination of SOA and virtualization resulted in a new computing resource provisioning innovation: cloud computing. Each of these levels is virtualized and supplied as a service in cloud computing, separating issues between service, platform, and infrastructure. Networking allows distinct service levels and application services to communicate with one another.

5.4 DATA CENTRE FACILITY INFRASTRUCTURE: IMPLICATIONS FOR ENERGY EFFICIENCY

Data centres range in size from closet-sized operations to massive complexes, and their primary job is to store, process, and move data. Data centre IT infrastructure must be properly planned for efficient and successful service provisioning. IT infrastructure, on the other hand, is just a power-consuming, heat-producing "critical black box" in which the energy necessary to provide services is dissipated as heat by IT infrastructure - processors, memory, input-output devices, and other components.

As a result, the supporting infrastructure's two primary functions are as follows:

1. Ensure that electricity is available to the IT and facilities at all times.
2. Remove the heat created from the data centre to maintain the desired temperature.

We'll talk about the energy implications of these goals throughout this section.

5.4.1 Power System

The fundamental difference between a conventional office or home environment and a data centre in terms of electricity is the electrical load's 'criticality.' At most cases, losing power is merely a nuisance, but losing power to important IT services (such as in a banking institution) can be severely disruptive, if not catastrophic. To minimize such an interruption, a data centre uses a UPS and a battery bank to ensure that the important IT load receives smooth and uninterrupted power. The UPS uses the batteries to provide 'ride-through' power to the important load if the electrical supply is lost. The goal of providing ride-through power is to provide support electrical generators (typically diesel-powered) enough time to come online before the mains power supply is restored.

The practice of having redundant pathways and systems from the substation level all the way down to the server power supply level is another element of most data centre electrical power systems, providing power availability in the case of a single point failure in one of the paths. This adds to the resilience by ensuring that if one path fails, the other will continue to power the IT infrastructure. Figure 5.1 shows an example of a common power path. AC power is provided to the main utility station and routed via switch gear to the data center's substation, as indicated above. In normal operation, the UPS functions as a filter, smoothing incoming AC and sending it to the PDUs while also charge the DC batteries with AC power. In the event of a power outage, the battery DC power is converted to AC and supplied to the IT equipment racks through PDUs until the backup generators can sustain the IT load.

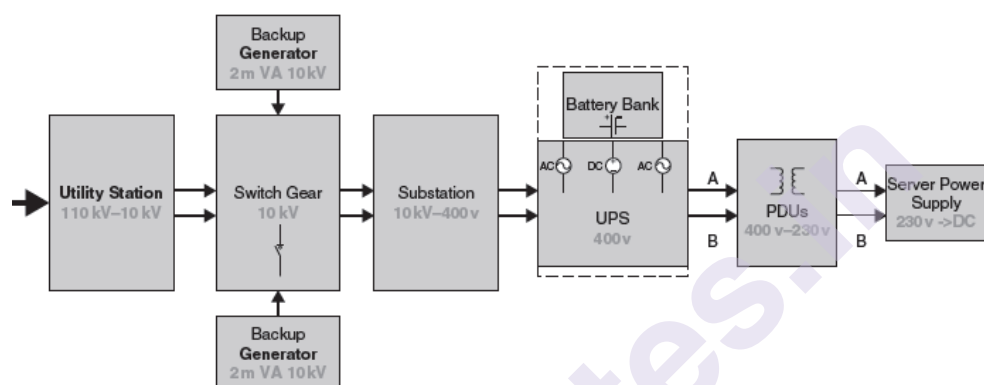


Figure 5.1 Dual power path to IT load

5.4.2 Cooling

IT servers and equipment are just heat-producing "black boxes" supplemented by heat-producing infrastructure such as UPSs, PDUs, and fan motors. Cooling ensures that components within IT equipment do not overheat, resulting in damage or performance degradation, and thereby affecting service. Convection, which is the transfer of energy between an item and its environment due to circular motion of a fluid (e.g. water) or a gas, is the most frequent method of cooling (e.g. air). Direct touch – conduction – is employed in some new cooling technologies to remove heat from servers and IT equipment, but it is not widely used yet. At the server level, air remains the most common medium for cooling IT equipment, with the interior components of the server transferring their heat to cool air pushed across the devices by server fans and then directed out the back of the chassis.

Server racks are typically arranged front-to-back to form cold aisles and back-to-back to form hot aisles in most data centres. Typically, one of three types of air-based systems are used in these facilities, all of which use a computer room air conditioner (CRAC) or computer room air handler (CRAH). The systems are classified according to their size and placement, such as 'room-based' big units, 'row-based' medium units, or 'rackbased' small units (see Figure 5.2).




Figure 5.2 Room, row and rack level cooling

Air-side economization, also known as direct-free air cooling or natural cooling represents current best practice in terms of energy-efficient heat removal and, perhaps surprisingly, is the simplest mechanical system of those used. Ambient outside air is brought into the DC and across the servers in direct-free air cooling, with the exhaust merely sent out into the open air. The only expense in normal operation is the fan energy used to move the air, while refrigerated backup is frequently included for circumstances where the ambient air temperature is higher than the intended inlet temperature.

5.4.3 Facilities infrastructure management

As previously said, the data centre is a complex ecosystem that must be effectively monitored, assessed, and managed in the traditional sense. The line between support and IT infrastructure management solutions is blurring as best practice moves toward a holistic strategy that includes both IT and facilities infrastructure. However, management has typically been compartmentalized and based on predefined maximum requirement service-level agreements (SLAs).

Facilities operational decision support tools and dashboards often need to integrate data from a variety of information systems, such as building management systems (BMSs), facilities management systems (FMSs), security, capacity, power, and even weather information systems, for effective data centre infrastructure management (DCIM) (Figure 5.3). While managing to define KPIs, the facilities manager must consider the ecosystem's interconnectedness. A data centre expert should be able to locate and display both static and dynamic information using DCIM technologies.

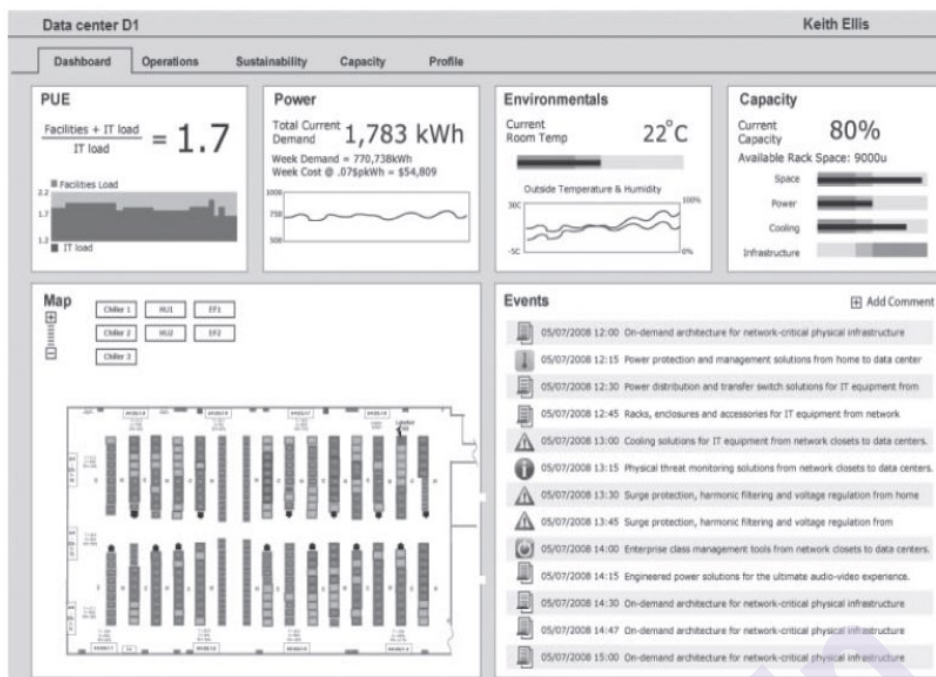


Figure 5.3 Through DCIM a holistic data centre view

5.5 IT INFRASTRUCTURE MANAGEMENT

5.5.1 SERVER POWER

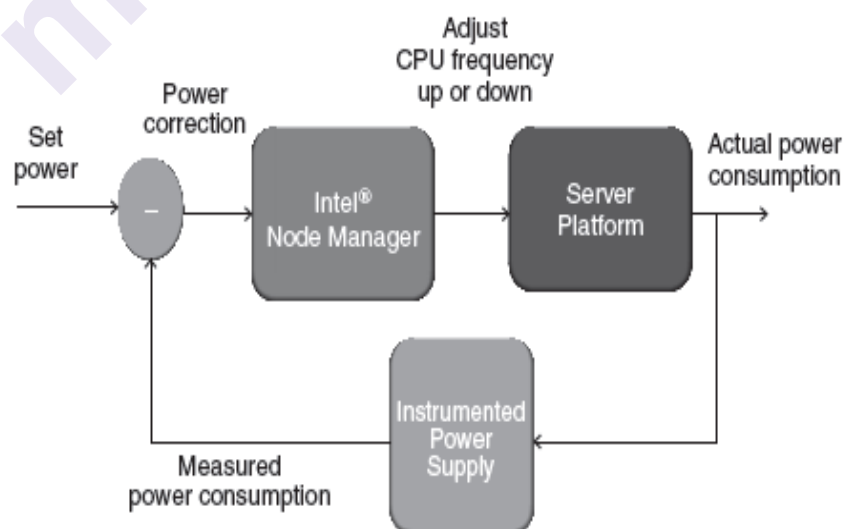
Provisioning has been the conventional method for data centre operators to meet service-level agreements (SLAs) for peak demand, not just daily but also seasonal peaks, and on top of that it starts with a large safety margin to provide for demand increase beyond the deployment equipment's estimated planning horizon. The total result is that most of the time, systems are significantly over-provisioned, with typical utilization rates in the single digits. Data centres with plenty of room for expansion and energy expenses that are a tiny percentage of the total cost of operations are implicit assumptions in this method of operation. They do not reflect the current crisis. Due to rising equipment demand on the data centre infrastructure, these assumptions are less likely to hold true. Unfortunately, even when the situation is clearly unsustainable, there are perverse dynamics that tend to maintain the status quo, such as a split incentive system in which server equipment is under the care of an IT organization with a charter to provide the best service levels possible, with overprovisioning as the preferred tool. Worse, departmental boundaries may prevent servers from being shared, resulting in widespread underutilization. On the other hand, the company's energy expenditures may be handled by the facilities department, which would have no voice in server management policy.

5.5.4.1 Need for power server management

The need for data centre services has skyrocketed as a result of new applications, particularly in the support of social media applications.

Because of improved technology and processes, the total energy used to power and cool servers in data centres has also increased, but at a far slower rate. Nonetheless, energy expenses are becoming more visible at the level of the chief information officer (CIO), giving momentum for globally coordinated power management plans. In this context, advancements in server power and energy management technology have mostly been driven by improvements in CPU fabrication processes during the last five years or more. When server systems are fully loaded, they are most efficient. With the power-proportional computing properties of today's servers, there's "pay-to-play" dynamic in effect: Even when servers are idle (i.e. delivering no work), they consume around half of the power required for full load. In theory, the best outcomes occur when applications have the ability to choose when to apply specific rules and what trade-offs these policies should entail.

In essence, this is a classic feedback process in action. In fact, there are existing feedback mechanisms in place for Intel-based servers: A feedback loop is implemented on servers equipped with Intel® Node Manager (Node Manager) technology, in which a server is given a consumption target. Embedded sensors in power supply that conform to the Power Management Bus (PMBus) measure real-time server power consumption. Firmware in the management engine (ME) microcontroller incorporated in the chipset can compare real power consumption to the consumption target and regulate processor voltage and speed as long as the control target power consumption is met. This method is known as power limitation, or power capping as it is more fondly known. The Node Manager feedback loop is depicted in Figure 5.4. The P-states of the Advanced Configuration and Power Interface (ACPI) standard pre-set voltage and CPU frequency. In conclusion, when this feedback control mechanism is placed under application control, it allows for the regulation of the equipment profile to fit workload requirement while minimizing the annoyance factor.



Source: Intel

Figure 5.4 Server power management control loop

5.5.4.2 Server power management in the data centre

Power management is a collection of IT processes and supporting technologies aimed at optimizing data centre performance while keeping costs and structural constraints in mind. For example, when racks are subject to power or thermal limits, increasing the deployable number of servers per rack makes power consumption more predictable and easier to plan for. The server infrastructure is the most energy-intensive component of a data centre, and it is a logical starting point for any thorough data centre power monitoring and control approach. The amount of power capping that can be done depends on the server architecture. The range is in the order of 20–30 % of a server's peak power consumption for current-generation servers. The combination of sophisticated server power monitoring and control technology with not only do IT processes allow for the setting of data centre energy consumption reduction goals, but they also allow for the setting of data centre energy consumption reduction goals sudden reduction in power.

5.5.2 CONSOLIDATION

A usage model for a system is officially described as a set of facts that describes how the system is used in a specific situation. The usage model data describes the user-system interactions at a level that indicates the system's advantages to the user. The emphasis is on policy-based applications. A policy is a plan of action or technique of action chosen from a variety of options and applied to specific circumstances. Three utilization models for integrated data centres are shown in Table 5.1. The usage models are arranged in ascending order of complexity. As the number of Intel Node Managers grows, so does the demand for integration utilization of technological features.

Table 5.1 Models for centralised server power control in data centres

Usage model	Benefits	Use cases
Perform real-time server power monitoring	Reduce stranded power by scheduling available data centre power to actual server power consumption	Real-time monitoring of power consumption Manage data centre hot spots Power and thermal scheduling Power trending and forecasting
Power guard rail: Impose power guard to prevent server power consuming beyond a pre-set limit	Deterministic power limit and guaranteed server power consumption ceiling	Maximize server count per rack and therefore CapEx return on investment (ROI) per available rack power when rack is under power budget with negligible per server performance impact
Static power capping: Operate servers under a permanent power-capped regime	Operate under impaired power availability conditions	Maximize per rack performance yield when rack under power budget Application power optimization Application performance compensation Business continuity: Continue the operation in the presence of power outages

5.5.3 VIRTUALIZATION

The potential for lowering energy use only through power-capture technologies is limited. If significant energy savings are to be achieved, power cuts must be deep and long-term. Furthermore, under a continuously capped regime, some energy savings are achievable, but these are limited by the capping range or the necessity to remove the capping policy to enhance performance yield. Traditional operating models, which have a hard connection between servers and the applications they execute, do not allow for additional degrees of freedom in virtualized cloud data centres. Cloud applications run in virtualized environments, allowing workloads to be dynamically consolidated. Due to the heavy load, there will be a specified nucleus of highly utilized servers that will efficiently process the workloads, allowing the remaining servers to be put into a low-energy sleep mode until needed. Dynamic power management policies take advantage of the additional degrees of freedom that virtualized cloud data centres provide, as well as the dynamic behaviors enabled by sophisticated platform power management technologies. Power-capping levels are allowed to change over time and become independent control variables. Not only may selective equipment shutdowns save energy usage, but they can also help with power management.

Traditional data centres lack a number of operational degrees of freedom that virtualized cloud data centres provide. To begin, applications hosted in virtualized cloud data centres run on virtualized operating system, which means the operating system does not run on bare metal but rather on a virtualization hypervisor. As a result, applications are no longer limited to a physical host and can be moved around within a pool of servers to improve the pool's overall power and thermal performance. Second, because of the loose coupling between applications and hosts, a group of hosts can be treated as a pooled resource, allowing group optimizations that would not be achievable with individual computers, such as shutting down some equipment during periods of low demand. As illustrated in Table 5.2, virtualized data centres enable three more advanced uses.

Table 5.2 Models for virtualized data centre power management

Usage model	Benefits	Use cases
Time-varying power capping: Adjust server performance profile to workload demand	Optimize infrastructure for quality of service (QoS) to match target service-level agreement (SLA)	Match capping set points to workload Provide support for multiple service classes
Manage data centre energy consumption for time-varying workloads	Cut electricity costs	Make a dynamic reconfiguration to achieve extreme power-proportional computing
Carry out integrated data centre power	Power optimization realized across server, communications and storage	Use server sensor data to optimize cooling equipment set points

The following section discusses some of the most critical metrics to consider while managing the data centre ecosystem holistically, effectively, and efficiently.

5.6.1 PUE and DCiE

PUE and its inverse data centre infrastructure efficiency (DCiE) are new measurements. 'Green Grid Metrics: Describing Data Centers,' a whitepaper published in February 2007 by the Green Grid, 'Energy Efficiency' is a term used to describe how In fact, DCiE was originally referred to as data centre efficiency in the original publication. However, there was some debate over whether the score implied IT efficiency. Because it does not, the name was changed to DCiE. This brings up an important point. PUE and DCiE are both indicators that indicate the effectiveness of a system. Only 'supporting infrastructure' uses power. All energy is dedicated toward the IT burden is said to produce 'productive work.' PUE is computed as shown in Figure 5.5. In terms of energy consumption, PUE has become the de facto metric in data centres. Total facility power (i.e. all power-consuming elements that make up the data centre eco-system) is divided by IT equipment power to get PUE. The statistic reveals how well or how efficiently one is supporting the IT load in terms of both time and energy. PUE should ideally be 1, indicating that no additional energy is used to support the IT load. As a result, the smaller the PUE value, and the higher the DCiE value, the better.

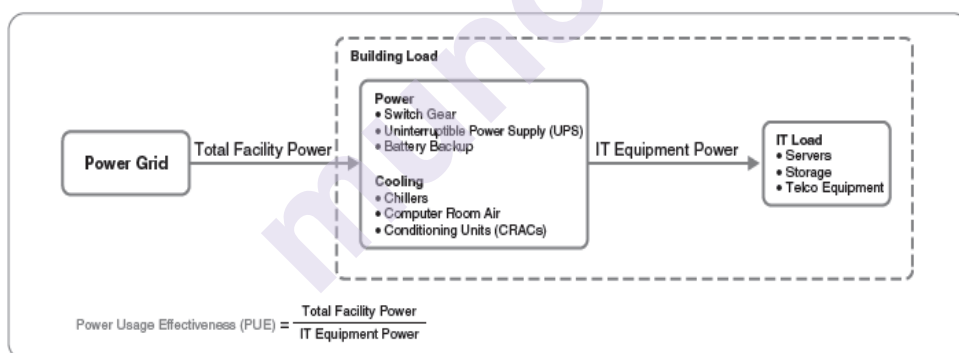


Figure 5.5 Calculation of PUE

5.6.2 POWER Vs ENERGY CONSUMPTION

While measurements such as 'watts per square foot of work cell' (Patterson et al., 2007) assist us in determining the size of a work cell, xUE measures assist us in understanding power density design and infrastructure provisioning to determine how well one's infrastructure and resources are employed to support one's mission the computer hardware It is, nonetheless, necessary to make informed energy decisions. As a result, a metre for calculating power and/or energy should be used. It may seem self-evident, yet data centre operators must use electricity or energy efficiently. As a primary statistic, they use energy consumption (kW

and/or kWh) in conjunction with other metrics PUE, for example. If the goal is to lower the amount of energy used to supply valuable IT services, relying solely on a statistic like PUE may lead to poor decisions. If, for example, the same level of IT service could be delivered using 50 kW less than the previous 100 kW while our facilities infrastructure remained largely unchanged, our total load would be 150 kW, 50 kW for the IT load plus 100 kW for supporting infrastructure, and PUE would be 3, implying that PUE would increase while we saved 50 kW on our previous total load of 200 kW. Of course, the Green Grid does not suggest focusing on any particular issue. There is no single statistic that should be used, and the industry opinion is that a set of metrics should be used to control the data centre as a whole.

5.7 Summary

The ICT sector offers a one-of-a-kind chance to help other industries improve their energy efficiency and reduce their carbon emissions. However, if the sector is to reach its full potential, it must first demonstrate its commitment to long-term ICT operations, goods, and services. We've discovered that the pervasiveness of ICT and the increasing digitalization of modern life has resulted in an increase in data centre deployment, which is now the fastest growing contributor to the overall ICT carbon footprint, and thus the 'greening' of these facilities is a key component in the sector walking the walk. Greening data centres necessitates a holistic approach that includes measuring, monitoring, and managing IT equipment and services, as well as the supporting infrastructure that keeps data centres running. The ability to modularize systems into separate servers or services at the software level has resulted from advancements in integrated circuit (IC) efficiency, which has been mirrored by marked improvements in server, storage, and network equipment efficiency, while innovations in software platforms and OS virtualization have resulted in the ability to use physical assets more efficiently.

5.8 List of References

1. Harnessing Green IT: Principles and Practices, San Murugesan, G. R. Ganadharan, Wiley & IEEE.
2. Green IT, Deepak Shikarpur, Vishwkarma Publications, 2014
3. Green Communications: Principles, Concepts and Practice- Samdanis et al, J. Wiley
4. Green IT for Sustainable Business Practice: An ISEB Foundation Guide, Mark G. O'Neill, The Chartered Institute for IT, 2010
5. About the Green Grid: <http://www.thegreengrid.org/about-the-green-grid>

5.9 Unit End Exercises

- 1] Write a note on data centre and enlist the associated energy challenges.
- 2] Explain the IT infrastructure of data centre.
- 3] Discuss on the implications for energy efficiency for the data centre facility infrastructure.
- 4] Explain the concept of IT infrastructure management.
- 5] Explain consolidation.
- 6] Write a note on virtualization.
- 7] What are the different metrics associated with data centre?

munotes.in

GREEN DATA STORAGE

Unit Structure

- 6.1 Introduction
- 6.2 Storage Media Power Characteristics
 - 6.2.1 Hard disks
 - 6.2.2 Magnetic Tape
 - 6.2.3 Solid state drives
- 6.3 Energy Management Techniques for Hard Disks
 - 6.3.1 State transitioning
 - 6.3.2 Catching
 - 6.3.3 Dynamic RPM
- 6.4 System-Level Energy Management
 - 6.4.1 RAID with power awareness
 - 6.4.2 Power aware data layout
 - 6.4.3 Hierarchical storage management
 - 6.4.4 Storage virtualization
 - 6.4.5 Cloud storage
- 6.5 Summary
- 6.6 List of References
- 6.7 Unit End Exercises

6.1 INTRODUCTION

The nature of the data, access patterns, and the value of the data all influence the complexity, design, and size of data storage solutions. Furthermore, the power consumption of various data storage methods varies significantly. Storage, for example, accounts for roughly 10% of overall data centre energy consumption in high-performance computing data centres, which often require a lot of central processing unit (CPU) resources but store less data. Storage systems utilize 25% of total energy in enterprise data centres, which typically handle a mix of compute and data storage. The majority of data is normally kept on hard disc storage systems. However, flash drives and solid-state drives (SSD)-based storage solutions are gaining popularity, especially for storing frequently accessed data. Tape-based storage systems are commonly used to store near-line or offline data. The many forms of storage medium are depicted in Figure 1.6.




Figure 6.1 Types of storage media

When moving from tape to SSD, the cost of storage and access speed rises, while overall storage capacity usually rises. In other words, a system often has slower and less expensive storage rather than quicker and more expensive storage. Massive arrays of idle discs (MAIDs) and disc virtual tape libraries (VTLs) are displacing tape in growing numbers. Hard-disk energy management and system-level energy management are two types of storage energy management. Hard discs, tapes, and SSDs are the three main storage methods that require more power than cassettes and SSDs. As a result, hard-disk energy management plays a bigger role in storage energy management. State transitioning, caching, and dynamic rotations per minute are the key strategies for energy management of individual discs (DRPM). System-level energy management is concerned with the overall management of a system's components or devices. This category includes energy management approaches for redundant arrays of cheap discs (RAID) systems, power-aware data layouts like MAID and Copan, tiered storages, virtualized storage, and cloud storage.

6.2 STORAGE MEDIA POWER CHARACTERISTICS

The power consumption of various devices is heavily influenced by their operational characteristics. First, we'll go over the different operating modes and their power characteristics in hard discs, then we'll go over the features of tapes, and finally we'll go over the characteristics of hard discs a summary of the properties of SSDs, a new type of storage media.

6.2.1 HARD DISKS

The most prevalent nonvolatile storage medium is hard drives. Disk platters on a spinning spindle and read-write heads hovering above the platters make up a hard disc drive (HDD). Data is encoded magnetically via the read-write heads. The power consumption of a hard disc in various operational (power) states can vary significantly, and energy management strategies try to take advantage of this. Magnetic hard discs' power consumption is proportional to their rotational speed and data access rate.

The revolving spindle consumes the most power, followed by the head assembly, which moves along the platters to required sectors or logical block addresses (LBAs), and the buffers, which are used for queuing requests and requesting data. The sum of the power spent by the spindle motor (P_{spindle}), the power consumed by the head movement (P_{head}), and the power consumed by other components (P_{other}), which is generally modest, has less variance, and includes the power consumed by buffers (P_{buffer}).

$$P_{\text{total}} = P_{\text{spindle}} + P_{\text{head}} + P_{\text{other}}$$

A typical strategy for storage energy management is state transitioning (the process of shifting a hard disc to different operational states based on its access pattern or idle period). Disk accesses are decreased through caching to enhance idle times, which are essential for state transitioning to work well.

The host system can access hard discs via a variety of bus types. According to the Hard drives are classified by bus type and advanced technology attachment (ATA). Serial advanced technology attachment is also known as integrated drive electronics (IDE) (SATA), SCSI, SAS, and fibre channel are acronyms for small computer system interface, serial attached SCSI, and fibre channel (FC), respectively. As a result, each of these hard drive classifications has varied power characteristics. Various hardware and operating profiles, such as the magnetic material employed, the number of platters, the amount of buffering, as well as SATA drives and the rotating speed are all factors to consider. SCSI, SAS, and FC drives, for example, are often used in enterprise computing and have high rotational speeds, resulting in higher active and idle power consumption than ATA, IDE, and SATA drives.

6.2.2 MAGNETIC TAPE

Another type of data storage medium is magnetic tape. It's made of a long, narrow strip of plastic with a thin magnetic coating. A tape drive employs a motor to wind magnetic tape from one end to the other and via tape heads to read, write, and delete data on magnetic cassettes.

Because of its low storage-per-bit cost compared to hard drives, magnetic tape is still a popular data storage media in several industries, such as radio and television broadcasting. Tape's capacity is still quite high (i.e. similar to discs) despite its lower areal density than discs. This is owing to its bigger surface area compared to disc surface. Magnetic tapes are a potential choice for backup and archiving because of their enormous capacity and low cost.

Furthermore, when compared to a hard drive, the average energy consumption of a tape drive is substantially lower because it only consumes power while it reads or writes data. In the idle or retention mode, there is no power usage. The cost of maintaining a tape library is negligible when spread out over thousands of tape cartridges.

SSDs have recently grown in popularity as a result of their improved performance and power economy. Unlike electromechanical HDD, which use spinning discs and movable heads to store persistent data, SSDs use solid-state memory (i.e. nonvolatile memory) and have no moving elements. As a result, SSDs consume less power and take less time to access than HDDs.

Nonvolatile not AND (NAND) flash memory is used in today's SSDs. Multilevel cell (MLC) and single-level cell (SLC) flash memory are the two forms of flash memory. MLCs store several bits in a single memory cell by allowing each cell to store different electrical states, whereas SLCs store a single bit in a single memory cell. A four-level MLC, for example, holds two bits in a single memory cell. MLCs are often less expensive and have more storage density than SLCs, although they are slower.

Mobile devices such as phones, digital cameras and sensor devices, notebook computers, and tablets have all embraced SSD as their primary data storage medium. Small size, low weight, low power consumption, strong shock resistance, and fast read performance are some of its characteristics. SSDs have begun to make inroads into markets ranging from laptops and PCs to enterprise-class server domains. Enterprise-class SSDs offer unique opportunity to improve the performance of I/O-intensive applications in data centres. SSDs are more appealing for high-end servers in data centres than hard drives because of their faster read performance, lower cooling costs, and larger power savings. When compared to both dynamic random-access memory (DRAM) and hard discs, SSDs have the lowest power consumption rate of active devices. For instance, consider the power consumption of a 128GB SSD consumes roughly 2 W, while a 1GB DRAM dual in-line memory module (DIMM) consumes about 1 W. For a 15000 RPM, 300GB (7200 RPM, 750 GB) hard drive, it is around 17.5W (16.6 W). This efficiency is primarily due to the nonmechanical components the nature of these gadgets.

Page units are used to read and write data on SSDs. A page read from an SSD takes around 20 times longer than a page read from a hard disc. In contrast to hard discs, there are no time differences between random and sequential reads. Because hard discs still have the benefit of cost and storage capacity, SSDs are ideally suited for caching data between memory and hard discs. However, due to SSD's low write performance and longevity issues, new and inventive solutions are needed to fully integrate SSDs into data centre high-end servers. Unlike traditional discs, whose read and write operations are symmetrical, the write process in non-enterprise SSDs is significantly slower than the read operation. Because flash memory does not allow overwrite, all write operations must be preceded by a block erase operation. A block typically covers 64–128 pages, and live pages from a block must be relocated to new pages before the wipe can be completed. A flash memory block can also only be wiped a certain number of times.

6.3 ENERGY MANAGEMENT TECHNIQUES FOR HARD DISKS

Various strategies and methodologies are being used to minimize the energy consumption of hard discs. State transitioning, caching, and dynamic RPM are three typical energy management approaches for hard discs.

6.3.1 STATE TRANSITIONING

Because the spindle motor consumes the majority of the power in a hard disc, state transitioning approaches strive to turn it off or maintain it in standby mode during idle periods. If there are no requests to be fulfilled, the disc goes into standby or off mode. The time overhead of spin down and spin up can greatly impair disc response time if the disc idle period is not long enough. Furthermore, the power used in transitioning itself may be greater than the power saved from a low-power state during the brief idle phase. If the disc has been idle for the specified amount of time, it enters standby mode. It can convert to off mode if it stays in standby mode for a certain amount of time without receiving any requests. The past data is employed in this method to forecast future access patterns. Different access forecast variations are based on past data. The majority of current state transitioning research and development is on predicting idle periods and limiting the performance impact of these transitions on disc responsiveness (due to the fact that transition times are typically around 8–10 seconds). Some state-transitioning approaches offer a guarantee of performance.

6.3.2 CATCHING

Enterprise storage solutions often feature large quantities of cache in addition to normal discs to speed up access for both read and write requests. Various strategies are also proposed for using the cache to aid in disc power control. These cache management strategies or algorithms are designed to reduce disc power consumption by reducing disc access or lengthening idle periods.

Large caches could be used to extend the idle times of discs, allowing more discs to convert to the sleep state and so boosting energy efficiency. This concept is at the heart of the cache management algorithms partition-aware least recently used (PALRU) and partition-based LRU (PBLRU). PALRU divides all discs into two types based on access patterns: priority (discs with fewer cold misses and longer idle times) and regular (discs with fewer cold misses and shorter idle times), and keeps two distinct LRU queues. When an eviction decision is made, the victims are chosen first from the usual queue elements. If the regular queue is empty, the algorithm uses the priority queue to choose elements. PBLRU, on the other hand, distinguishes between discs by dynamically altering the amount of cache blocks assigned per disc. It divides the cache into many partitions (one per disc) and modifies the partition sizes based on workload factors on a regular basis. PALRU spends 14–16 percent less

energy and PBLRU consumes 11–13 percent less energy than typical LRU, according to simulation results using online transaction processing (OLTP) traces. In the Cello96 trace (file system trace), on the other hand, PALRU saves less than 1% energy over LRU, while PBLRU is 7.6–7.7% more energy efficient than LRU.

Because write requests in business storage systems are almost never written directly to target discs (they are cached instead), another option for reducing disc power consumption is to use write offloading. Write offloading allows for complete volume spin downs on a regular basis, resulting in significant power savings. In write-dominated application environments, write offloading can save 45–60 percent of the energy used.

6.3.3 DYNAMIC RPM

Another option for hard disc energy management is dynamic RPM, which changes the rotation speed of a hard disc based on workload. It is assumed that multispeed hard discs are available, and that power usage increases as the rotational speed increases.

The rotational speed of the disc is changed in dynamic RPM based on the required disc reaction time and the performance requirement. A quick response time that exceeds the set or expected threshold is a waste of time. The objective is to restrict the amount of performance wasted by lowering the disk's rotational velocity to a level that still produces acceptable results. The possibility of constructing a single disc that can vary speeds in a cost-effective manner limits practical adoption of this strategy, however modeling studies show that a dynamic RPM scheme can save up to 60% in power.

6.4 SYSTEM-LEVEL ENERGY MANAGEMENT

The energy consumption of storage systems could be successfully regulated at the system level, taking into account media type, data properties, data access patterns, and overall system operation. RAID with power awareness, power-aware data layout, HSM, storage virtualization, and cloud storage are all common strategies for regulating energy consumption at the system level.

6.4.1 RAID WITH POWER AWARENESS

Most enterprise storage options have redundancy as a standard feature. Different levels of redundancy are provided by mechanisms such as RAID. Scheduling redundancy-related operations for power management could be used in such systems.

EERAID (energy-efficient redundant and inexpensive disc array) is a RAID engine that adaptively schedules requests to various discs in a RAID group to reduce the energy consumption of RAID discs. Controlling the mapping of logical requests to a RAID stripe, for example, maximizes the disc idle duration of a selection of discs, allowing these discs to spin down more quickly. A windowed round-robin scheduler can be used for

RAID 1, which sends a window of requests to one RAID disc before switching to the other RAID discs – and vice versa. A transformable read scan should be utilized for RAID 5. The basic notion is that when a read request is made for a stripe that is actively spinning down on disc, the stripe is reconstructed using other data and parity blocks. In addition, in the design of EERAIID, a power-aware de-stage algorithm is provided to satisfy write demands.

PARAID (power-aware redundant array of inexpensive discs) dynamically adjusts the number of powered-on discs to meet the changing load. PARAID also maintains skewed data architecture to combat the problem of large penalties caused by requests for data on spun down discs. Free space on active and idle discs is specifically utilized to store redundant copies of data that exist on spun down devices. A gear is defined by the number of active and idle discs; a gear upshift entails adding more powered-on discs to meet greater performance demands, while a gear downshift entails turning down additional drives in response to a decrease in system load. When compared to a power-unaware RAID 5, the prototype built on a Linux software RAID driver saves roughly 34% of power.

Hibernator is a storage power consumption optimization disc array. It makes the assumption that multispeed discs are available and tries to dynamically allocate them. Multiple layers of discs, each rotating at a distinct speed, must be created and maintained. The number of discs in each layer and the speed of the discs themselves are varied based on performance. The data is exchanged across layers using a disc speed determination algorithm and efficient procedures. An energy model based on trace-driven simulations has been developed. For file system-based workloads, savings of around 65 percent can be realized.

6.4.2 POWER AWARE DATA LAYOUT

Another technique to skew a disc access pattern (i.e. change a disk's idle times) is to control disc access by optimizing data layouts. Popular data concentration (PDC) is a strategy that divides data into categories based on file popularity and then migrates the most popular files to a subset of discs, extending the idle periods of the remaining discs. Maximizing idle time allows for more transitions to the standby state, resulting in higher power savings. One drawback of this strategy is that accessing unpopular files may necessitate turning on or spinning up a disc, which can take up to 8–12 seconds before actual data can be accessed.

Another method that involves data movement is MAID. MAID, unlike PDC, attempts to transfer files based on their temporal proximity. MAID uses a small number of discs as dedicated cache discs and use classic temporal locality techniques. The remaining discs are powered up as needed. However, this technique has the drawback of requiring tens of seconds to retrieve files that have not been viewed recently. MAID is better suited for disk-based archive or backup solutions due to its substantial performance penalty.

GreenStore, a solution that leverages application generated cues to better manage the MAID cache, overcomes MAID's constraint of applicability to online storage systems. By utilizing application hints, cache misses are reduced, making this solution more suitable for online situations. When compared to traditional non-MAID storage solutions, this opportunistic scheme for application hint scheduling consumes up to 40% less energy, whereas standard schemes for scheduling application hints on typical MAID systems only achieve energy savings of about 25% when compared to non-MAID storage.

6.4.3 HIERARCHICAL STORAGE MANAGEMENT

HSM, sometimes known as tiered storage, is a method of data layout management that is frequently employed in industry. IBM, VERITAS, Sun, EMC, Quantum, CommVault, and others use it in their storage systems. Data is transferred between storage tiers in HSM based on data access patterns. Price, performance, capacity, power, and function are just a few of the qualities that change significantly between storage tiers. HSM keeps track of how data is accessed, anticipates how it will be used in the future, and saves the bulk of cold data on slower devices (such as tapes) before copying it to faster devices (hard disks) when the information becomes hot. The quicker devices serve as caches for the slower ones.

HSM is usually invisible to the user, and the user does not need to know where the data is stored or how to access it. The storage of often accessible data on hard drives and infrequently accessed data on tapes is an example of two-stage HSM. If data isn't accessed after a certain amount of time, it's moved to a tape, and when it's accessed, it's moved back to a hard disc. The data is moved automatically, without the user's involvement.

A three-stage HSM can be built using FC drives, SATA discs, and tapes. If the data becomes cold (that is, if it is not accessed for an extended period of time), it will first be migrated from high-speed and high-cost FC discs to lower-speed but lower-cost SATA drives, and then from SATA discs to tapes that are even slower and cheaper than SATA if the data is not accessed for an extended period of time.

By storing infrequently accessed data on low-power devices, data can be moved from hard discs to tapes, or from FC drives to SATA discs and then to tapes, lowering storage costs and lowering storage system power usage.

A three-stage HSM can be built using FC drives, SATA discs, and tapes. If the data becomes cold (that is, if it is not accessed for an extended period of time), it will first be migrated from high-speed and high-cost FC discs to lower-speed but lower-cost SATA drives, and then from SATA discs to tapes that are even slower and cheaper than SATA if the data is not accessed for an extended period of time.

By storing infrequently accessed data on low-power devices, data can be moved from hard discs to tapes, or from FC drives to SATA discs and then to tapes, lowering storage costs and lowering storage system power usage.

Integration of SSDs into the storage hierarchy is a growing trend. Hard disc power management techniques such as state transitions can be included into HSMs made up of SSDs and drives to save power. Hard disc idle times can be extended and access to discs optimized by using the SSD tier above the hard disc tier as a layer of big cache. If an SSD is utilized to cache the most frequently requested files as well as the first portion of all other files, the hard drives can be left in standby mode for the majority of the time, saving energy. SSD will be used to access the most frequently requested data. This will give the hard drive some time to spin up from standby mode of operation.

6.4.4 STORAGE VIRTUALIZATION

Another important approach for lowering storage power usage is storage virtualization. Storage virtualization allows access to storage to be aggregated into fewer physical storage devices, lowering both storage hardware and energy costs. Furthermore, due to workload consolidation, those devices will have less idle periods, considerably improving the energy efficiency of storage systems.

Storage virtualization is widely used in data centres to manage many network storage devices, particularly in a storage area network (SAN). It adds a layer of abstraction or logical storage between the hosts and the physical storage devices, making storage system administration easier and more flexible by masking the complexity and heterogeneity of storage systems. Storage pools, which are the aggregation of physical storage devices, are used to produce logical storage. The user has no visibility into the virtualization process. It shows the user a logical space and manages the mapping of physical devices to the logical area. The mapping information is normally saved as metadata in a mapping table. These metadata will be collected in response to an I/O request in order to transform the logical address to the actual disc location. Figure 1.7 depicts a storage virtualization scenario.



Figure 6.2 Storage virtualization

Virtualization simplifies management by presenting the user with a single monolithic logical device from a central console, rather than several heterogeneous devices dispersed throughout a network. By concealing the

real storage location from the host, it allows for non-disruptive online data movement. Changes to the physical location of data can be made while I/O requests are still being processed.

By allowing numerous hosts to use a single storage device and by permitting data migration, virtualization improves storage utilization. Because of the increased use, there are fewer physical storage devices, which mean less power consumption. Furthermore, because more tasks are distributed across a single device as a result of increased utilization, each device becomes more energy efficient as a result of fewer and shorter idle periods.

6.4.5 CLOUD STORAGE

As shown in Figure 1.8, cloud storage refers to online storage typically provided by third parties, rather than storing data on local storage devices. Third parties, or web hosts, often host multiple data servers (repositories) that form data centers. Users store or access data to or from data servers using the internet through a web interface and will pay the cloud storage provider for the storage they use. In general, the fees charged by service providers are much lower than the cost of maintaining local storage for most individual users, small and medium businesses, and even enterprises.

There is a reduction in IT and energy costs as users do not have to purchase and manage or their own local physical storage, perform storage maintenance such as replication and backup, "avoid over-provisioning, worry about running out of storage , etc. All these complex and time-consuming tasks are transferred to the cloud storage provider. The convenience, flexibility and manageability offered by cloud storage, along with the affordable, making the cloud storage popular.

Through the storage systems cloud service provider offers advanced innovatory techniques for minimizing the power consumption. Also the implementation of storage virtualization techniques allows integrating different user's workload to a single storage device thereby improvising the storage efficiency and further reducing the device rest time. Cloud storage provides several opportunities for improving and creating the greener storage efficiency.




Figure 6.3 Cloud storage

6.5 SUMMARY

We have learned that hard discs, cassettes, and solid-state drives are all used in storage systems, and each has its own set of power and performance characteristics. Energy management in storage systems can be done at the disc and/or system level. State transitioning, caching, and dynamic RPM are common disk-level energy management strategies. RAID with power awareness, power-aware data layout, HSM, storage virtualization, and cloud storage are all techniques for system-level energy control. It may be important to investigate ways to acquire and use semantic information or indications from higher levels of computer systems, such as applications, for improved energy management as part of future research on energy management in storage systems. Another difficult area of research might be maximizing an entire system's energy consumption rather than just the energy consumption of the storage system, as this would necessitate a thorough understanding of energy consumption interactions across multiple subsystems.

6.6 LIST OF REFERENCES

1. Tools: <http://www.thegreengrid.org/library-and-tools.aspx>
2. http://www.apc.com/prod_docs/results.cfm?DocType=Trade-Off%20Tool&Query_Type=10
3. <http://www.intel.com/itcenter/itatintel/index.htm>
4. <http://ivi.nuim.ie/research/dcsdpn.shtml>
5. <http://www.microsoft.com/environment/greenit/>
6. Bureau of Energy Efficiency: www.bee-india.nic.in
7. Energy Management Working Group mailing list: eman@ietf.org

6.7 UNIT END EXERCISES

- 1] What are the different metrics associated with data centre?
- 2] Write a note on storage media power characteristics.
- 3] What are the energy management techniques for hard disks?
- 4] Write a note on system level energy management.

GREEN NETWORKS AND COMMUNICATIONS

Unit Structure

- 7.1 Introduction
 - 7.1.1 Background on green network communications and management
 - 7.1.2 Next generation networks challenges
 - 7.1.3 Advantages of energy efficient networks
 - 7.1.4 Objectives of green networking
 - 7.1.5 Green-Networking Technology's Core Components
- 7.2 Objectives of Green Network Protocols
 - 7.2.1 Energy - optimizing protocol design
 - 7.2.2 Objectives of green network protocols
- 7.3 Green Network Protocols and Standards
 - 7.3.1 Carbon Emissions Reduction Strategies
 - 7.3.2 The EMAN Working Group's contributions
 - 7.3.3 Standardization Bodies' Contributions
 - 7.3.4 Contextual Information to Improve Energy Efficiency
- 7.4 Summary
- 7.5 List of References
- 7.6 Unit End Exercises

7.1 INTRODUCTION

The introduction of the Internet of Things and the expansion of the number of online applications have resulted in higher network capacity demands than previously. In the United Kingdom, more than 70% of people who have broadband at home consider it important, and the average customer there spends nearly half of his or her waking hours using telecommunications products and services (Department for Culture, 2009). This expansion can be ascribed to application customization to match user needs, high levels of quality of service (QoS), and broadband service affordability. However, the global telecommunications market's continued

expansion has raised concerns about future network success if services are offered in the same way they are now. In this context, 'success' refers to a low environmental impact, continuing market penetration through persistent affordability for end users and operators, and high levels of QoS. However, there are already signs of increased telecommunications use having a detrimental influence on the environment, decreasing affordability due to growing electricity costs, and possibly lower levels of QoS due to changes in the devices on which applications are run and the nature of service supply. In terms of environmental impact, the IT industry was estimated to be responsible for roughly 10% of global electric power consumption in 2009. In 2007, it was projected that the information and computer technology (ICT) sector was responsible for 2–3% of worldwide carbon emissions. While negative environmental repercussions are unlikely to be the limiting factor in future network expansion, there are social and government-imposed commitments to prevent further environmental damage. To avoid future damage to the environment and the Internet's reputation, optimizing the effectiveness of IT operational strategies has become a priority. As a result, there is an increasing demand for green networking and communications.

7.1.1 BACKGROUND ON GREEN NETWORK COMMUNICATIONS AND MANAGEMENT

The IT industry has been chastised for its role in carbon emissions and for failing to respond to severe climate impacts (Phippa, 2009). As a result, efforts have been made in this area, and IT energy efficiency is now a top concern, as indicated by documents like Smart 2020: Enabling the Low Carbon Economy in the Information Age (The Climate Group, 2008) and Digital Britain (Department for Culture, 2009). Carbon emissions will be reduced as a result of government restrictions and programs. The Climate Change Act 2008 (Department of Energy and Climate Change, 2008), the UK Low Carbon Transition Plan (HM Government, 2009), the Carbon Reduction Commitment (CRC) Energy Efficiency Scheme 2010 (Department of Energy and Climate Change, 2010), and the Climate Change Levy (HM Revenue and Customs, 2011) are all in effect in the United Kingdom. The Bureau of Energy Efficiency enforces comparable initiatives in India, such as the Energy Conservation Act 2001 (Ministry of Power, Government of India, 2011). Overall, the goal of this program is to reduce carbon emissions from all parts of daily living. These initiatives, on the other hand, are not particular to the creation or deployment of energy-efficient IT, and for the time being, network efficiency is a separate business that is not even under the jurisdiction of telecom authorities.

The Telecommunication Standardization Sector of the International Telecommunication Union (ITU-T) standardizes telecommunications operations on an international level in order to create a fair and competitive environment. The Office of Communications (Ofcom), which oversees operations to ensure fairness to customers by promoting competition, protecting against offensive material, governing licencing procedures, researching the market, and addressing complaints in accordance with ITU-T requirements, regulates the UK

telecommunications industry. In order to achieve green IT, the International Telecommunication Union published the report NGNs and Energy Efficiency in 2008, which acknowledges the negative effects of next-generation networks (NGNs) on climate change while also examining the efficiency of telecommunications networks and applications (International Telecommunication Union, 2008). The International Telecommunication Union Telecommunication Standardisation Sector (2006) set Principles for the Management of Next Generation Networks, one of which is to strengthen network autonomy in order to maximize performance in response to real-time dynamics and trends. While communication carbon cost reduction is not explicitly defined as an ITU-T management principle, it is implied in the specification of the need to meet next-generation networking requirements by delivering services to 'any place, any time, and on any device, through any customer-chosen access mechanism,' as well as assisting network operators and service providers to 'conduct their business... In the future, solutions should be energy efficient by default to meet this criterion in compliance with ITU-T goals.

The Telecom Regulatory Authority of India (TRAI) is in charge of overseeing the operation of India's telecommunications networks. With the rapid roll-out of wireless technology in distant and hard terrain regions, Indian networks have a specific need for sustainable services. Sustainability is linked to energy efficiency: If the number of bits associated with each transmission is reduced, the financial cost of using network services for a population with lower disposable income will be lower, and operators will be able to provide lower-cost services with less network resource consumption.

7.1.2 NEXT GENERATION NETWORKS CHALLENGES

Given their transmission of data with a variety of QoS requirements and tolerance for less-than-optimal services, NGNs represent a challenge in providing energy-efficient solutions. Applications that may be transmitted across NGNs include:

1. those with real-time interactivity requirements and the ability to accommodate slight loss (e.g. voice)
2. those with real-time interactive requirements but the inability to cope with loss (e.g. online multiplayer games)
3. those without real-time requirements but the inability to cope with any loss (e.g. file transfer); and
4. those without real-time requirements but the ability to cope with any loss (e.g. file transfer).

As a result, energy-efficient communication capabilities must be able to accommodate these varied QoS demands. In addition, NGNs use a variety of carrier types to enable a wide range of applications, with traffic passing via numerous technologies on its way from source to destination. A transmission between communicative end points could, for example, travel

between nodes in a data centre or network core connected via wired lines or across wireless networks to a mobile device. As a result, QoS must be supported and redefined for contexts with varying levels of application support capability.

These characteristics of NGNs drive how energy-efficient networking solutions should be provisioned – where network intelligence happens autonomously in response to real-time dynamics, context should be collected to drive the energy efficiency process and assert appropriate actions for each network type and in response to the nature and requirements of the transmission being sent. To optimize efficiency of operation, level of service achieved, and diversity of solutions, next-generation green-networking solutions must take into account the characteristics of client devices, networks, and applications, the configurations available for each, the level of service commonly achieved across each network portion, and the ability to support application QoS requirements.

7.1.3 ADVANTAGES OF ENERGY EFFICIENT NETWORKS

Energy efficiency solutions are then developed for use in NGNs, and the management tactics employed is influenced by the energy limits and efficiency objectives of telecommunications operations. The display screen backlight dims when power saving is enabled on a notebook computer, for example, as part of a battery conservation approach. When power conservation is used in a wireless sensor network, an intermediary node may be able to 'shut down,' allowing just a limited number of probe packets to be broadcast in order to assess whether it needs to 'awaken,' and become partially or fully functional. The authors consider a variety of domains, as shown in Figure 7.1, for which energy efficiency is a limiting force on operational ability (e.g. delay-tolerant networks), for which environmental concerns arise due to the volume of emissions (e.g. data centres), and for which intelligent energy management (EMAN) is important (e.g. mobile devices) in the provision of green-networking services. Because of the goal for 'always-available' services and a wide range of devices that may be networked using the Internet Protocol, intelligent EMAN is becoming increasingly crucial in smart homes (IP). When devices are available on-demand, the QoS achieved will be higher; as a result, users may be more likely to leave devices powered on (at least in standby mode) and disable low-power options for convenience, presenting an opportunity for intelligent and autonomous device management to improve efficiency. Environmental difficulties develop in a data centre due to the large number of devices in plants that are ready to service client demands, as well as the accompanying plant management costs (such as lighting and air conditioning) incurred while maintaining a suitable working environment. In rural farming areas, energy efficient networking is especially vital (Singh, 2006); wireless solutions are easier to implement and less expensive to roll out in these areas. Efficient utilization of wireless resources contributes to the network's longevity and operational ability, as well as farmer utility and satisfaction with services delivered. As a result, the motivations for improving energy efficiency

vary by operational environment, while the priority and desire to reduce carbon emissions remain consistent.

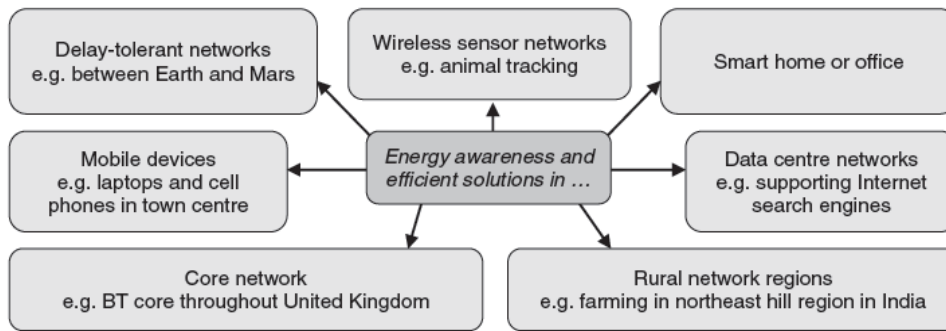


Figure 7.1 Energy-efficient network requirements for better sustainability and cost savings in domains

7.1.4 OBJECTIVES OF GREEN NETWORKING

The following are some of the goals of green network communication and management systems in a variety of domains:

- reducing the carbon footprint of delivery networks
- improving operational sustainability in wireless networks
- lowering the financial cost of transmission for operators
- allowing application QoS to be achieved within network resource constraints
- lowering network load and thus per transaction power consumption
- bridging the digital divide between urban and rural areas

Regardless of whether or not sustainability has improved, the overall requirement of green IT across domains is to reduce the number of bits per transmission in order to reduce energy demands, power costs, and carbon emissions. From this perspective, energy efficiency goals include minimizing power consumption in wired networks, as well as increasing operational sustainability in wireless networks. The efficiency demands of two example domains, the data centre and the wireless sensor network, are contrasted. Node power resources are limited in wireless sensor networks: According to a survey conducted by Sensys Networks (2007), the average battery lifetime is between 23 and 35 days. As a result, the objectives in this scenario place a premium on sustainability in order to extend the network's operational life. When responding to client requests, data centre networks, on the other hand, have a high degree of redundancy to minimize reaction time and maximize performance. However, with expected server power densities of 20 000 W/m², the carbon cost per square metre is considerable. As a result, one of the primary goals of NGN equipment is to reduce energy usage in this environment. In the development of green networking solutions, both sustainability and cost reduction needs are considered equally important, as both are a result of more efficient operation. Cost per transmission will be lowered, and sustainability will be increased, as communication efficiency improves.

7.1.5 GREEN-NETWORKING TECHNOLOGY'S CORE COMPONENTS

Networks require two essential skills (Figure 7.2) to offer sustainable and lower-cost services through green technology: energy awareness and energy efficiency. The network's capacity to measure energy cost per packet and recognize if power restrictions are becoming a limiting influence on its ability to operate or if carbon emissions are rising above a threshold is referred to as energy awareness. Energy efficiency refers to a network's ability to minimize carbon emissions compared to those experienced previous to its implementation and to extend the network's lifetime while preserving QoS.

Energy awareness provides networks with capabilities that were previously unavailable for achieving efficiency goals. The capacity to reorganize communication, including operations at each stack tier and management functions in general, is enabled by empowering this ability. The traditional focus of network communication protocol standards has been on achieving reliability for applications with stricter QoS requirements or a faster response time for those without. Improved efficiency mechanisms should, however, be included into a cross-layer strategy in next-generation green networks (NGGN) so that these goals can be met with efficient operation.

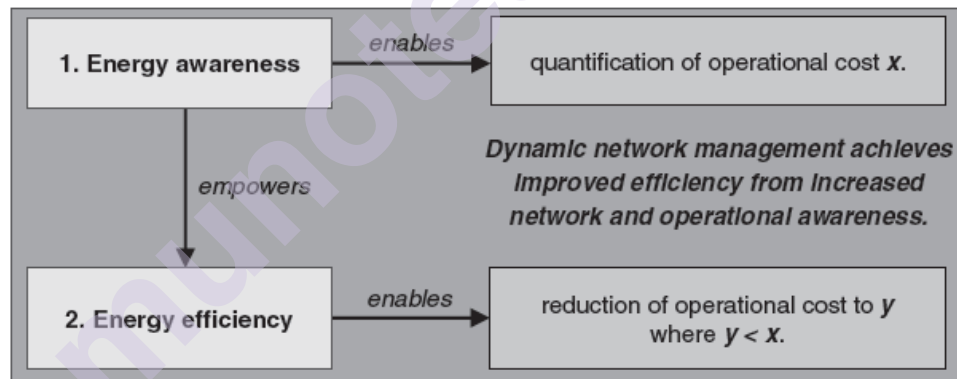


Figure 7.2 Next-generation green network components

7.2 OBJECTIVES OF GREEN NETWORK PROTOCOLS

It is vital to understand protocol overhead in terms of mandatory fields in packet headers and control packets currently utilized to govern transmissions in order to empower the network with energy awareness and efficiency capabilities. The goal of this section is to get an understanding of how protocols can be streamlined while maintaining application QoS, as well as to comprehend the mandatory content carried inside protocol headers. This leads to the discovery of techniques to optimize protocols so that their level of dependability is maintained by reducing the amount of bits associated with each, lowering the cost of transmission during any transaction, and therefore improving energy efficiency.

7.2.1 Energy - optimizing protocol design

The design of protocols employed can improve network efficiency. Communication efficiency can be improved by reducing the amount of bits connected with a transmission and reducing network stress. There will be fewer processing operations required at nodes, less finite power resources consumed during transmission, and less carbon emitted, as well as reduced network congestion, fewer re-transmissions, and an overall better optimal process, if fewer bits are communicated. The number of bits involved in network protocols can be lowered by (i) reducing the number of overhead packets per protocol, (ii) reducing the number of obligatory bits per protocol, (iii) reducing re-transmission attempts, and (iv) maximizing the number of successful data packets transmitted.

The following are the four optimization goals:

Objective 1: minimizing the number of overhead packets O per protocol:

$$\text{minimize } O \quad (1)$$

$$n \in N^*(i, j; G)$$

where O is the amount of overhead packets pushed from each node n on path (i, j) between source and destination devices in network G , and n is the set of all nodes traversed over sub-links on path (i, j) . Control and management packets transmitted across the network in support of the protocol architecture are referred to as "overhead" in this situation. When a link between nodes that want to communicate is needed, a broadcast packet is delivered, as in the Ad hoc On-Demand Distance Vector (AODV) protocol. When a message is received at an intermediary node with a route to the destination, it conveys this information to the source node, which then begins to use the route. Power consumption can be reduced by reducing the amount of overhead packets required to support protocol functioning.

Objective 2: minimizing the number of mandatory M bits per protocol:

$$\text{minimize } M \quad (2)$$

$$n \in N^*(i, j; G)$$

where M is the number of obligatory bits associated with a protocol for packets pushed from each node n that uses the protocol, and n is the number of nodes that use the protocol. Bits that must be delivered with application data in packets wrapped at each stack tier are known as mandatory bits. Fewer resources will be required to support packet transportation if the number of obligatory bits associated with a protocol is reduced, resulting in a more optimized, and hence efficient communication. Parallel to optimizing the number of bits associated with

a network transaction, sufficient detail and capability must be preserved in order to ensure operational performance.

Objective 3: minimizing retransmission attempts R :

$$\text{minimize } R_{n \in N^*(i,j;G)} \quad (3)$$

where R denotes the number of re-transmission attempts associated with traffic sent from each network node n . When appropriate protocol measures have been applied in the event that application packets have been lost or received improperly, re-transmissions refer to data packets delivered more than once via the network. In response, one or more re-transmissions may be sent.

Objective 4 : maximizing the number of successful data packet sends S :

$$\text{maximize } S_{n \in N^*(i,j;G)} \quad (4)$$

where S denotes the number of packets successfully sent from each of the n nodes.

Protocol design can have an impact on objectives 1 and 2. The traffic volume travelling over links, nodes, and client devices is used to calculate network expenses in general. The number of active ports at nodes on the end-to-end network, remaining node memory resources, dependability mechanisms connected with protocols conveying traffic via nodes, and overhead required to enable the protocol to perform its control and management role all influence the volume. In an ideal world, all of the objectives in Equations 1–4 will be met at the same time. Due to attempts to achieve minimization across conflicting parameters at the same time, this requirement set poses a restricted optimization difficulty. Using the procedures specified in Equations 1 and 2 to reduce protocol overhead may have an impact on the number of successful data packet sends and subsequent re-transmissions. As a result, achieving minimization and maximizing of all attributes as specified in Equations 1–4 will be impossible.

7.2.2 Objectives of green network protocols

The following goals are included in the development of green network protocols:

1. Improving cross-layer detail use between protocols unpacked at various stack layers
2. Minimizing or eliminating redundancy in header detail
3. Removing support for earlier versions of protocols to optimize protocol architecture.

The most significant gains in energy efficiency are expected from a multi-layer approach allowing problems and/or inefficiencies to be addressed in a consistent manner at each layer. Cross-layer compatibility ensures that any attributes deleted from a protocol header are not required at other layers. Similarly, for maximum efficiency, qualities added for greater energy intelligence should be used at a maximum number of layers.

These goals are applied to the procedures evaluated in Section 7.2.1 in the following way to increase their efficiency:

- **Achieving objective 1**

In the case of the IP, header detail may be better utilized between stack tiers. The IP and, for example, the RSVP header both use the time to live (TTL) field. Because the network layer is traversed at each node, the attribute does not need to be included in all headers, but only those needed lower in the stack, and it is appended when the previous header layer is unpacked.

- **Achieving objective 2**

At this tier, the ToS field in the IP header may be deemed superfluous. It specifies the packet's priority, acceptable latency, throughput volume, and needed level of reliability. Instead, the TTL attribute, which is also supplied in the header by default in a method designed for energy efficiency, can be used to obtain this information. The TTL, on the other hand, may be retrieved from the ToS field and does not need to be appended to the header. Because of the level of detail kept in these fields, just one characteristic is required, not both.

It's also debatable if the inclusion of source and destination addresses in a variety of packet headers used at different stack tiers is necessary. In MAC, IP, AODV, UDP, TCP, and SNMP packet headers, for example, source and destination addresses are supplied. However, with optimized UDP, the source and destination addresses may be omitted from the packet header. This information will be conveyed by the routing protocol, thus it is not necessary to duplicate it at the transport layer.

For increased efficiency, the inclusion of a checksum in each protocol header can also be explored. For the purpose of increasing communication energy efficiency, it may be able to optimize the inclusion of a checksum in all protocol headers on an application-by-application basis, particularly in the event when the application can live with a modest degree of inaccuracy.

The header fields given for IPv6 may have redundancy. For example, the hop limit field could be replaced entirely by the traffic class field. The traffic class field can be used to specify the permissible delay associated with a packet stream (traffic classes are

not specified in RFC 2460), obviating the necessity for both fields in the packet header.

In the instance of ICMPv6, there may be an opportunity to reduce the protocol's redundancy: While this protocol uses fewer error message types than ICMPv4, it may be feasible to limit the range of error codes. There are seven potential error codes indicating reasons why the destination is unreachable, for example, in the destination unreachable message. The options 'Beyond Scope of Source Address,' 'Address Unreachable,' and 'Reject Route to Destination,' for example, might be substituted with the single error code 'No Route to Destination'.

- **Achieving objective 3**

The Internet Group Management Protocol (IGMP) Version 3 established in RFC 3376, which also supports packet types associated with prior versions of the protocol, is an example of provision for support of updated versions of the protocol.

7.3 GREEN NETWORK PROTOCOLS AND STANDARDS

This section serves as a gateway between the identification of protocol costs and the methodologies recommended to provide green-networking solutions, and it includes a study of current green communication protocols and operational management. The present research gap in energy-efficient networking standards can be determined through the exploration in this area.

7.3.1 CARBON EMISSIONS REDUCTION STRATEGIES

Business for Social Responsibility (BSR, 2009) recommends techniques for reducing carbon emissions at every stage of the business life cycle, from product development to distribution. They propose that carbon reductions can be achieved by:

1. Enabling cleaner sourcing and manufacturing
2. Lowering emissions in transit
3. Enabling cleaner warehouse operations
4. Reducing transit distances
5. Removing nodes or legs
6. Reducing total volume and/or mass shipped
7. Consolidating movements
8. Contributing to other reductions
9. Increasing recycling and reuse

These carbon-reduction approaches are not limited to telecommunications networks and take into account carbon emitted during resource

transportation, development and production costs, and on-site day-to-day operations.

7.3.2 THE EMAN WORKING GROUP'S CONTRIBUTIONS

The EMAN Working Group is involved in the conversion of "work in progress" Internet draughts into formal RFC publications (their key contributions are described in Table 7.1).

Table 7.1 Internet draughts of "work in progress" contributions

'Work in progress' Internet draft	Contribution
MIB for energy, efficiency, throughput and carbon emission (Sasidharan, Bhat, and Shreekantaiah, 2010)	Defines MIB attributes required to calculate the carbon emission of network elements, with attributes including power consumed whilst performing packet throughput when idle, when operating with full power and to operate with half power
Definition of managed objects for energy management (Quittek <i>et al.</i> , 2010a)	Defines MIB structures required to appreciate the energy characteristics associated with network transactions, including a power state MIB, energy MIB and battery MIB
Energy monitoring MIB (Claise <i>et al.</i> , 2011)	Defines a number of non-operational and operational states in which nodes can exist to optimize energy efficiency, including standby, ready, reduced-power and full-power modes
Benchmarking power usage of networking devices (Manral, 2011)	Defines a power usage calculation for network devices, with attributes including the number of active ports and their utilization
Requirements for power monitoring (Quittek <i>et al.</i> , 2010b)	Defines requirements when calculating the energy consumption cost of network devices, which includes consideration for monitoring granularity and information required (state, state duration and power source)

In general, these proposals establish management information base (MIB) structures that are intended to provide networks with energy awareness so that they can attain efficiency. Carbon emissions are calculated using energy consumption, operational efficiency, and usage of each device characteristic in the MIB for Energy, Efficiency, Throughput, and Carbon (Sasidharan, Bhat, and Shreekantaiah, 2010). Quittek *et al.* (2010b) adds to Sasidharan, Bhat, and Shreekantaiah (2010) by outlining prerequisites for performing energy calculations. This entails ensuring that all network components are monitored, and that the attributes collected include the current state and time spent in each state, total energy consumed at a device and since the last monitoring interval, and current battery charge, age, state, and time when the device was last used. Claise *et al.* (2011)'s Energy Monitoring MIB captures attribute parameters such as power cost per packet, duration of power demand periods, and maximum demand in a window. Compliance with MIB monitoring processes is also addressed in this Internet draught, which includes support for both reading and writing context from and to MIBs. This standard also recommends improved operational efficiency modes, as well as 12 power states that can be applied to nodes in response to gathered context. These tactics can be equated to enabling cleaner warehouse operations by enhancing

understanding of the real-time environment and ensuring timely and suitable responses to it when compared to the BSR's principles.

7.3.3 STANDARDIZATION BODIES CONTRIBUTIONS

The Environmental Engineering group of the European Telecommunications Standards Institute (ETSI) defines systems for monitoring and controlling telecommunication infrastructure in response to gathered context and specified alarm conditions. As a result, their draught outline the alerts, events, and metrics that are required to offer the amount of management that is required. The European Telecommunications Standards Institute (ETSI, 2010) defines the minimum range of events that should be monitored on a backup generator in the AC Monitoring Diesel Back-Up System Control and Monitoring Information Model, with alarms raised if an undefined stop, start failure, fuel leakage, or battery charger failure occurs. The monitored features of a DC power system control are defined in ETSI, Environmental Engineering (2009). When circumstances such as battery failure, battery overheating, and low-voltage output are detected, alarms are triggered. The breadth of context that must be collected on an application-specific approach, as well as the customizing of alerts in relation to the domain in which management is applied, are further highlighted in these draughts. Integration of alarms, such as those proposed by ETSI, relate to cutting emissions in transit by ceasing operations when environmental circumstances are insufficient to support them, as opposed to the BSR's plans.

The Energy Efficient Ethernet (EEE) Study Group, which is part of IEEE 802.3, is working to reduce the amount of electricity needed to operate Ethernet technology. A low-power state for activation during idle periods and times of low utilization (low-power idle (LPI)) is one of the main contributions in IEEE Standard 802.3az. This option is used in conjunction with the status of the link and the observed traffic flow. When data arrives for transmission via an Ethernet link, the standard includes an alarm signal that can be used to wake up connections that have been put to sleep.

7.3.4 CONTEXTUAL INFORMATION TO IMPROVE ENERGY EFFICIENCY

ETSI defines alarms and measurements to control the operation of power systems, and the IEEE defines strategies to optimize the power required to operate Ethernet technology. The EMAN working group has proposed MIB structures specific to the challenge of improved communication efficiency; ETSI defines alarms and measurements to control the operation of power systems; and the IEEE defines strategies to optimize the power required to operate Ethernet technology. In addition, independent researchers, such as Ye, Heidemann, and Estrin (2004) and Rhee et al., provide solutions for use in certain domains and/or in response to a specific operational difficulty at a specific stack layer (2005).

Table 7.2 shows the range of context that may be required in a solution applied across domains to drive the optimization process for a variety of domains that could benefit from improved efficiency networking solutions

due to a desire to improve sustainability and/or reduce operational costs due to a currently high volume of carbon emissions. Context attributes are gathered in order to promote intelligent decision making in terms of the level of detail necessary on each individual node inside the domain as well as across the network in the larger environment. Optimization solutions in a variety of environments that use distinct contexts and to which a variety of contrasting assessments and actions can be applied are accomplished by studying the issue domain in this way. Furthermore, in addition to the standalone solutions identified in the literature, an integrated context-aware management solution that is cross-layer compatible across domains can be developed with a potential deployability and sustainability improvement in a manner similar to the TCP/IP and Open Systems Interconnection (OSI) protocol stacks that have underpinned the Internet's success to date.

Table 7.2 To achieve energy efficiency, domain-specific context data is required.

Network domain	Context used in each domain (per node)	Context used in each domain (in the wider environment between client and destination devices)
Data centre	<i>At an individual server within the data centre, context includes server utilization, packet arrival rate (packets/second), power consumption rate (Watts/second), job completion rate (seconds), operational state (per node and per port), processing delay (seconds) and page faults (faults/page)</i>	Bandwidth availability (bits per second), temperature (°C), power consumption rate (Watts per second) and operational state of neighbours (per node and per port)
Delay-tolerant network	<i>At an individual spacecraft deployed in deep space, context includes tilt of solar panel (degrees), propagation distance from neighbours (seconds), critical activities, temperature (°C), line-of-sight connectivity with neighbours (true or false), residual battery capacity (units), received signal strength (dB) and operational state (per node and per port)</i>	Wind speed (miles per hour), location of neighbours (x, y and z co-ordinates), residual battery capacity at neighbours (units), strength of signal arriving at neighbours (dB), operational state of neighbours (per node and per port), time of day, time of year, bit error rate (packets per second) and bandwidth availability (bits per second)
Mobile device	<i>For a mobile phone or laptop, context includes backlight (% brightness), residual battery capacity (units), application type of service, device type, memory capacity (bits), device-critical activities, packet-sending rate (packets per second) and location (x, y and z co-ordinates)</i>	Time of day (hours, minutes and seconds), bandwidth availability (bits per second) and location of neighbours (x, y and z co-ordinates)
Core	<i>At an individual router/switch in the network core, context includes throughput (bits per second), utilization (%), operational state (per node and per port), energy</i>	Bandwidth availability (bits per second), retransmission count at neighbours (packets per second), residual memory capacity at neighbours (bits) and bit error

Table (continued)

Network domain	Context used in each domain (per node)	Context used in each domain (in the wider environment between client and destination devices)
Rural region	<i>At an individual networked device (client device or intermediary router), context includes residual battery capacity (units), location (x, y and z co-ordinates), retransmission count (packets per second), packet transmission rate (packets per second), power cost per packet (Watts per packet) and packet arrival rate (packets per second)</i>	Temperature (°C), bandwidth availability (bits per second), residual battery capacity at neighbours (units) and time of day (hours, minutes and seconds)
Smart home and office	<i>At an individual networked device, context includes use of solar panel (true or false), device critical activities, operational state (per node and per port), time spent in state, energy cost per packet (Watts), time of last node sleep (hours, minutes and seconds) and sleep duration (seconds)</i>	Bandwidth availability (bits per second), time of day (hours, minutes and seconds), location of nodes (x, y or z co-ordinates) and operational state of neighbours (per node and per port)
Wireless sensor network	<i>At an individual sensor, context includes residual battery capacity (units), node location (x, y and z co-ordinates), operational state (per node and per port), propagation distance from neighbours (metres), temperature</i>	Temperature (°C), time of day (hours, seconds and minutes), location of neighbours (x, y and z co-ordinates), residual battery capacity at neighbours (units) and residual memory capacity at neighbours (bits)

7.4 SUMMARY

When an operation is supported by finite-resource technology, energy-efficient networking is also investigated from the standpoint of lowering the cost of communication and enhancing device sustainability. In the development of energy-efficient network solutions, these needs are given equal weight, allowing for enhanced efficiency in both wired and wireless communications. The variety of domains in which energy-efficient networking solutions can be employed to increase performance in terms of both application QoS and user quality of experience is explored from this perspective.

7.5 LIST OF REFERENCES

1. Bureau of Energy Efficiency: www.bee-india.nic.in
2. Energy Management Working Group mailing list: eman@ietf.org
3. European Telecommunications Standards Institute (ETSI) Environmental Engineering: <http://www.etsi.org/>
4. WebSite/technologies/EnvironmentalAspects.aspx

7.6 UNIT END EXERCISES

- 1] What are the objectives of green networking?
- 2] What are the different core components of green network technology?
- 3] Enlist the objectives of green network protocols.
- 4] Write a note on standards and protocols of green network.

munotes.in

ENTERPRISE GREEN IT STRATEGY

Unit Structure

- 8.1 Introduction
- 8.2 Approaching Green IT Strategies
- 8.3 Business Drivers of Green IT Strategy
 - 8.3.1 Costs reduction (including energy and operational costs)
 - 8.3.2 Regulatory and legal
 - 8.3.3 Sociocultural and political pressure
 - 8.3.4 New market opportunities
 - 8.3.5 Enlightened self-interest
 - 8.3.6 A responsible business eco-system
- 8.4 Business Dimensions for Green IT Transformation
 - 8.4.1 Economy
 - 8.4.2 Technology
 - 8.4.3 Process
 - 8.4.4 People
- 8.5 Organizational Considerations in a Green IT Strategy
- 8.6 Steps in Developing a Green IT Strategy
- 8.7 Metrics and Measurements in Green Strategies
- 8.8 Summary
- 8.9 List of References
- 8.10 Unit End Exercises

8.1 INTRODUCTION

Green information technology (green IT) strategies are essentially business strategies that include carbon emission reduction and other environmental issues as part of their design and implementation. Green IT methods that work take into account an organization's goals, industry context, and sociocultural climate. Alvin Toffler (1975) and others predicted and argued for the importance of this vital convergence of business, technology, and society. This synergy is now represented in the requirement for enterprises to expand holistically, particularly in the context of the environment.

As a result, decision-makers in a corporation must align and synchronize their business, environmental, and societal goals. The primary concept of

this chapter is the development of green IT strategies, which is based on this alignment and the business drivers for becoming and maintaining green. As a result, a green IT strategy is a strategic business approach that does not distinguish between environmental and technological concerns considerations based on business objectives. The alignment of a company's operations it's preferable to think of methods with environmental conscience as a crossroads between the company's business and the environment's long-term viability. Business issues, leadership and decision-making and critical infrastructure are all part of green IT strategies technology (including information and communications), thinking and business architecture as well as soft issues such as people, their morale, and their well-being motivation.

8.2 APPROACHING GREEN IT STRATEGIES

Philipson (2009) has improved RMIT University's original framework for green IT (Molla, 2009) into a comprehensive green IT framework that can be utilised in practise to model an organization from an environmental standpoint. These models and frameworks for green IT were aided by a procedural model for sustainable information systems management (Murugesan, 2007, 2008a, 2008b; Schmidt et al., 2009; Worthington, 2009), which provided vital input into the development of a green IT strategy.

The duration of a green IT strategy's impact is an important factor to consider while establishing one. Simple measures such as turning off monitors and PCs when not in use can be implemented immediately if the firm simply perceives IT as a producer of carbon footprint (Figure 8.1). Other methods will be used in a more strategic approach to carbon footprint reduction, and the process will take longer. Given the immediate need to reduce carbon emissions and comply with regulations, businesses are frequently forced to deploy tactical solutions to fulfil immediate demands and demonstrate the benefits of green IT, which may later become the standard or be replaced by longer-term solutions.

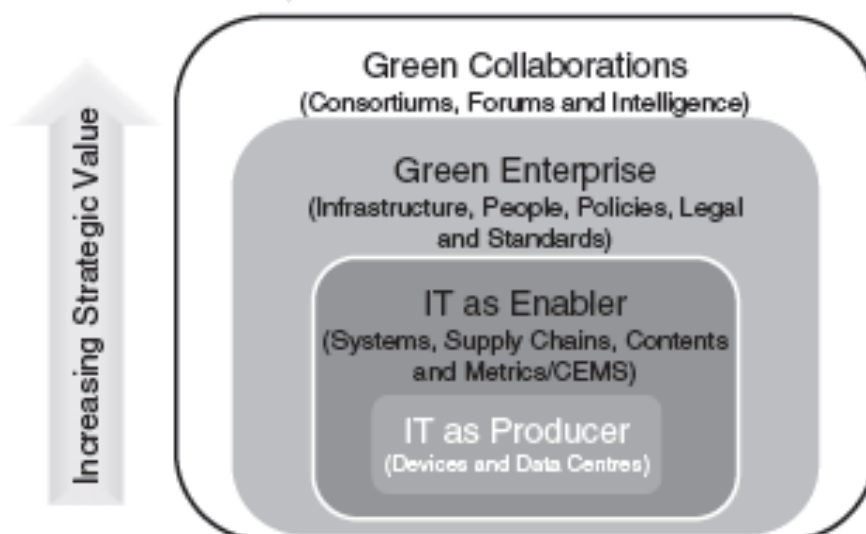


Figure 8.1 Green businesses go well beyond green IT

A strategy that cuts across all echelons and silos of a company results in effective green strategies. Individual understanding, leadership, vision, knowledge of the organization's structure and dynamics, awareness of the organization's operational nuances, and people's (i.e. stakeholders') attitude toward change are all factors that contribute to such strategies. The green strategic mentality recognizes these critical first efforts, particularly for the added visibility they provide, and adopts a long-term, holistic, all-encompassing approach. Green ICT can be used to provide tactical solutions for businesses to achieve quick wins and meet regulatory requirements, as well as enable longer-term strategic solutions across a group of organizations that make up a green business ecosystem. The impact of green IT strategies on an organization is depicted in Figure 8.2. The impact of a green IT strategy within and across an organization might range from operational to strategic-exploratory in nature.




Figure 8.2 Over the next few years, the impact of green IT strategy on businesses will change

• Operational (Immediate)

An organization's simple, quick action in the area of green IT. Switching off computer monitors while not in use, or not printing on paper as much as possible, are examples of instant acts, often known as low-hanging fruits. Simply remind users to turn off laptops when not in use, or design an internal system that charges the users' cost centre for paper usage.

• Tactical (<1 year)

On a tactical level, it will take about a year for the organization to build up its ability to reduce carbon emissions. For example, replacing the organization's existing computer monitors with green, flat-screen monitors or replacing mobile gadgets and networking equipment within a year are examples of tactical actions. Managers can also create recycling initiatives for their departments that encourage employees to recycle paper and reduce printing.

• Strategic Initial (<3 years)

Green IT projects will have a three-year impact based on strategic goals. The organization's senior leadership, including a dedicated "C-level" role (such as a chief sustainability officer), would be involved in these initiatives. These green IT initiatives are developed and authorized by the board of directors, have a sizable budget, and are based on a comprehensive approach to greening that includes the organization's data centres, facilities, supply chains, disposal methods, and even sales and marketing. Green policies, software, and measurements will provide

enormous value to a green enterprise transformation, and that value will be maximized by sticking to a 3–5-year schedule for putting such strategies and plans in place.

- **Strategic (<5 years)**

This green IT plan is a continuation of the previous three-year strategy, although it covers a wider range of topics. In addition to re-engineering activities over a three-year period, this strategy would also entail a total mindset shift among employees at all levels, a reorganization of the business architecture, and the implementation of significant oversight measures for the board. During this time, the physical infrastructure, such as buildings and data centres, will be significantly upgraded. Carbon data is used strategically not just to collect and report information, but also to identify risks and opportunities related with greening, as well as to analyze trends and patterns in terms of internal carbon savings and external carbon credits and trading. As a result, green IT plans encompass organizational capacity planning, resourcing and skills (human resource) strategies, technology acquisitions, risk management, and governance. Renewable energy sources are investigated and consumed, with fully automated, system-based monitoring, reporting, and monetization in place.

- **Strategic Exploratory (<8 years)**

Over an eight-year period, a green IT strategy must continuously study carbon reduction opportunities and attempt to connect them with the company, which will change over that time period. As a result, a long-term strategy would have strategists imagining the future in terms of technology and business and incorporating it into the green IT strategy. Such investigations are vital, especially for large and multinational corporations as well as government bodies, because they produce think tank-based output that allows organizations to prepare for a variety of prospective technologies. Such firms will, for example, have the resources to develop prototypes and assess the effects of nanotechnologies and bio mimicry (technologies that mimic nature to achieve the greatest carbon results) on their carbon emissions.

8.3 BUSINESS DRIVERS OF GREEN IT STRATEGY

Green IT strategies must be pursued and implemented for compelling reasons. Green IT business drivers can be classified into six categories, as shown in Figure 8.3 : (i) costs reduction (including energy and operational costs), (ii) regulatory and legal, (iii) sociocultural and political pressure (iv) new market opportunities, (v) enlightened self-interest, and (vi) a responsible business eco-system are all factors to consider. A combination of one or more of these drivers provides the necessary impetus for green IT strategies, policies, design, implementation, and practice. This section discusses the impact of these drivers on green IT strategies.

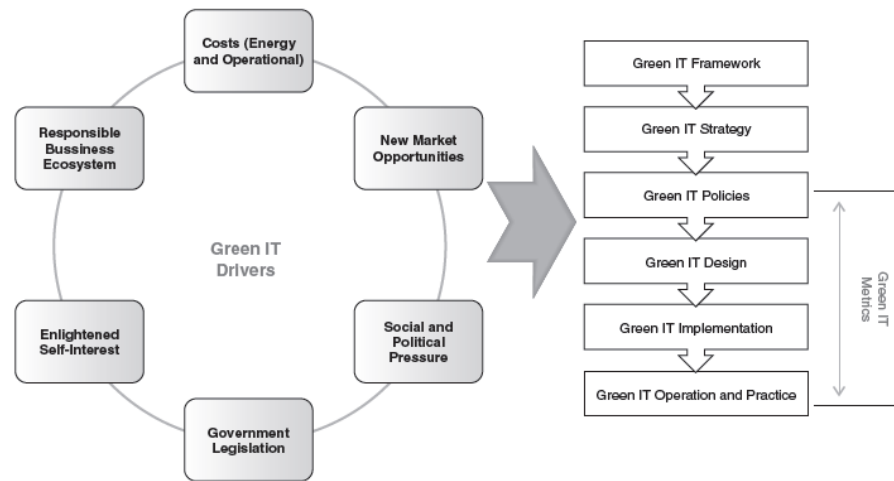


Figure 8.3 Green IT strategies for Business drivers

8.3.1 Costs Reduction (Including Energy and Operational Costs)

Cost savings are a good motivator for a company to have a complete green IT strategy. Cost savings could come from reducing energy usage (increasing energy efficiency), reducing the use of raw materials and equipment, recycling equipment and garbage, and optimizing storage and inventory as a consequence of a green effort. While cost-cutting activities can help to reduce carbon emissions, firms undergoing green transformations must be cognizant of the investment that will be required as a result of their efforts. Optimizing a business process, for example, may eliminate the requirement for a desktop machine, but it may be necessary to replace that desktop with a mobile device. Virtualizing a data centre will require some initial investment from the firm to achieve virtualization while boosting resource efficiency and lowering cooling expenses. Costs connected with a green enterprise transformation initiative, as well as anticipated cost savings as a result of the transformation, must be included in at the organizational level.

8.3.2 Regulatory and Legal

Many green enterprise transformation programs are fueled by government rules and regulations. The regulatory component has the biggest relative relevance – 70%, according to Unhelkar – when compared to other elements such as organization self-initiation, customer demand, and societal pressure (2011). Regulatory legislation such as the National Greenhouse and Energy Reporting (NGER) and the Carbon Pollution Reduction Scheme (CPRS) compel businesses to disclose carbon emissions if they exceed a particular threshold. Some basic calculators are also provided by regulatory authorities to aid in greenhouse gas calculations. The NGER calculator OSCAR is an example of such a calculator. These calculators are used to calculate an organization's total carbon emissions, which can be used to assess whether obligatory reporting is required. Green information systems, in addition to basic calculators, collect external regulatory data and store, analyze, and publish

the results, allowing a company to track and improve its performance. These organization-specific green information systems are more advanced than the regulatory bodies' basic calculators.

8.3.3 Sociocultural and Political Pressure

When an organization's society recognizes the environment as important and is engaged in safeguarding it, sociocultural and political pressures become strong driving forces. As a result of society's recognition of the importance of the environment, the organization is under pressure to adapt. For example, the growing popularity and observance of Earth Hour (the last Saturday in March in most countries), in which almost all large edifices around the world turn off their electrical power for all non-essential items for an hour, or Earth Day (the United States' 22 April and the United Nations' 20 March) have a corresponding impact on many large businesses' sustainability strategies. As a result of this 'groundswell' of beliefs, political viewpoints fluctuate as well. As a result, the company is being compelled to rethink its business priorities and operations in light of the changing environment. While the size and form of such CSR benefits for a firm might vary depending on the nature of the business and are difficult to measure their significance should not be overlooked.

8.3.4 New Market Opportunities

Global environmental awareness, corresponding legislation, and sociocultural and political pressure on businesses have created opportunities for new markets that did not exist or could not be imagined just a few years ago. These new markets, for example, can develop and provide products and services that help other organizations achieve their green objectives. As a result, we're discussing not only "green businesses," but also "green as a business offering." Carbon emissions management software (CEMS), for example, is a new breed of software that is now accessible. The creators of these new software applications have discovered a previously untapped market. Markets that are projected to grow in the carbon economy include smart metre to assess carbon emissions, chances to apply new standards for emission optimization, and innovative architecture and design of low-carbon gadgets.

8.3.5 ENLIGHTENED SELF-INTEREST

When an organization recognizes the need for, and benefits of, being environmentally responsible and creates or adopts a green plan on its own, it is acting in its own self-interest. It could include a variety of interests, such as the organization's desire to pursue a genuine common good, the need for business leadership to achieve personal satisfaction or maintain or raise employee morale, or simply the decision makers' understanding that a self-interest approach that also helps the environment can reduce costs and increase customer satisfaction. This driving also includes the desire to gain brand awareness based on environmental sustainability or an understanding of its impact on business continuity. As a result, this driving can be translated into self-motivation, and it has the potential to be a powerful green business driver. Incentives-driven compliance (IDC), a

variant of this driver, incorporates innovation and self-motivation into its environmental approach for greater carbon compliance.

8.3.6 A Responsible Business Eco-System

If a large organization changes its direction and priorities, all of its associated smaller entities must adjust their priorities as well. When a large corporation embarks on major environmental sustainability initiatives affecting its entire supply chain, it affects an entire ecosystem comprised of business partners, suppliers, customers, and internal user organizations, as well as the industry and the corresponding business consortiums in which the corporation operates. Environmentally responsible activities and strategies are constantly pushed on these many stakeholders and associations. This occurs as a result of the numerous physical and electronic contacts that occur in the course of daily company operations. There are additional considerations of superimposition of long-term trends upon short-term markets, in addition to the impact of partnering organizations in a business ecosystem. Such superimpositions result in significant company transformations and restructuring in favor of environmentally friendly ideas.

8.4 BUSINESS DIMENSIONS FOR GREEN IT TRANSFORMATION

Once the drivers for a company's green initiatives have been identified and documented, it's time to talk about which parts of the company are likely to be affected by the changes. Because the changes brought about by green IT initiatives impact an organization, comprehending them is an important aspect of a green IT strategy. As shown in Figure 8.4, an organization alters or transforms along four main lines or dimensions: economy, technical, process, and people. These business transformation dimensions apply to any type of transformation and can also be thought of as factors that change as an organization evolves.

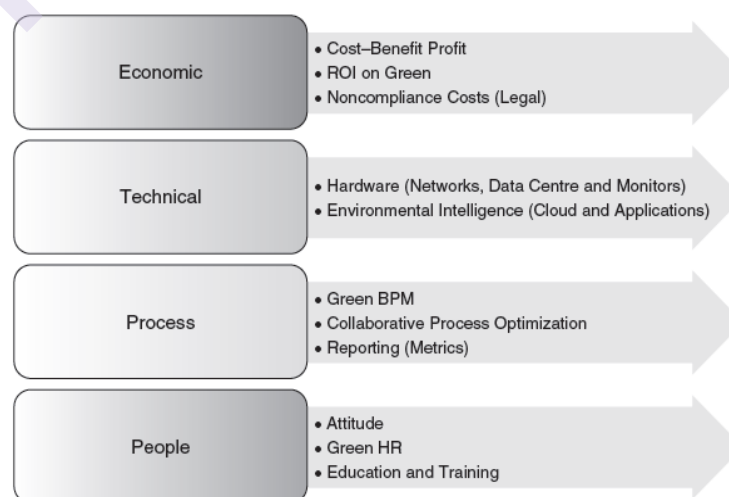


Figure 8.4 In the execution of a green IT strategy, there are four factors to consider: economy, people, process, and technology

One of the most important aspects in an organization's decision to incorporate environmental policies and procedures is cost. The costs of green transformations, as well as the return on those costs, are the first things that come to mind for leaders and people in charge of the green transformation. As a result, this is a key axis along which green transformation occurs in a company. The cost-benefit analysis and a financial return on investment (ROI) analysis are two examples. In today's economy, economic expansion is frequently accompanied by an increase in carbon emissions. This is especially true in developing economies, when many industries are growing, resulting in increased emissions across the board rather than just in one company. For example, as the Copenhagen conference in 2009 demonstrated, this economic factor causes friction between the 'developed' and 'developing' worlds. The disparity between the developed world's resource usage and the requirement for BRIC (Brazil, Russia, India, and China) and other developing countries to minimize their carbon footprint in order to save the environment for future generations can lead to economic and legal snags.

8.4.2 Technology

In this context, technology refers to the hardware, network infrastructure, software, and applications of a company. This is also the more obvious and well-known element of green IT. The early, apparent features of change along this dimension include turning off monitors, virtualizing servers, and avoiding printing on physical paper. The data centre is then organized (including the physical facility, the rack system, and the actual servers themselves) and operated in a long-term strategic manner. In this dimension, emerging information technologies such as service orientation, software as a service (SaaS), and cloud computing are creatively employed to lower a company's overall carbon emissions. Environmental intelligence is created when business intelligence (BI) tools are extended and augmented with carbon data.

8.4.3 Process

The 'how' things are done within an organization is dealt with by the process dimension of an organization. The fundamental rethinking and drastic rebuilding of business processes to produce major gains in critical, modern performance measures such as cost, quality, service, and speed is known as business process re-engineering. Green business process management (BPM) is the process of modeling, studying, and optimizing a company's processes in order to improve its environmental credentials. This effort entails optimizing existing operations and establishing new environmentally conscious processes that will not only cut carbon emissions but also improve customer satisfaction. The process dimension of a business is likely the most obvious, and it is frequently used to assess a firm's level of environmental responsibility for its green ICT. This is due to the fact that the process dimension has immediate and quantitative

effects on the carbon footprint of a corporate operation. It also has an impact on the collaboration's clients, vendors, and business partners.

8.4.4 People

People are the most challenging and possibly most complex aspect of a green enterprise transition. While the human side of an organization's behavior has been extensively researched, the focus of this discussion is on individual attitudes and the sociocultural setting in which they act in the context of the environment. The same societal impetus that propels an enterprise toward green IT also poses hurdles when it comes to implementing that change. In addition to addressing the people component in the framework of corporate leadership for green IT transformation, there is a big difficulty in addressing the individual employee and customer at the grassroots level. To ensure the success of an enterprise-wide green strategy, it is best led from the top of the business. Leadership in this area, such as that provided by top directors and chief officers, can make or break an environmental endeavor. Senior management engagement in bringing about a change in the people dimension is critical - and it must be done early in a green effort, even though such involvement from senior leadership demands a significant time, money, and other resource commitment.

8.5 ORGANIZATIONAL CONSIDERATIONS IN A GREEN IT STRATEGY

The following are key practical issues to recognize and consider when developing a complete green IT strategy for businesses:

- Build the green IT strategy on the conviction that reducing carbon footprints and achieving business results are not mutually exclusive – and, in fact, can be leveraged.
- Recognize that the concept of a green enterprise is holistic and subjective, encompassing personal, individual, and attitude difficulties.
- Create numerous provisions to address the ambiguity of carbon emission regulations and norms.
- Shift away from businesses' desire to focus solely on obtaining quick advantages and toward a more comprehensive strategy. The allure of so-called low-hanging fruits, such as turning off physical carbon-emitting technology (e.g. monitors and data servers) right away, may overshadow more comprehensive carbon reduction.
- CEMS can be integrated with an organization's existing packages and systems.
- Incorporate cost–benefit analysis into measures related to green projects, focusing on environmental initiatives' payback.

- Manage the risks associated with the usage of technology-based efforts in the area of green initiatives, such as cloud computing, business intelligence, and knowledge management.
- Consider the design and construction of buildings and accompanying infrastructures (when possible). If attention is made to an infrastructure's initial design and construction from a carbon reduction standpoint, it has the potential to significantly reduce its emissions.
- Think of data centres as specialized structures that hold data and computer servers, as well as an organization's network equipment, and that require strategic attention during the planning and installation stages.
- To modify current attitudes and establish a route for change, provide education and training (attitude and culture) to the workforce, but also consider green human resource assistance that fosters attitude changes.
- As the organization strategies for a green transformation, consider and plan for technology (hardware, servers, and network) changes that will inevitably occur. Reusing and recycling current hardware, as well as developing plans for replacing it with new, more carbon-efficient gear, are all factors to consider.
- Consider upgrading applications and systems in two areas: first, to enable the inclusion of carbon data into existing applications and systems, and second, to plan for a new CEMS that is solely dedicated to collecting, storing, analyzing, and reporting on carbon data.
- Carry out green process re-engineering and management, utilizing previously used re-engineering methodologies such as identification, modeling, and process optimization.
- Include green metrics and measurements in the identification and modeling of an organization's "as is" and "to be" states.
- Make sure that legal compliance is a key component of any green IT strategy. Legal requirements might range from local and state legislation to national carbon legislation. International consortiums and summits also impose legal obligations that must be integrated into the ERBS.

8.6 STEPS IN DEVELOPING A GREEN IT STRATEGY

This section outlines the stages involved in creating a green IT strategy. Depending on the type, size, location, and industry vertical, the final strategy document may differ, and those differences must be taken into account. The major phases that can be completed iteratively (three times) to establish a green IT strategy are listed below (Figure 8.5).

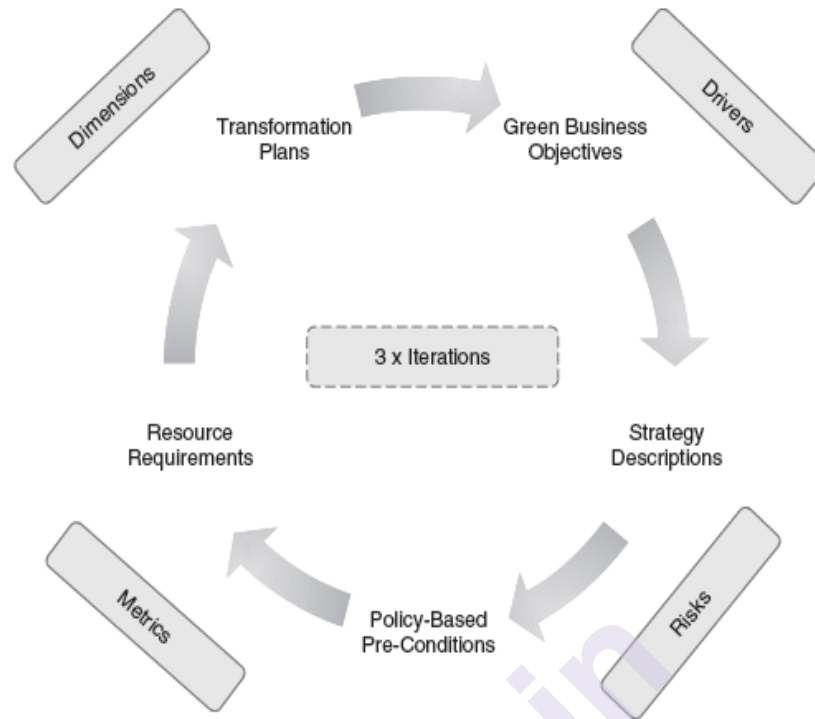


Figure 8.5 Developing a green IT strategy: a step-by-step guide

- **Aligning green IT strategy with business objectives:** The key objectives for a company undergoing a green transformation are green business objectives. These goals must be in line with an organization's carbon and commercial objectives. The primary drivers influencing the organization can help you grasp the business goal of going green. If, for example, cost reduction is the primary motivator, this will be reflected in the green IT strategy's business purpose and the KPIs associated with it.
- **Descriptions of the strategy:** This section explains how the company will go green. The economic, technical, process and people dimensions of green transformation will be incorporated into the strategy description. The amount of time needed for implementation will also be considered. The development of green IT strategy descriptions might easily take 1–3 months (or more, depending on the size of the firm). Green enterprise transformation programs have an impact on the organization as a result of strategy definitions. This also necessitates familiarity with the present strategic strategy, if one exists. SMART (specific, measurable, achievable, realistic, and timely) goals can be used to describe strategies.
- **Green enterprise transformation plan and timeline:** The final and most critical step in building a green IT strategy is to create a green enterprise transformation plan. A transformation plan is a project plan that includes activities, roles, and deliverables, as well as a delivery schedule. Typically, this transformation project plan serves as a blueprint for change. This plan is typically divided into two parts: a high-level road map that specifies important work areas,

deliverables, and dates, and a detailed, task-by-task project plan that employs all known project and program management methodologies.

- **Iterations and hazards:** Developing a green IT strategy should not be a one-way street. Instead, it should be created through an iterative process that includes all drivers, dimensions, risks, and metrics. To get at a final, comprehensive, and actionable green IT plan, you may need three revisions. These iterations, which can last anywhere from three to six months, may also include studying industry trends and new advances in green IT. On the basis of these patterns, green policies must be updated. These iterations also imply a strategy for putting the policies into action, which should be based on iterations as well. During practice, it is intended that the policies would be refined iteratively. The concepts of continual improvement in processes, people, and technologies that give lean business endeavor energy also apply here.

8.7 METRICS AND MEASUREMENTS IN GREEN STRATEGIES

Metrics for an organization's green IT performance can be based on internal ROI targets and/or legal reporting obligations. While the ISO 14000 family of standards can serve as a good starting point for green IT Key Performance Indicators (KPIs), CEMS can be used to automate, monitor, and report on carbon emissions and footprints.

The following are some common KPIs that must be implemented in a company that is pursuing green initiatives:

- **Result in terms of money:** Reduce energy consumption by 10% per year for three years; expand green services (for example, the inclusion of one detailed insurance service dedicated to green).
- **Expertise:** For all warehoused data, employ virtualized data servers; for recording, re-posting, and controlling emissions, use smart metres.
- **The procedure:** Reduce or re-engineer particular processes through optimizing supply chain management.
- **Individuals:** At all levels, individuals should be trained in green IT. Reduce your carbon footprint by telecommuting once a week.

The KPI groupings outlined here can be expanded further, and depending on the dimension to which they belong, they will have their own peculiarities. A KPI that is solely focused on carbon reduction without regard for cost considerations, for example, may not be acceptable to an organization's economic component. A technologically advanced energy-efficient cooler, on the other hand, might use less energy and hence have a reduced operational cost, but the capital costs of such a cooler will be

included in the economic dimension. As a result, KPIs should be tied to business efficiency and carbon reduction efficiency on a regular basis. Carbon savings in many sectors of a company, like as production, sales and marketing, research and development, and administration, must be linked to cost savings.

8.8 SUMMARY

The notion of green IT strategy is also detailed in the concept of green IT strategy from an immediate, tactical approach to a long-term strategic approach, starting with the basic premise that carbon reduction measures should not be distinct from an organization's commercial goals. The six factors that give an organization the impetus it needs to transform into a green business, as well as the four dimensions (economic, technical, process, and people) along which such transformation might occur were explored. The steps and procedure (iterations) for building a green IT strategy for a company were outlined. The future of green IT strategies lies in developing and integrating them into an organization's regular business processes, monitoring the results, and demonstrating the carbon savings that will serve as the foundation for growing carbon economy activities such as carbon trading.

8.9 LIST OF REFERENCES

1. IEEE P802.3az Energy Efficient Ethernet Task Force: <http://grouper.ieee.org/groups/802/3/az/public/index.html>
2. International Telecommunication Union Telecommunication Standardisation Sector (ITU-T): www.itu.int/ITU-T/
3. The Office of Communications: www.ofcom.org.uk
4. Telecom Regulatory Authority of India (TRAI): www.trai.gov.in/

8.10 UNIT END EXERCISES

1. What are the different strategies for approaching green IT?
2. Explain business drivers of green IT strategy.
3. Enlist the business dimensions for green IT transformation.
4. What are the different organizational considerations in a green IT strategy?
5. Discuss different steps in developing a green IT strategy.
6. Write a note on metrics and measurements in green strategies.



SUSTAINABLE INFORMATION SYSTEMS AND GREEN – METRICS

Unit Structure

- 9.0 Objectives
- 9.1 Introduction
- 9.2 Multilevel Sustainable Information
- 9.3 Sustainability Hierarchy Models
 - 9.3.1 Sustainability Frameworks
 - 9.3.2 Sustainability Principles
 - 9.3.3 Tools for Sustainability
- 9.4 Product Level Information
- 9.5 Individual Level Information
- 9.6 Functional Level Information
- 9.7 Organizational Level Information
- 9.8 Measuring the Maturity of Sustainable ICT
- 9.9 Let us Sum Up
- 9.10 List of References
- 9.11 Bibliography
- 9.12 Unit End Exercises

9.0 OBJECTIVES

After going through this unit, you will be able to:

- To get knowledge on green and sustainable informatics and metrics as a tool to monitor and drive sustainable behavior.
- To explore the multilevel nature of green metrics with examples at the regional, organizational, functional, product/ service, and individual levels.
- Energy efficiency within data centers and assessing the maturity of sustainable information and communication technology (SCIT) within organizations.
- To develop a sustainable strategy for local government, with relevant KPis

9.1 INTRODUCTION

The first wave of green Information Technology (green IT), greening of IT, aims to reduce 2% of global emissions from IT by reducing the footprint of ICT by actions such as improving the energy efficiency of hardware (processors and disk drives), and waste from obsolete hardware. The second wave of green IT, greening by IT, also called Green IT 2.0 or Sustainable IT is shifting the focus towards reducing the remaining 98% by focusing on the innovative use of IT and Information Systems (IS) in business processes to deliver positive sustainability benefits beyond the direct footprint of IT, such as monitoring a firm's emissions and waste to manage them more efficiently. Relevant and accurate data, information, metrics, and key performance indicators (KPIs) are critical to supporting sustainable practices, and the need to manage this information has led to the emergence of green information systems (GIS) as a field. Green IS is the engine driving both the strategic and operational management of sustainability.

The chapter starts with an overview of the need for multilevel sustainable information, and introduces the wider context of sustainability frameworks, principles, and tools. It then examines the information requirements and methods utilized at multiple levels including regional (regional sustainability plan), organization (GHG emission reporting), business function (data center energy efficiency), individual (commute tracking), and product and service (life cycle assessment (LCA)). It also examines how the sustainability capability of an organization can be examined to determine its effectiveness and highlight areas for further research.

9.2 MULTILEVEL SUSTAINABLE INFORMATION

The goal of **sustainable development** is to meet the needs of today, without compromising the needs of tomorrow. This means we cannot continue using current levels of resources as this will not leave enough for future generations. Stabilizing and reducing carbon emissions is key to living within environmental limits. Sustainability requires information on the use, flows and destinies of energy, water and materials including waste, along with monetary information on environment-related costs, earnings and savings.

9.3 SUSTAINABILITY HIERARCHY MODELS

- Natural laws : The laws imposed upon us by Mother Nature such as the Law of Conservation of Mass and Energy, Law of Entropy and Laws of Thermodynamics.
- High-level conceptual rules for sustainability
- Principles : General and sector-specific guidelines that detail sustainable practices and actions.

9.3.1 Sustainability Frameworks

To improve the understanding and communication of sustainability issues and to provide high-level definitions of sustainability, several frameworks have been proposed. Natural Capitalism, The Natural Step, Ecological Footprint, and the triple bottom line (TBL) are popular frameworks and are described in this section. These frameworks are used to reach a shared mental model of sustainability, providing concrete definitions, and scoping for concepts and terms. The frameworks can be used to develop high-level visions and planning for sustainability activities.

a) Natural Capitalism

The concept of Natural Capitalism is to look at the environment as a system that yields a valuable flow of goods and services (e.g., fish, trees, and other such things). *Natural capitalism* is a system, where business and environmental interests overlap, and in which businesses can better satisfy their customers' needs, increase profits, and help solve environmental problems all at the same time. Natural capitalism requires the following shifts in business practices: Radically increase productivity in the use of natural resources; shift to biologically inspired production models and materials; move to a service-and. flow' business model; and reinvest in natural capital.

b) Ecological Footprint

The ecological footprint is the measure of environmental impact of actions on the Earth's natural resources and ecosystem functionality. Ecological footprint analysis is widely used around the Earth in support of sustainability assessments. It enables people to measure and manage the use of resources throughout the economy and explore the sustainability of individual lifestyles, goods and services, organizations, industry sectors, neighborhoods, cities, regions and nations.

c) Triple Bottom Line (TBL)

The TBL framework allows an organization to focus on not only its economic bottom line, but also its environmental and social 'bottom lines'. TBL expands the scope of responsibility for an organization from shareholders and owners to also include stakeholders - anyone who is influenced, either directly or indirectly by the actions of the organization. The triple bottom line (TBL) is a framework.

9.3.2 Sustainability Principles

Below the high-level frameworks are a broad range guidance on the sustainability practices needed to achieve the objectives set out in higher level frameworks. Principles are typically created via a group or community process. Two well known general principles mentioned in this section are 3Rs and Earth charter.

a) Reduce, Reuse and Recycle (3R's)

The 3Rs classify the strategies according to the desired order of use. It is better to reduce then reuse, and better to reuse than recycle. The aim of the waste hierarchy is to extract the maximum practical benefits from products and to generate the minimum amount of waste.

b) Earth Charter

Earth Charter is an international declaration on the principles needed to build a sustainable and peaceful global society. The Earth Charter proposes that environmental protection, human rights, equitable human development.

9.3.3 Tools for Sustainability

At the bottom of the hierarchy are the sustainability tools that provide guidance and best practices. Sustainability tools come in many forms, from methodologies used to help determine the environmental impact of a steel water bottle, to certifying the design and construction of a building. Sustainability tools can be general purpose, such as the GHG Protocol for emissions reporting, or specific to a sector, like the Leadership in Energy and Environmental Design (LEED) for green building certification. These tools provide sustainability information needed to make sustainable decisions and drive sustainable behavior. Many of these tools can be classified as IT for green or as an opportunity for IT to 'green' processes

9.4 PRODUCT LEVEL INFORMATION

Understanding the impacts of a product or service requires an analysis of all potential impacts associated with a product, process, or service for its entire life cycle. This is achieved using a technique known as life cycle assessment.

Life-Cycle Assessment

Life-cycle assessment (LCA), also known as life cycle analysis, is a technique to systematically identify resource flows and environmental impacts associated with all the stages of product and service provision. LCA provides a quantitative cradle-to-grave analysis of the products or services' global environmental costs (i.e., from raw materials through materials processing, manufacture, distribution, use, repair and maintenance and disposal or recycling). The demand for LCA data and tools has accelerated with the growing global demand to assess and reduce GHG emissions from different manufacturing and service sectors.

9.5 INDIVIDUAL LEVEL INFORMATION

Primary sources of emissions for an individual will include the life cycle of products and services they purchase and use, the construction (LCA of material and construction) and operation of their residence (energy use,

water use and waste disposal and recycling) and private travel (especially emissions from long-haul flights). Studies have shown that energy conservation improves if residents receive real-time feedback on their energy usage; one study in Ontario showed the average household reduced its electricity usage by 6.5%. Another interesting development is the utilization of smart phones as a low-cost means to determine the impacts associated with commuting.

9.6 FUNCTIONAL LEVEL INFORMATION

An organization typically comprises multiple functions or operating divisions and departments including marketing, finance, operations, human resources and research and development. Sustainability issue at the functional level is related to business processes and the value chain; they involve the development and coordination of resources through which the overall organization's objectives and goals can be executed efficiently and effectively. Monitoring energy usage for an IT action can involve metrics around building energy usage (lights and heat), staff computer energy usage and in particular energy used within data centers. Here, we examine energy usage within the data center and the role of metrics.

a) Data Centre Energy Efficiency

The power needs of data centers may range from a few kilowatts for a rack of servers in a closet, to several tens of megawatts for large facilities. Power usage in a data center goes beyond the direct power needs of servers to include networking, cooling, lighting, and facilities management. The US Environmental Protection Agency (EPA) estimates that servers and data centers are responsible for up to 1.5% of the total US electricity consumption or roughly 0.5% of US GHG emissions for 2007. With electricity costs the dominant operating expense of a data center, it is vital to maximize the center's operational efficiency to reduce both the environmental and economic costs.

b) Emerging Data Centre Metrics

When assessing the financial health of a business, one should not look at one metric in isolation. The same is true for assessing the efficiency of a data center. Whilst PUE and DCiE have proven to be effective industry tools for measuring infrastructure energy efficiency, there is a need to measure the operational effectiveness of the data center. To this end, several metrics are under development to measure dimensions including resource utilization and environmental impact. Each of these metrics provides data center operators more visibility into where opportunities for further efficiency improvements exist.

c) Carbon usage effectiveness (CUE)

Measures data center to specific carbon emissions, which are emerging as an extremely important factor in the design, location,

and operation of a data center. CUE, combined with PUE, can assess the relative sustainability of a data center to determine if any energy efficiency and/or sustainable energy improvements need to be made. **Note** : CUE does not cover the emissions associated with the data center or the building itself. CUE is defined as:

$$CUE = \text{CO}_2 \text{ emitted (kg CO}_2 \text{ eq)/unit of energy (kWh)} \times \frac{1}{\text{(Total Data Centre Energy/IT Equipment Energy)}}$$

d) Energy reuse effectiveness (ERE)

Many data centers are now recovering waste energy from their operations and reusing. Since PUE does not consider these alternate uses for waste energy, the ERE metric is used.

e) Data center energy productivity (DCeP)

DCeP quantifies useful work compared to the energy it requires. It can be calculated for an individual IT device or a cluster of computing equipment. DCeP is a sophisticated metric where useful work and energy are compared relative to a user-defined time limit.

$$DCeP = \frac{\text{Useful Work Produced}}{\text{Total Data Centre Energy Consumed over time}}$$

f) Data center computer efficiency (DCcE)

DCcE :- and its underlying sub metric, server compute efficiency (ScE), enables data center operators to determine the efficiency of their compute resources, which allows them to identify areas of inefficiency. The metric can reveal unused computer resource with a data center making it easier for data center operators to discover unused servers (both physical and virtual) and then decommission or redeploy them.

g) Environmental consumer chargeback

Carbon dioxide and environmental damage are gaining acceptance as viable chargeable commodities. Using environmental chargeback models, data center operators can "chargeback" the environmental impacts (i.e. CO2 emissions), in addition to the financial costs, of their services to the consuming end users. Environmental chargebacks can have a positive effect on environmental impacts by linking consumers to the indirect impacts of their service usage, allowing them to understand the impact of their actions.

Energy consumption metrics do not tell the full story of the impacts of data centers on the environment. To understand the full Environment's burden of a data center a full life cycle assessment of the data center facilities and IT equipment is needed. These additional costs should not be under-estimated.

9.7 ORGANIZATIONAL LEVEL INFORMATION

Organizational information derives its meaning from the sense-making frameworks that characterize specific **organizations**. Many **organizational** members need **information** to fulfill their responsibilities, so other members gather data and convert it into **information**. Sustainability strategy and direction are determined at the organization level based on macro concerns; the overall sustainability goals of the organization, determining the types of businesses and activities in which the organization should be involved and defining organizational responsibilities.

Organizations need to determine their goals and objectives for sustainability. Typically, organizational sustainability also involves one or more of the following:

- Develop significant capabilities
- Keep pace with industry or stakeholder expectations.
- Meet minimum compliance requirements and reap readily available benefits.

Performance metrics and KPIs are used to measure environmental, social and economic impacts, and to ensure the delivery of strategic sustainability-

related objectives. Sustainability performance measures and KPIs help organizations to establish progress against sustainability goals and to ensure that they cover their environmental, social and economic impacts.

Reporting Greenhouse Gas Emissions

A **greenhouse gas (GHG)** is a **gas** that absorbs and emits radiant energy within the thermal infrared range. **Greenhouse gases** cause the **greenhouse effect**. The primary **greenhouse gases** in Earth's atmosphere are water vapor, carbon dioxide, methane, nitrous oxide and ozone. The GHG Protocol is the most widely used international accounting tool for government and business leaders to understand, quantify and manage the emissions of all six majors. GHGs: carbon dioxide (CO₂), methane (CH₄), nitrous oxide (N₂O), hydrofluorocarbons (HFCs), perfluorocarbons (PFCs) and sulfur hexafluoride (SF₆). The GHG Protocol is based on five principles: relevance, completeness, consistency, transparency, and accuracy.

9.8 MEASURING THE MATURITY OF SUSTAINABLE ICT

IT organizations need to develop a sustainable information and communication technology (SICT) capability to deliver sustainability benefits both internally and across the enterprise. To assist organizations, understand the maturity of their SICT capability, several tools for

measuring SICT maturity have been developed including the G-readiness framework which provides a benchmark score against SICT best practices, or the Gartner Green IT Score Card which measures corporate social responsibility (CSR) compliance. In the remainder of this section, we examine the SICT-Capability Maturity Framework (SICT-CMF) from the Innovation Value Institute (IVI). The SICT-CMF Maturity provides a comprehensive assessment to determine current maturity level and a set of practices to increase SICT capability; performing an assessment allows an organization to identify capability gaps and identified opportunities to improve SICT performance.

a) A Capability Maturity Framework for SICT

The SICT-CMF assessment methodology determines how SICT capabilities are contributing to the business organization's overall sustainability goals and objectives. This gap analysis between what the business wants and what SICT is achieving positions the SICT-CMF as a management tool for aligning SICT capabilities with business sustainability objectives. The SICT-CMF gives organizations a vital tool to manage their sustainability capability. The framework provides a comprehensive value-based model for organizing, evaluating, planning and managing SICT capabilities. The framework focuses on the execution of four key actions for increasing SICT's business value :

- Define the scope and goal of SICT.
- Understand the current SICT capability maturity level.
- Systematically develop and manage the SICT capability building blocks.

Assess and manage SICT progress over time

b) Defining the Scope and Goal

Firstly, the organization must define the scope of its SICT effort. As a prerequisite, the organization should identify how it views sustainability and its own aspirations. Typically, organizational goals involve one or more of the following:

- Develop significant capabilities and a reputation for environmental leadership.
- Keep pace with industry or stakeholder expectations.
- Meet minimum compliance requirements and reap readily available benefits.

Secondly, the organization must define the goals of its SICT effort. It is important to be clear on the organization's business objectives and the role of SICT in enabling those Objectives.

c) Capability Maturity Levels

The framework defines a five-level maturity curve for identifying and developing SICT capabilities:

d) Initial

SICT is ad hoc; there is little understanding of the subject and few or no related policies. Accountabilities for SICT are not defined and SICT is not considered in the system's life 'cycle.

e) Basic

There is a limited SICT strategy with associated execution plans. It is largely reactive and lack consistency. There is an increasing awareness of the subject, but accountability is not clearly established. Some policies might exist but are adopted inconsistently.

f) Intermediate

A SICT strategy exists with associated plans and priorities. The organization has developed capabilities and skills and encourages individuals to contribute to sustainability programs. The organization includes SICT across the full system's life cycle, and it tracks targets and metrics on an individual project basis.

g) Advanced

Sustainability is a core component of the IT and business-planning life cycles: IT and business jointly drive programs and progress. The organization recognizes SICT as a significant contributor to its sustainability strategy. It aligns business and SICT metrics to achieve success across the enterprise. It also designs policies to enable the achievement of best practices.

h) Optimizing

The organization employs SICT practices across the extended enterprise to include customers, suppliers and partners. The industry recognizes the organization as a sustainability leader and uses its SICT practices to drive industry standards. The organization recognizes SICT as a key factor in driving sustainability as a competitive differentiator.

This maturity curve serves two important purposes. Firstly, it is the basis of an assessment process that helps to determine the current maturity level. Secondly, it provides a view of the growth path by identifying the next set of capabilities an organization should develop to drive greater business vary from SICT.

i) SICT Capability Building Blocks

Whilst it is useful to understand the broad path to increasing maturity, it is more important to assess an organization's specific capabilities related to SICT. The SICT framework consists of the following four categories:

- Strategy and planning, which includes the specific objectives of SICT and its alignment with the organization's overall sustainability strategy, objectives, and goals
- Process management, which includes the sourcing, operation, and disposal of ICT systems, as well as the provision of systems based on sustainability objectives and the reporting of performance

j) Assessing and Managing SICT Progress

With the initial assessment complete, organizations will have a clear view of current capability and key areas for action and improvement. However, to further develop SICT capability, the organization should assess and manage SICT progress over time by using the assessment results to achieve the following goals:

k) Develop a roadmap and action plan.

Add a yearly follow-up assessment to the overall IT management process to measure over time both progress and the value delivered from adopting SICT.

Agreeing on stakeholder ownership for each priority area is critical to developing both short-term and long-term action plans for improvement. The assessment results can be used to prioritize the opportunities for quick wins - that is, those capabilities that have smaller gaps between current and desired maturity and those that are recognized as more important but might have a bigger gap to bridge.

9.9 LET US SUM UP

- Introduces green and sustainable informatics and metrics as a tool to monitor and drive sustainable behavior.
- Describe the multilevel nature of green metrics with examples at the regional, organizational, functional, product/ service, and individual levels.
- Provide a overview of the metrics and standards for greenhouse gas reporting, life cycle impact analysis, energy efficiency within data centers and assessing the maturity of sustainable information and communication technology (SCIT) within organizations.
- Illustrates applications of the Natural Step Framework to develop a sustainable strategy for local government, with relevant KPIs.

9.10 LIST OF REFERENCES

- Harnessing Green IT, authored by San Murugesan and G.R. Gangadharan
- The Green Computing Book Tackling Energy Efficiency At large Scale, Wu-chun Feng
- Sustainable ICTs and Management Systems for Green Computing, Wen-Chen Hu

9.11 WEBLIOGRAPHY

- <https://www.lancaster.ac.uk/data-science-of-the-natural-environment/blogs/green-computing-a-contribution-to-save-the-environment#:~:text=Green%20computing%20is%20the%20environmentally,that%20reduces%20their%20environmental%20impact.>
- <https://www.techopedia.com/definition/14753/green-computing>
- <https://www.intechopen.com/chapters/75768>
- <https://datascience.foundation/datatalk/green-computing-the-future-of-computing>

9.12 UNIT END EXERCISES

1. Explain the hierarchy of sustainability models.
2. What is the role of sustainability frameworks, principles, and tools?
3. What are multi-level models?
4. What are the key metrics for data center energy efficiency?
5. What is LCA? Explain the four stages of LCA.
6. Discuss the maturity of SICT capabilities.

ENTERPRISE GREEN IT READINESS

Unit Structure

- 10.0 Objectives
- 10.1 Introduction
- 10.2 Readiness and Capability
- 10.3 Development of the G-Readiness Framework
- 10.4 Measuring an Organization's G-Readiness
- 10.5 Let us Sum Up
- 10.6 List of References
- 10.7 Bibliography
- 10.8 Unit End Exercises

10.0 OBJECTIVES

- Green IT.-readiness, or *a-readiness*, framework
- Defines the homological structure of the framework
- Offers mechanisms to explain the organizational capability in sustainable management of the IT infrastructure

10.1 INTRODUCTION

The pairing of information technology (IT) and the environment is now referred as Green IT. The emerging notion of green IT raises questions such as how green IT is and how IT can enable a business's green strategy. In the context of 'green' and 'IT', a few perspectives have emerged. To refer to the preservation of environment, whilst many use the terms green, some use the terms sustainability and eco-sustainability. Likewise, although the terms IT, information systems (IS) and information and communications technologies (ICTs) are used interchangeably, sometimes these terms are used to refer to different things. The concept of IT is complex. There is a difference between the technical IT infrastructure and the IT human and managerial capability. The IT technical infrastructure is commonly defined as a pyramid of four layers: the physical infrastructure (e.g., cooling, ventilating and power delivery), IT network and communications technologies, shared services (e.g.,

enterprise-wide databases and electronic data interchange (EDI)) and business applications that utilize the shared infrastructure.

The IT human capability pertains to 'the experiences, competencies, commitments, values and norms of the IT personnel delivering the IT products and services'. The managerial capability comprises the management of all IT activities, including strategic foresight concerning changes in the business, IT, and wider environment. It recommends three approaches - the tactical incremental, strategic, and deep green approaches. Whilst the tactical incremental approach comprises preserving existing IT infrastructure and policies and incorporating simple measures to achieve moderate green goals, the strategic approach involves conducting an environmentally oriented audit of a company's IT infrastructure use and developing a comprehensive green plan and initiatives to address broader green goals. The deep green approach expands the measures of the strategic approach to fundamentally redesign the IT infrastructure in a way that neutralizes greenhouse gas emissions.

10.2 READINESS AND CAPABILITY

Although the term readiness has been used to explain organizational change, IS business process re-engineering (BPR) and innovation implementations, it has also become very popular in the e-commerce and e-government areas.

- Readiness as a precursor condition for the implementation of change, IS or digitization innovations and .
- Readiness as an indicator of the agility of a business and a capability that needs constant building, rebuilding, and upgrading.

It can also include functional skills and cultural perceptions to manage change and innovation. Capabilities emerge over time through complex interactions among tangible and intangible resources. Overall, though, capabilities could be input, transformational and output based. The capabilities refer to firms' physical, capital and human resources. In particular human resource capabilities include the training, experience, judgement and managers and workers. Transformational capabilities transform inputs into outputs, and include innovation to generate new processes, products, and services as well as organizational culture, learning and adaptation. Output capabilities refer to firms' tangible products and services and intangible output.

Extending these views to green IT, capabilities represent the tangible and intangible assets, resources, and processes by which firms deploy eco-sustainability considerations in building and managing their IT infrastructure. Thus, a firm's capability to green its IT can be understood by looking at the permeation of eco-sustainability in a company's IT and IT department's input, transformation, and output. In terms of output capability, there are a number of technologies that are considered to be

green and there is an increasing push for green data centers. Thus, green IT covers not only technical considerations but softer practices as well.

10.3 DEVELOPMENT OF THE G-READINESS FRAMEWORK

G-readiness refers to the maturity of environmental considerations, whether they're part of a coherent set of IT management policies, they have been adopted into IT human and managerial operational practices or they have been built into concrete IS. It demonstrates the comparative levels of green IT development among enterprises and serves as a benchmark for understanding IT's contribution to support an enterprise's environmental sustainability strategy. G-readiness consists of five components: attitude, policy, practice, technology and governance. The attitude dimension of G-readiness is an element of the IT human infrastructure and represents the green IT input capability. The policy and Governance components are elements of IT managerial capability and represent its transformational capability. The technology and practice components are elements of the IT technology infrastructure and represent its output capability. G-readiness is an organization's capability as demonstrated through the combination of green IT attitude, policy, practice, technology and governance in applying environmental criteria to its IT technical infrastructure as well as within its IT human infrastructure and management across the key areas of IT sourcing, operations and disposal to reduce IT business process.

Proposition 1:

Firms with IT people and managers who believe that environmental issues are important are more likely to have a relatively well-developed green IT policy, and a relatively matured green IT governance mechanism.

Green IT Policy

Organizations should develop their green IT policy aligned with its overall environmental policy and initiatives. Green IT policy encompasses the frameworks the organization puts in place to apply environmental criteria in IT-related activities. It defines the extent to which green issues are encapsulated in an organization's procedures guiding the sourcing, use and disposal of the IT technical infrastructure, the activities of the IT human infrastructure and the use of IT in the wider enterprise. The maturity of green IT policy reflects whether environmental considerations are systematically permeating the IT activity value. Policy captures an organization's intent to green IT. However, not all policies are expected to be smoothly implemented, nor are all practices expected to be policy led. Thus, we arrive at the following propositions:

Proposition 2 :

Firms with a relatively well-developed green IT policy are more likely to have greener technical IT infrastructure.

Proposition 3 :

Green IT practices are more likely to flourish in firms _the.it have a relatively well-developed green IT policy.

Green IT Governance

Green IT governance is the operating model that defines the administration of green IT initiatives and is closely related to the policy construct. Roles, responsibilities, accountability, and control for green IT initiatives need to be clearly established. Businesses should decide whether the responsibility for green IT initiatives should be delegated to the CIO or the environmental sustainability manager. Green IT governance also includes allocating budget and other resources to green IT initiatives and defining metrics for assessing the impacts of green IT initiatives. Indeed, governance capabilities will require standard administrative processes for developing green IT initiatives to be put in place. This leads to the following two propositions :

Proposition 4 :

Green IT practices are more likely to flourish in firms with a relatively well-developed green IT governance mechanism.

Proposition 5:

Firms with relatively well-developed green IT governance are more likely to have a high level of green IT technical infrastructure

Green IT Practice

Green IT practice pertains to the actual application and realization of eco sustainability considerations in IT infrastructure sourcing, operation and disposal, Organizations are likely to vary in the practice of analyzing the green track record of IT hardware, software and service providers. They are also likely to vary in their practice in operating the IT and network-critical physical infrastructure in data centers and beyond data centers throughout the organizations in an eco-friendly manner. Green IT practices would contribute positively to greening the IT technical infrastructure.

Proposition-6:

Firms that apply eco-considerations in IT infrastructure decisions are more likely to have a green IT technical infrastructure.

Green IT Technology .

The technological dimension refers to technologies and IS for

- reducing the energy consumption of powering and cooling col12Qrate IT assets
- optimizing the energy efficiency of tl:le IT technical infrastructure
- reducing IT-induced greenhouse gas emissions

- supplanting carbon-emitting business practices and analyzing a business's total environmental footprint

The SAP Environmental Compliance application is designed to help 'organizations ensure compliance with environmental laws and policies and reduce associated costs, efforts and risks on (the) plant and . corporate level. It streamlines all environmental processes by seamless integration with operations control data, production control systems, and components from SAP software for environment, health and safety, enterprise asset management, materials management, the SAP Manufacturing Integration and Intelligence (SAP MII) application, business intelligence and knowledge management'.

10.4 MEASURING AN ORGANIZATION'S G-READINESS

To illustrate the utility of G-readiness in practice, in this section we present a brief case study of a green IT consultancy firm that works with client organizations to determine their current 'G-readiness' capability, with a view to assisting them in improving this. Data were collected via semi structured interviews with the firm's two principals and by document analysis, including green IT consultancy reports developed by the firm and information published via the company Web site.

a) Consultancy Service

Connection Research, Australia is a market research firm specializing in sustainability issues and, sustainability issues for IT. The firm has two principles, two employees and approximately six part time consultants. The firm's key customers include trade associations, government agencies and medium-large organization. Connection Research provides market data to clients collected via direct methods such as surveys and through comparing such data with data from the Australian Bureau of statistics. Assessing the G-readiness of organizations is one of connection Research's key areas of business.

To develop its framework, Connection Research 'initially administered a survey based on the framework to 500 organizations around the world and asked them to rate themselves on a 1-5- point scale. Since the survey was first administered. It has been repeated •by Connection Research in Australia, England, India, and the United States. The framework and survey instrument are publicly available to encourage as much use as possible and, thus, build comprehensive database. The firm then uses the results of the survey to benchmark its clients against other organizations. Essentially an organization completes the survey and compares itself across each dimension to organizations across different industry sectors and also different countries.

Based on the client firm's initial result in the survey, Connection Research provides additional services to client organizations. There are two versions of these services. The first is a 'Green. IT Report Card' of approximately four pages, which provides a breakdown of the _ client's performance across the five dimensions and pillars of green IT compared to other organizations in its industry and at the national total. The report card includes a brief reflection on the client's performance across each area and offers broad suggestions of how performance could be improved. The second set of services is the 'Readiness Index'. This product involves a more comprehensive consultancy process where more individualized services are offered to the client company. These services are offered in greater detail and, generally via face-to -face methods. The product provides customized advice to the client about its objectives for green IT and explains how this relates to each of the five dimensions of green IT. This consultancy service .can focus on calculating the index across all five dimensions, or it can focus - on one dimension. Typically, the service results in a report, prepared for the individual client, reflecting on their performance in terms of industry and national averages.

Connection Research also provides green IT training services to client organizations based on the G-readiness framework, Report Card and Readiness Index.

b) Calculating the G-Readiness Index via a Survey Instrument

A simple application of G -readiness is its use by practitioners to measure the current green IT capabilities of an organization and to identify areas that need improvement. Such an assessment is inherently subjective. If it is done by a group of managers for a single organization, .it first :requires developing a shared understanding of G-readiness items. In this study, measurement items for the five G-readiness constructs - attitude, policy, governance, practice and technology were generated via a literature review. These items formed an initial instrument to survey CIOs of large organizations m Australia, New Zealand, and the United States to test the validity and reliability of the developed model and instrument.

10.5 LET US SUM UP

- Presents a green IT.-readiness, or a-readiness, framework, to display the input, transformational and output capabilities in greening IT. .
- Defines the homological structure of the framework and offers a series of propositions linking the G-readiness dimensions.
- Offers mechanisms to explain the organizational capability in sustainable management of the IT infrastructure and in the IT department's role to promote enterprise-wide sustainability.

10.6 LIST OF REFERENCES

- Harnessing Green IT, authored by San Murugesan and G.R. Gangadharan
- The Green Computing Book Tackling Energy Efficiency At large Scale, Wu-chun Feng
- Sustainable ICTs and Management Systems for Green Computing, Wen-Chen Hu

10.7 WEBLIOGRAPHY

- <https://www.lancaster.ac.uk/data-science-of-the-natural-environment/blogs/green-computing-a-contribution-to-save-the-environment#:~:text=Green%20computing%20is%20the%20environmentally,that%20reduces%20their%20environmental%20impact.>
- <https://datascience.foundation/datatalk/green-computing-the-future-of-computing>

10.8 UNIT END EXERCISES

1. Discuss different strategies and approaches for green IT and their managerial implications.
2. Describe the strengths and weaknesses of the G-readiness framework.
3. For your study, identify an organization. Using the framework suggested in this chapter, analyze it to assess its G-readiness, and compare your findings with the examples discussed in this chapter or presented elsewhere.

SUSTAINABLE IT SERVICES: CREATING A FRAMEWORK FOR SERVICE INNOVATION

Unit Structure

- 11.0 Objectives
- 11.1 Introduction
- 11.2 Factors Driving the Development of Sustainable IT
- 11.3 Sustainable IT Services (SITS)
- 11.4 SITS Strategic Framework
- 11.5 Let us Sum Up
- 11.6 List of References
- 11.7 Bibliography
- 11.8 Unit End Exercises

11.0 OBJECTIVES

- Reviews the current state of the development of sustainable IT Services
 - Discusses the factors driving the development of green and sustainable IT
 - Examine the sustainability dimensions of IT
-

11.1 INTRODUCTION

Over the past two decades the terms green computing and . green IT has come into common usage to define strategies, processes and practices focused in reducing the environmental impacts of information technology. From its inception with the US Environmental Protection Agency's (EPA) 1991 Green Lights program for energy-efficient lighting and the 1992-Energy Star program for computers and monitors, green IT has evolved significantly. One of the objectives of green IT is to reduce energy consumption and thus the cost of computing. More recently, the linkage of energy use. with _carbon generation and elevated concerns about IT s impact on the environment and potential contribution to climate change has enabled reducing energy use to be positioned as 'green'. The increasing energy costs of data centers comprise a primary driver behind the current manifestation of the green IT movement. Enterprise-scale data centers can easily account for half of a technology company's energy bill and corresponding carbon footprint. The rapid uptake of cloud-based computing models and rising energy costs · has exacerbated this trend with the costs for data center computing and cooling expected to rise as

organizations become more data intensive. Since cost reduction has a direct impact on business value, green IT at its core will likely remain focused on cost reduction for the foreseeable future. In addition to energy costs, other drivers for the continued development of green IT are

- The continuing rapid growth of the Internet use,
- Increasing data center power
- Density,
- Increasing cooling requirements,
- Restrictions on energy supply and access
- The growing awareness of the environmental impact of computing
- Regulatory compliance

The relationship between energy use, carbon footprints and associated costs as the primary drivers of the first wave of green IT is direct and simple: Energy cost reduction also reduces the organization's environmental footprint. Green IT issues impact energy efficiency, IT budgets, legal compliance and risk reduction, data center design, server virtualization, supply chain optimization, product design, manufacturing, e-waste, air pollution, water pollution, hazardous materials, recycling and reuse. To gain from customer or brand equity, businesses should make their green initiatives and credentials visible to customers and other stakeholders, and recent research suggests that firms should be developing products and services that can be positioned on their environmental and social responsibility credentials.

The second wave of green IT is service based and will enable firms to embrace the development of sustainability solutions as a business strategy. Sustainable IT services (SITS) move beyond the first-wave green IT issues to focus on the long-term importance of an organization's IT unit to the firm, its customers and society at large. This will be accomplished by aligning IT with corporate environmental and social responsibility strategy. The IT unit will become a key enabler of corporate sustainability (CS) strategy through the development and implementation of an innovation platform for new IT service-based business models. SITS is service innovation that redefines markets and transforms the competitive landscape and forces companies to change the way they think about opportunities and value creation. SITS-based solutions have the potential to revolutionize business models by turning environmental and social responsibility

problems into business opportunities for the firm, its customers and society at large. As green IT and SITS are complementary and sequential concepts on the same developmental roadmap, it is useful to conceptualize them as domains within the larger sustainable IT construct. Sustainable IT is the overarching concept that encompasses green IT as its first wave and SITS as its second wave. The sustainable IT construct extends from its initial focus on energy use and legal compliance through

its higher order developments in terms of new processes and new business models that derive from a SITS' innovation platform.

11.2 FACTORS DRIVING THE DEVELOPMENT OF SUSTAINABLE IT

The development of sustainable IT, and its green IT and SITS domains, is the direct result of several market forces and enabling factors. For green IT, the primary drivers are energy efficiency and legal compliance. Currently, political and social forces are focusing on the environmental costs associated with IT, especially the potential impact on climate change. Although it is likely that a cap-and-tax or carbon-trading regime will not be fully implemented in the United States, the EPA, some state and local governments, nongovernmental organizations (NGOs) and consumer groups are continuing to pressure companies in terms of their environmental strategies, especially their carbon footprints.

The development of SITS recognizes the migration from product-oriented green IT to service development as IT organizations transition from internally focused cost and compliance issues to externally focused sustainability-driven innovation. SITS brings the IT organization into alignment with the corporation's CS strategy. SITS involves the integration of software, hardware, telecommunications networks, data, maintenance, technical support, customer service and consulting that is necessary to design, deploy, operate and maintain SITS. Increasingly, SITS will evolve to cloud-based solutions. This is already happening as firms such as IBM, HP and Oracle are launching SITS-based services that were initially developed for internal use. SITS is a force for disruptive change for IT organizations. It will require IT managers to shift their priorities from internal, efficiency to customer and societal-focused value creation.

a) The Sustainability Dimensions of IT

IT performance requirements are precise and unforgiving. The technology must be capable to meet demand, availability and capacity management requirements. IT systems must work as expected and need to be ready all the time. IT must be cost effective, deliver customer and business value and generate positive return on investment (ROI). Sustainable IT must enable business to meet these requirements or it will not be pursued. Sustainable IT is 'everything you need to keep IT going indefinitely'.

b) Corporate Sustainability, Social Responsibility, and IT

CS and corporate social responsibility (CSR) are very similar constructs with areas of overlap and points of difference. Both have similar conceptualizations of the economic, environmental, and social responsibility dimensions of sustainability. CS researchers are more environmental purists and take systems view where the economy and society are part of a greater ecological system. They

focus on the 'intrinsic value' of the environment for its own sake independent of the potential benefit for humans or business. The emphasis on 'climate change' is a good example of the CS orientation. CSR researchers view the underlying dimensions individually and look at the environment as a source of benefits for society, or 'use value'. Environmental and social impacts should be minimized, mitigated, avoided or resolved favorably, but not necessarily proscribed. Both approaches embrace stakeholder interests and encourage collaboration.

To be recognized as socially responsible, firms must address all of these issues in the course of their businesses and, hence, CS and CSR are converging and essentially equivalent. We will, therefore, use this broadly integrated CS concept as the underlying discipline for the discussion of sustainable IT and conceptualization of SITS.

Many IT professionals tend to define green IT from narrow economic and ecological perspectives. Mostly they ignore the broader CS social responsibility dimensions or engage to some degree with stakeholders as necessary. The terms green IT and green computing, by evoking this narrow perspective, obscure the true potential of information technology to become the innovation engine for sustainability-focused services. To improve IT's alignment with overall CS efforts, the emphasis needs to be on service innovation, especially for solutions that address ecological, ethical-social, and philanthropic dimensions of sustainability as a means for achieving economic success. This second wave of sustainable IT reflects a shift in customer requirements from the tangible cost benefits of IT as product to the more intangible benefits of sustainable IT as a service for implementing socially responsible business models such as creating services to support the economical provision of sustainable water management in less developed countries.

IT organizations have been able to develop innovative green IT solutions. The problems are well defined and the benefits, especially in terms of cost reduction, are tangible and achievable in relatively short time. IT organizations have a high degree of familiarity with the technology and its applications, implementation, and operation which are internal to the organization. For the broader CS-focused IT services, the organization, however, is likely to be poorly prepared to understand the full range of social responsibility issues, much less design services to address them. The existing communications gap between the IT organization and its customers, both internal and external, and the questionable alignment between IT and business strategy will be challenged to new levels as IT managers try to ordain the complexities of CS. Service innovation is the future of corporations, and information technology will provide the infrastructure for its development and delivery. SITS will need to become a core competence of sustainable IT organizations.

With the on-going shift in the IT industry from product-oriented to service-oriented business models such as cloud-based IT services, the underlying technology for developing SITS is at hand IT organizations need to develop the business case for its implementation. We define SITS from a total societal value perspective as aggregate value available to society from the systematic integration and alignment of the individual IT service components for the purpose of creating superior societal value

a) Developing a Service-Dominant Logic

Service-Centered view of business that has the following requirements:

- Identify or develop service-oriented core competencies, knowledge and skills that represent potential competitive advantage.
- Identify potential customers that can benefit from the competencies.
- Engage in collaborative relationships that involve customers in developing customized, competitively compelling value propositions to meet specific needs.
- Monitor service-outcomes to improve customer collaboration and firm performance.

To implement such a service-centered approach, it is important for IT managers to understand the conceptual differences between products and services, the types of value created by each and how they are delivered. It involves understanding the distinction between the goods-dominant (G-D) logic paradigm and the service-dominant (S-D) logic paradigm. G-D logic views 'services' from a product perspective as 'add-ons' or intangible products that may be offered on an after-sale basis. In this view, services are designed and conceptualized as outputs, essentially intangible products, that 'add value' that is created and then delivered to the consumer for compensation. It is an 'arms-length' transaction. Therefore, G-D logic is based on a value-in-exchange conceptualization. Conversely, S-D logic envisions services to be product independent, although goods can play a role in service provisioning. The value proposition is service centered. Service value is always co-created with customers. It is based on the value-in-use concept. Co-creation involves the process of proposing value, the acceptance of the proposal and the realization of the proposal by two service systems.

Since the customer is actively engaged with the 'provider' in the co-creation of value, the outcome of the service experience should, by definition, be more customer oriented and satisfying. The distinctions between G-D logic and S-D logic are useful for assessing the evolving nature of sustainable IT from a product

orientation to a service orientation. The first wave of sustainable IT with its focus on green computing is an artefact of G-D logic that focuses on the product dimensions of business value to reduce costs and increase benefit (lower carbon footprint) without directly engaging customers or other service systems on a co-creation of value basis. Green IT by nature is a traditional design-build-sell approach that creates value-in-exchange. SITS, in contrast, must engage both internal and external customers and other members of its business ecosystem be successful. Its foundation is collaboration that facilitates the co-creation value propositions on a value-in-use basis. In short, for sustainable IT to be viable in the long term, it must become a true service that can create sustainable: customer value.

b) Business Value, Customer Value and Societal Value

IT operations have a demonstrable impact on the bottom line of the corporation. Business value is defined as the overall value that results from IT products and services that is realized by the corporation. To optimize business value, IT spending must be aligned with corporate business strategy. Increased revenues or lower costs that can improve ROI are evidence that business value is being created. Successful investments in innovative IT solutions, such as green IT, are primary drivers of business value. Although business value ultimately depends on creating value for the customer, business value objectives are primarily about generating returns for the corporation. Therefore, business value tends to focus on short-term, cost-based solutions that are relatively quick and easily quantifiable. Green IT fits well in this paradigm with metrics that include energy cost reductions, headcount reduction, productivity increases, uptime, risk avoidance and cost avoidance. This internal focus can overlook the long-term best interests of the customer, society, the business and ultimately the IT organization as well.

Alternatively, the SITS concept requires a continuous engagement between two service systems to co-create value to meet economic, ecological, and societal requirements. This situation describes the S-D logic co-creation of value-in-use for an IT service. However, in markets characterized by technological innovation customer value can change rapidly, whether that value is co-created. Some customers may be willing to look at their long-term needs in a societal context, but for most IT customers cost and performance are the primary value drivers. The definition of social value is characterized by subjective interpretations of value. If the organization is dealing with renewable energy systems water and wastewater systems, pollution, recycling and other tangible issues, it is likely that well-defined parameters and strategies can be developed. If the issues are world hunger, social justice, environmental justice, income inequality or neglected social services, the opportunity for highly subjective interpretations of value

c) SITS as Service Science

Service science's purpose is to develop a discipline that supports the application of scientific rigor to the practice of services and the development of innovative business models. This is particularly important for IT services since the transition of IT organizations from product orientation to service orientation has just begun. Historically, most technology related research has been focused on product development and manufacturing innovation. The producer was separated by space and time from the consumer as products were produced away from the market. Product standardization supported the goals of efficiency and profit maximization. Conversely, service involves the interaction of producers, consumers and value network partners to co-create value through the collaboration of service systems. This is exactly the transition that is taking place as sustainable IT moves from product-centric green IT to service-centric SITS. IT managers are key players in this transition. They need to think of service that can be manifested in two ways:

- IT products and infrastructure as provisioning agents for the development and delivery of high-value services and
- pure play services

If one conceives of value being created in terms of a flow of services delivered from a common platform (such as cloud-based services) rather than one-off sales of tangible products with add-on minor services, then it is possible to conceive of a scalable service system. This is a complementary relationship with IT services representing recurring revenue opportunities with high profit potential that leverage the IT system platform. As hardware becomes commoditized and software can be hosted anywhere, pure play services can be developed. This is happening with smart phone wireless applications and with third party-hosted software. As the IT organization moves from value in exchange product models to value-in-use service models, the importance of intangible value becomes apparent. Intangible value necessitates an understanding of the knowledge, emotions and experience factors that are associated with the design and delivery of services. Service is knowledge intensive, and the co-creation of value is predicated on a two-way exchange of knowledge. In addition, these collaborations with customers have high emotional valence that sets expectations for the service.

11.4 SITS STRATEGIC FRAMEWORK

Although the need for the development of strategies for the development and delivery of SITS has been apparent for many years, there is no extant body of literature on strategies or best practices. The primary driver of sustainable IT, both green IT and SITS, is CS, especially as it applies to a

firm's impact on the economy, environment and society at large. Therefore, IT managers will need to work across the functional areas of the organization to ensure an integrated organization-wide SITS strategy that aligns with overall CS strategy. SITS development and deployment will likely to be slower than that of green IT. It will take time for firms to understand a market that is early in its formation. CS requirements and outcomes are harder to measure customer and societal value generated. Due to the front-end costs and lack of visibility about returns, it may be difficult for small companies to offer SITS solutions, although there will eventually be start-ups or spin-offs with SITS business models. At present, it is the larger IT companies in which SITS is starting to get traction. They are packaging applications developed for internal use and marketing them to their customers.

The SITS Value Curve

Firms are changing their policies and practices to lower costs and minimize their environmental impact, but such efforts are not usually integrated with CS projects have focused on customer safety, legal and regulation compliance, product design, sustainable processes, stakeholder engagement, strategic social responsibility, sustainable organizational culture, and efforts to protect the brand. What is needed is a value roadmap with sequential steps that results in an integrated approach to develop an innovation platform for sustainable IT. One approach is for companies to view their current activities against a CS value curve to obtain a perspective on what the next steps should be: The SITS Value Curve, adapted from IBM Business Value Institute's CSR Value Curve, depicts a path from cost-risk mitigation efforts to the development of an integrated core strategy.

It conceptualizes the value migration path that starts with energy efficiency, cost savings and compliance activities and migrates through subsequent steps to the strategic apex of the sustainable IT innovation platform. When organizations have successfully met the goals of each step at the lower levels of the curve, they can continue their journey upward towards more strategic activities and higher levels of value creation in a sequential fashion. Our conceptualization recognizes green IT initiatives as a starting point along a value migration path that culminates with a SITS-driven sustainable IT innovation platform that supports overall CS efforts. The following strategy development milestones depict the firm's maturation from focusing on costs and compliance at the low end to embracing the organizational change necessary to create a high-value SITS-based growth platform for the implementation of CS innovation. The major milestones along the curve are as follows:

- **Energy efficiency and cost savings** : The starting point for SITS value creation is efficiency. The initial step along the curve recognizes the linkage between energy use by IT products and operations and overall carbon dioxide generations. The emphasis is on reducing the energy costs and carbon footprints associated with computing, especially in data centers.

- **Legal, compliance, standards, and risk management** : CS-associated requirements for product and service design, production, operations, and distribution are well established in Europe and increasingly in North America. These regulations and standards are spreading worldwide as they are adopted by national governments, manufacturers, and marketers.
- **Sustainable product design** : This value milestone includes clean technology, both 'clean tech' and 'green tech'. Clean tech includes renewable energy, transportation, water and waste water, air pollution, materials, smart manufacturing, agriculture, recycling, and waste. It also includes design for environment (DfE), design for recycling (DfR), e-waste minimization and cloud-based service applications (mostly green IT).
- **Sustainable processes** : Includes IT systems, energy-efficient operations, the green supply chain, product development, engineering, manufacturing, operations, distribution, and marketing. All these functions rely on IT, and their processes can affect the firm's environmental and social impact.
- **Sustainable organizational culture**: Creating a common culture based on sustainability empowers employees to become more aware of issue opportunities and the actions required for achieving a desired result. It involves setting clear goals and Objectives and developing a code of conduct to guide business activity.
- **SITS Innovation Platform**: SITS integrates the value innovation potential of all the sustainable IT dimensions into a platform for market growth'. SITS creates growth by identifying new market opportunities collaborating to co-create value and creating innovative solutions to drive business performance.

Integrating Sustainable IT and Business Strategy

- **Understand the competitive context** : Competitive context refers to the dynamics of the industry and its key players. It refers to the quantity and quality of business resources, the rules that govern competition, competitors' capabilities, the size and sophistication of demand, the firm's value chain members' availability and capability and the key stakeholders' characteristics and capabilities. The IT organization needs to understand the relationships, how and where sustainability issues will impact them and the potential for developing partnerships and alliances.
- **Choose sustainability issues to address** : IT managers must choose high-impact sustainability issues that intersect with its key business initiatives. Ideally the choice will align with CS initiatives to create shared value that provides a meaningful benefit for the environment while providing value for the business. To ensure that the IT organization can meet its business goals, sustainability issues that

are in the firm's strategic interest". should be given priority. General social responsibility issues might be important to society but are not affected by or affect the IT organization. Value chain sustainability issues that directly impact the IT organization might present a business opportunity. Competitive context sustainability issues that can significantly affect the competitiveness of the IT organization and the firm should be given the highest priority. A method for evaluating and prioritizing sustainability issues needs to be developed.

- **Create a sustainability agenda** : The IT organization should engage the firm's stakeholders to identify sustainability opportunities. Porter and Kramer advise managers to choose between responsive sustainability and strategic sustainability for agenda development. Responsive sustainability appeals to good citizenship by responding to stakeholder demands. However, responsive sustainability often deals with generic issues of little strategic value. It does not impact the competitive context. Strategic sustainability raises the bar to focus on issues that directly impact the competitive context and transform value chain activities to enhance sustainability whilst supporting business strategy. However, some stakeholders may not appreciate the firm acting in its own best interests.
- **Create a sustainability dimension for all value propositions** : To address sustainability strategy, the IT organization will need to ensure that sustainability principles are at the heart of its value proposition. IBM has been successful in organizing its sustainability efforts around the concept of a 'smarter planet'. This positioning provides a unifying theme for the value proposition.

11.5 LET US SUM UP

- Examines the sustainability dimensions of IT: service, temporal organizational, economic, environmental and societal sustainability, cost,
- Discusses the corporate sustainability, social responsibility and IT.
- Summarizes the service-dominant logic with business, customer societal values
- Discusses the steps to integrate sustainable IT with business strategy.
- Explores the best practices from major IT to non-IT companies.
- Recommends a set of Guidelines for sustainable IT services development.

11.6 LIST OF REFERENCES

- Harnessing Green IT, authored by San Murugesan and G.R. Gangadharan
- The Green Computing Book Tackling Energy Efficiency At large Scale, Wu-chun Feng

11.7 WEBLIOGRAPHY

- <https://www.techopedia.com/definition/14753/green-computing>
- <https://www.intechopen.com/chapters/75768>
- <https://datascience.foundation/datatalk/green-computing-the-future-of-computing>

11.8 UNIT END EXERCISES

1. Outline the current state of development of sustainable IT services.
2. What are the driving factors for the development of green and sustainable IT?
3. How are corporate sustainability, social responsibility and IT related to each other?
4. How should companies move from a goods-dominant logic to a service-dominant logic by utilizing business, customer and societal values?
5. How can an organization integrate sustainable IT with its business strategy?
6. What is a sustainable IT services innovation platform? What capabilities should it have?
7. On what will the major SITS oriented applications be focused in five years?
8. How can companies implement sustainable IT services development practices?

GREEN ENTERPRISES AND THE ROLE OF IT

Unit Structure

- 12.0 Objectives
- 12.1 Introduction
- 12.2 Organizational and Enterprise Greening
- 12.3 Information Systems in Greening Enterprises
- 12.4 Greening the Enterprise: IT Usage and Hardware
- 12.5 Inter-organizational Enterprise Activities and Green Issues
- 12.6 Let us Sum Up
- 12.7 List of References
- 12.8 Bibliography
- 12.9 Unit End Exercises

12.0 OBJECTIVES

- Provides an overview of various green initiatives.
- Presents a comprehensive overview of various green enterprise activities and functions and their role with IT.
- Discusses the roles of IT and IS in greening enterprises.

12.1 INTRODUCTION

Enterprises are becoming greener, and IT plays major roles in greening an enterprise. In this chapter we provide an overview of various green initiatives within and between organizations. Based on the value chain perspective and focusing on various organizational functions, and green activities. We then outline organizational aspects of greening which include purchasing and logistics relationship amongst organizations. The general greening strategies, practices (routines), policy and process characteristics across the internal and external value chains set the stage for evaluating and integrating IT considerations. Whether for the greening of physical and usage aspects of actual capital equipment or the application of IT tools and software for greening purposes, various green value chain dimensions, in-bound logistics, purchasing, production, distribution and outbound

logistics, design and marketing, inter-organizational relationships and reverse logistics (closing the loop) are addressed in this chapter.

We identify some practices which organizations can complete with various greening policies such as environmental management systems (EMS), green procurement; design for the environment, life cycle analysis, green marketing and green logistics and transportation, which are all evaluated here. For planning and controlling each of the value chain functions of organization, IT can effectively be used. For example, IT could be harnessed for life cycle assessment (LCA), environmental management system, environmental performance measurement systems and design for environment and even green supplier and customer management. In terms of capital equipment and the tangible physical product, green IT means that purchasing greener technologies, using IT in service settings, manufacturing these systems (if it is an IT company) or managing the end of life of IT equipment all have implications for greening the enterprise. E-commerce business practices, which are heavily dependent on IT, have significant inter organizational implications.

Green IS referring to improving the flow and management of information, whilst green IT refers more to the hardware and other infrastructure that can be better managed and designed from an environmental perspective. It involves the greening of existing enterprises, focusing on the production process, and the promotion of green enterprises in the production of environmental goods and services. Green enterprise development results in greener, environmentally friendly, safer, and more productive workplaces. The business value argument is necessary to help organizations more effectively manage the resources they target towards these efforts. The bottom-line and competitive dimensions of the organization are influenced by green IT and IS.

Elimination or reduction of waste reduces the cost. Consumption of excessive resources such as energy and paper essentially mean there is inefficiency and increased costs. The most common example is less energy usage but IS practices such as the paperless office are also examples. Organizations may realize extra revenue through green IT and IS practices. By-products and former waste products may find alternative uses instead of non-value-adding disposal to landfills. IT products and materials that are returned may be recycled, remanufactured, and resold, potentially as 'green' products. Thus, what may have been viewed as a cost center for organizations may become a revenue generator or profit center. Business continuity has many dimensions but means having the resources to remain in business and deliver products or services to customers, making sure your organization and partners are resilient. If organizations are unsustainable, they may use up their resources, materials may not be as easily available and/ or their costs can dramatically increase. By managing supply chains in a sustainable way, such as by making sure that hazardous materials are not in IT products, it is more likely that long-term continuity will occur. To operate effectively and with little stakeholder conflict, organizations need to develop their 'right to do business'. This value-adding dimension is related to furthering the

reputation and legitimacy of organizations that have green IT and IS practices in place. Extending the life of IT products by donating to charity can also improve an organization's image and reputation. Appropriate justification and management of these systems can be not only environmentally beneficial but also economically and, competitively valuable.

12.2 ORGANIZATIONAL AND ENTERPRISE GREENING THE GREEN ENTERPRISE: A VALUE CHAIN PERSPECTIVE

The aim of a green approach to Value Chains is :

- i) to ensure the sustainable use of natural resources and to increase the share of renewable and recycled resources at the input side of the value chain,
- ii) to maximize material- and energy efficiency at each stage of the process,
- iii) to reduce negative environmental impacts as outputs at all points of the chain.

Green Value Chain Development also refers to the promotion of green mark opportunities, where economic benefits from the use of renewable resources are maximized while environmental harm is minimized. Green value chain interventions also include supporting the creation of an enabling environment for green investment, skills training in green technologies, green entrepreneurship, and business development, and greening the workplace.

Within this systemic perspective of the organization, activities begin with procurement and inbound logistics functions that introduce materials and services into the organizational system. These materials are transported from various vendors. The policies for selection of vendors, including transportation and delivery services, are a central issue for purchasing agents. Thus, the selection of material, services and suppliers may become critical for the purchasing function to help guarantee the environmental performance of their supply chain and materials. Relationships with suppliers including environmentally oriented selection, development and supplier management are needed for effectively greening the procurement function. These materials are then stored, or the services are utilized, and may be managed under the auspices of the purchasing function.

An important strategic activity for an enterprise is the design of products and processes. Each of the major functions will be profoundly influenced by the design of the product and the process. This influence should be managed with the cross-functional inclusion of various functions and suppliers in the design of processes and products. Designs should now address new ecological themes as life cycle analysis and design for the environment concepts. The production (transformation) function in a

typical manufacturing organization is composed of assembly and fabrication (service organizations would focus on transformation in different ways such as information transformation and locational transformation). Within this function, environmental issues such as closed-loop manufacturing, total quality environmental management, de-manufacturing and source reduction make some form of value-adding contribution, even though some of them also influence other functional areas. Outbound logistics includes such activities as transportation determination, packaging, location analysis, warehousing and inventory management (for finished goods and spare parts). In these areas, the environmental implications from transportation mode selection and the amount and type of packaging are critical environmental dimensions of organizations.

Activities such as using recyclable containers and green fleets and efficiently designing and managing warehousing and inventory play roles in managing an enterprises ecological footprint. Marketing's role is important for activities within this stage of organizational functions. Green marketing has proven to be effective component of green enterprise strategy for many organizations.

Within this reverse channel, the product may have reusable, recyclable, or manufacturable characteristics. The reverse logistics function may feed directly back to an organization's internal supply chain or to an external vendor, starting the cycle again. As each of these supply chain activities consumes energy and generates some waste, reductions in energy usage and waste generation are issues that need to be addressed throughout the supply chain. IS and the movement of information rather than the movement of materials (information substitution) can greatly enhance this reduction.

12.3 INFORMATION SYSTEMS IN GREENING ENTERPRISES

Organizational IS can be separated into four major categories of systems: the transaction processing system (TPS), decision support system (DSS), management information system (MIS) and executive support system (ESS). In addition to these, recently a new category of IS, environmental management information system (EMIS) that helps to manage environmental aspects is emerging. TPS are focused at the operational level of an organization and deal with real-time and very short-term information requirements. MIS and DSS are typically focused at the middle management (tactical-planning) level of an organization, MIS's role is to summarize and aggregate in reporting and communication systems the company's operational information. Typically, this information is aggregated to intermediate time scales ranging from weekly to yearly information. DSSs are typically analytical tools that help middle-level managers make relatively routine decisions and use information from various levels of the organization. ESSs focus on the needs and requirements of upper management and aid their focus on

strategic management of the organization. Information is typically aggregated from the lower-level systems and summarized in longer term perspectives ranging into years and focusing on nonroutine-type decisions. Environmental software to support green IS and IT management range from auditing and managing emissions to analyzing energy and minimizing waste. More recent aspects of software systems include the utilization of Life-Cycle-Analysis (LCA), which goes beyond the organizational boundaries. Design for environment (DfE) databases also play a role, which may integrate supplier and vendor data and processes. The management of hardware, software, and its usage is critical to organizational greening.

A) Environmental Management Information Systems

EMISs include hardware, software, people, procedures, and tasks that manage environmental information and support environmental and other managers in managing environmental issues within and between organizations. We will focus on software and systems that are used to manage an organization's environmental and greening activities. The usage of such systems also extends to intra-enterprise activities. We begin this section with an overview of the software and database dimensions of EMIS

B) Software and Databases

Environmental software programs cover everything from auditing and managing emissions to analyzing energy and minimizing waste. Traditional environmental software as part of IS has focused on internal practices. More recent aspects of software systems include the utilization of LCA which goes beyond organizational boundaries. Design for environment databases also plays a role and helps integrate supplier and vendor data and processes.

C) ERP/EMIS

EMIS have typically been designed as stand-alone, isolated, software tools such as waste administration packages. Integration and adequate co-operation of environmental and ERP functionality require an emphasis in their common base: the physical flow model.

Environmental harm directly connected to production is related to physical flows entering and leaving the production process. Enterprises are obliged to deal with data on energy and raw materials consumption, on discharge and on emissions. From a costing point of view, physical flows are also important, even if they were not accounted for traditionally, because they are imposed gradually to specific costs like taxes, tariffs and levies. The essential units for most EMISs are physical ones: mass units (kg) and energy units (MJ), each measured over a definite unit of time. The time unit depends on the availability of data, and on the scope (instrumental, operational, tactical and , strategic) that is considered. Also, other operational units are often required, like volume, surface, number of items and so on, with appropriate conversions between units also easily

accessible. Of course, the relationship between physical data and environmental data is usually dependent on industry, and appropriate accounting of materials and their environmental implications is needed. Another unique aspect of environmental information and the IS design is the need to consider multiple stakeholders: supply chain partners (i.e. suppliers and customers), regulatory authorities, consumer organizations, insurance companies and the community. EMIS waste and emissions' data would require similar controls and accuracy, as do product and material data in standard ERP systems.

D) Materials and Products Information

Material and product data are available through a bill of material (BOM). Although BOMs and supporting inventory control systems provide some insight into the composition of materials flows, this usually is not sufficient from an environmental point of view. This product and material information is on technical aspects, quality, product specifications and so on. Environmentally relevant data, however, only partially overlap the BOM and material requirements planning (MRP) data. BOM material and product composition can be represented on different levels and in different ways. These levels may include elements, compounds, materials and parts or component assemblies (e.g. bolts or motors).

A BOM usually includes no weight of components or only incomplete information on weight. In addition, components are reported and not materials or substances as needed for a LCA inventory. Supporting processing materials (solvents and lubricants) are also not included in these BOMs. An example is a printed circuit board listed as a component; this would include data on the supplier, type number, dimensions and functionality which can be obtained from

BOM. Yet, physical properties like mass, the presence of hazardous materials and components and the coarse composition may be missing.

Given regulatory policy such as the 'Restrictions on Hazardous Substances' (ROHS), manufacturers are required to know the prohibited substances in their products. BOM represents a single product or assembly, or a group of products, which can aid recognition of any RoHS restricted substances in the final product. A component's testing report is a crucial document in determining where RoHS restricted substances exist within its manufacture.

Clearly, this type of information would be necessary in an environmentally integrated ERP system. Integrating a BOM with a material safety data sheet (MSDS) is one way of furthering both environmental and product information.

E) Process and Workflow Information

Process flow charts are part of ERP systems, especially in ISO 9000-certified organizations. This information may be used to describe the

material and energy flows in an LCA. Work orders may contain a relatively complete data set about materials, components, required resources and the job task order. Using this information, process flow charts can be generated to track production related data for resource. and energy use.

F) Supply Chain Information

Primarily, transportation and logistics-type information could be useful for environmental management reasons. Supplier information can be used to calculate the transportation effort to the company. In addition, the kind of delivery is registered in the conditions of the supplier. Transportation methods within a company can be obtained from the work order. The distribution methods from a company can be obtained from the sales and distribution functionality of an ERP system.

G) Information about Waste

This information includes types and quantity of waste, and the product that generated the waste. Logistics and procurement-scheduling modules may also provide valuable information from an environmental perspective. Transportation planning will identify weights, modes of transport, distances and other information that can be used to evaluate the carbon footprint of the organization's logistics function.

H) Integrating Environmental and LCA Information with ERP

As mentioned, some ERP systems have environmental, health and safety (EH&S) modules included in them. For example, the SAP R/3 system has an integrated EH&S module and provides operational solutions for hazardous substance management, product safety, dangerous good management, and waste management. Yet, the characteristics and linkages necessary for a full corporate analysis linking them to various functional and supply chain decisions have not been advanced.

There have been some structural systems models that show a more extensive linkage between EMIS, LCA and E-RP. The functionality of an ERP system should include various managerial and functional decision levels. Using the managerial decision-making hierarchy to link these two has been recommended. Others have advocated that ERP systems need integration of mass and energy flow models for effective integration. This integration is necessary since, as we shall see in this chapter, business process models are required for systems analysis and design approaches for implementing enterprise-wide IS.

A consortium of German industry and schools has been investigating the development of standards called Publicly Available Specification (PAS) 1025 for the standardization of an exchange format between ERP and EMIS systems. This exchange is eventually meant to be a two-way integration since ERP data are useful for EMIS and EMIS data can be used for ERP. such as in cost-accounting modules

D) Electronic Environmental and Sustainability Reporting

Environmental communication and environmental and sustainability corporate reporting are tools to communicate a company's (environment~) performance, demonstrate its management systems and present its responsibility, among other items. Within the electronic commerce field, the most common language used for communication and reporting standardization is Extensible Markup Language (XML). In fact, many EMS are currently under design to arrive at the corporate report. XML-based documents have a precisely defined structure. XML-based reports are also a suitable tool for Internet and cross-media reporting.

12.4 GREENING THE ENTERPRISE : IT USAGE AND HARDWARE

A driving force in greening IT from a manufacturing perspective is the adoption of environmentally friendly materials, procedures and practices in the design, manufacture, and delivery of IT systems. Electronic products require significant 'clean' requirements. Higher purity requirements mean a greater need for energy and materials. Large quantities of energy, materials and chemicals are consumed during the production phase, not all of which will be contained in the final products. Modern electronics and computing equipment consists of over 1000 materials, including lead and cadmium in computer circuit boards, lead oxide and barium in the monitors. CRTs and mercury in switches and flat screens. Thus, the management of hardware and its usage is critical to organizational greening.

a) Environmental Information Technology Standards

Green IT standards have emerged recently. One of the major efforts towards greening IT product is the Electronic Product Environmental Assessment Tool (EPEAT) that helps identify greener computers and other electronic equipment. It is a powerful tool for enhancing an enterprise's sustainability performance. EPEAT is a ranking system that helps purchasers in the public and private sectors evaluate, compare and select desktop computers, notebooks and monitors based on their environmental attributes. EPEAT evaluates products according to 51 specific criteria into three tiers of environmental performance - Bronze, Silver and Gold. EPEAT is managed and operated by staff contracted from the Green Electronics Council. EPEAT is a three-tiered, point-based system. These criteria cover the entire life cycle of a product, from a reduction of the toxic materials used in production, to the energy it uses whilst in operation, and the recyclability of its materials at the end of life. Other standards such as Energy Star and IEEE Standard 802.3az-2010 influence design. A new standard, IEEE Standard 802.3az-2010, requires a dual-power state for Ethernet connections, high power and lower power, that can cut down the energy used during those idle times. These standards provide valuable guidelines for manufacturers' eco-design (design for the environment) and for meeting and achieving green IT goals.

b) Green Management of Data Centers

A green data center is a repository for the storage, management, and dissemination of data in which the mechanical, lighting, electrical and computer systems are designed for maximum energy efficiency and minimum environmental impact. Internally for enterprises, there are many activities that can help green the usage of IT is one of the most critical energy users within an enterprise with around-the-clock usage is the data center. File servers and centralized data storage devices are predominant. The management of these centers is critical to the overall greening of IT whether the enterprise is service or manufacturing oriented. Data centers in the United States may apply to be certified as green data centers. The most widely used green building rating system is the Leadership in Energy and Environmental Design (LEED). Developed by the U.S. Green Building Council, it is available in several categories. Depending on ratings, data centers may receive a silver, gold or platinum certification. The platinum certification is given to data centers with the highest level of environmentally responsible construction and efficient use of resources.

12.5 INTER - ORGANIZATIONAL ENTERPRISE ACTIVITIES AND GREEN ISSUES

Data centers may also be certified under the National Data Center Energy Efficiency Information Program by Energy Star, part of an initiative by the U.S. Environmental Protection Agency and the U.S. Department of Energy. The program certifies buildings and consumer products for energy efficiency.

One organization's waste (by-product) may be another organization's raw material. Environmental information plays important roles in the planning, design, construction, management and retrofit of eco-industrial parks, in industrial symbiosis and in by-product exchange network synthesis. One of the key approaches to improve environmental performance and production management efficiency is to integrate members' environmental information into a park management system. A recent survey of 17 of these systems shows that their major purposes (functionalities) are

- (i) opportunity identification,
- (ii) opportunity assessment,
- (iii) barrier removal,
- (iv) commercialization and adaptive management and
- (v) documentation, review and publication.

Inter-organizational environmental IS is also recommended for establishment within industrial parks and zones to provide integrated and reliable data. Such data would include eco-industrial park member surveys, detailed information about inputs and outputs of materials, environmental monitoring and other data. Life cycle analysis and material flow analysis tools and information will help eco-industrial parks guide eco-industrial development.

12.6 LET US SUM UP

- Provides an overview of various green initiatives within and between organizations
- Presents a comprehensive overview of various green enterprise activities and functions and their role with IT.
- Discusses the roles of IT and IS in greening enterprises.

12.7 LIST OF REFERENCES

- Harnessing Green IT, authored by San Murugesan and G.R. Gangadharan
- The Green Computing Book Tackling Energy Efficiency At large Scale, Wu-chun Feng

12.8 WEBLIOGRAPHY

- <https://www.techopedia.com/definition/14753/green-computing>
- <https://www.intechopen.com/chapters/75768>
- <https://datascience.foundation/datatalk/green-computing-the-future-of-computing>

12.9 UNIT END EXERCISES

1. Identify and discuss the four major ways in which organizations can gain value by greening an enterprise.
2. What are the major elements of the value chain, and how does 'closing the loop' relates to the value chain?
3. What are the major categories of information systems within an organization? Provide example of greening enterprise activities at each level.
4. How do enterprise resource planning systems integrate with greening enterprise activities?
5. How can green inter-organizational activities be supported with IT and IS activities
