Planning for Mobile Devices

Ch6

Topics

- Mobile device infrastructure policies
- Mobile device enterprise solution issues
- Disaster recovery in the mobile device infrastructure
- Mobile device backup and recovery
- Planning for new mobile technologies

Basic Mobile Device Concepts

- Organizations need to be able to manage mobile devices.
- Requires a formalized structure that provides for policy, security, resource management, and long-term lifecycle planning..
- Integrating the different aspects of mobile technologies into a unified approach to using and managing mobile devices in the business enterprise environment is called mobile device management (MDM).

History
Mobile Devices in the Enterprise

- Research In Motion Limited's BlackBerry first major smartphone to impact enterprise.
  - BlackBerry infiltrated the corporate market at a significantly faster pace than other devices.
  - BlackBerry had the first devices that could be centrally managed by an organization's IT department, and provided secure corporate email services to its users.
- Primary use to contact the employee more efficiently and exchange secure corporate communications.
- Weren't many user-friendly apps or things the user could do on the device for entertainment, such as play videos or music.
- At that time, mobile devices were really more of an extension of the normal IT paradigm.

Devices infiltrated the Enterprise.

- Many users eventually began to use personal devices to check company email and so on…
  - Corporate IT realized too late that mobile devices belonging to the user, but used to access organizational data, were here to stay.
- A new paradigm in IT management has developed, not only in dealing with mobile devices (both personal and corporate owned) in the enterprise.
- Employers have found that they can actually use the introduction of mobile devices into the workplace to their advantage, simply because they make users more productive.
  - Can also reduce IT costs by taking advantage of the fact that users want to use their own personal devices.

Mobile Evolution

- Several years ago, mobile devices might simply have been laptops.
- Changed with smartphones and then tablets.
- Now, corporate infrastructures include not only laptops but tablets, smartphones, and a host of other smart devices that connect to an IP network and process data.
- Most smart devices have applications that are able to actually produce work-related content (called content creation).

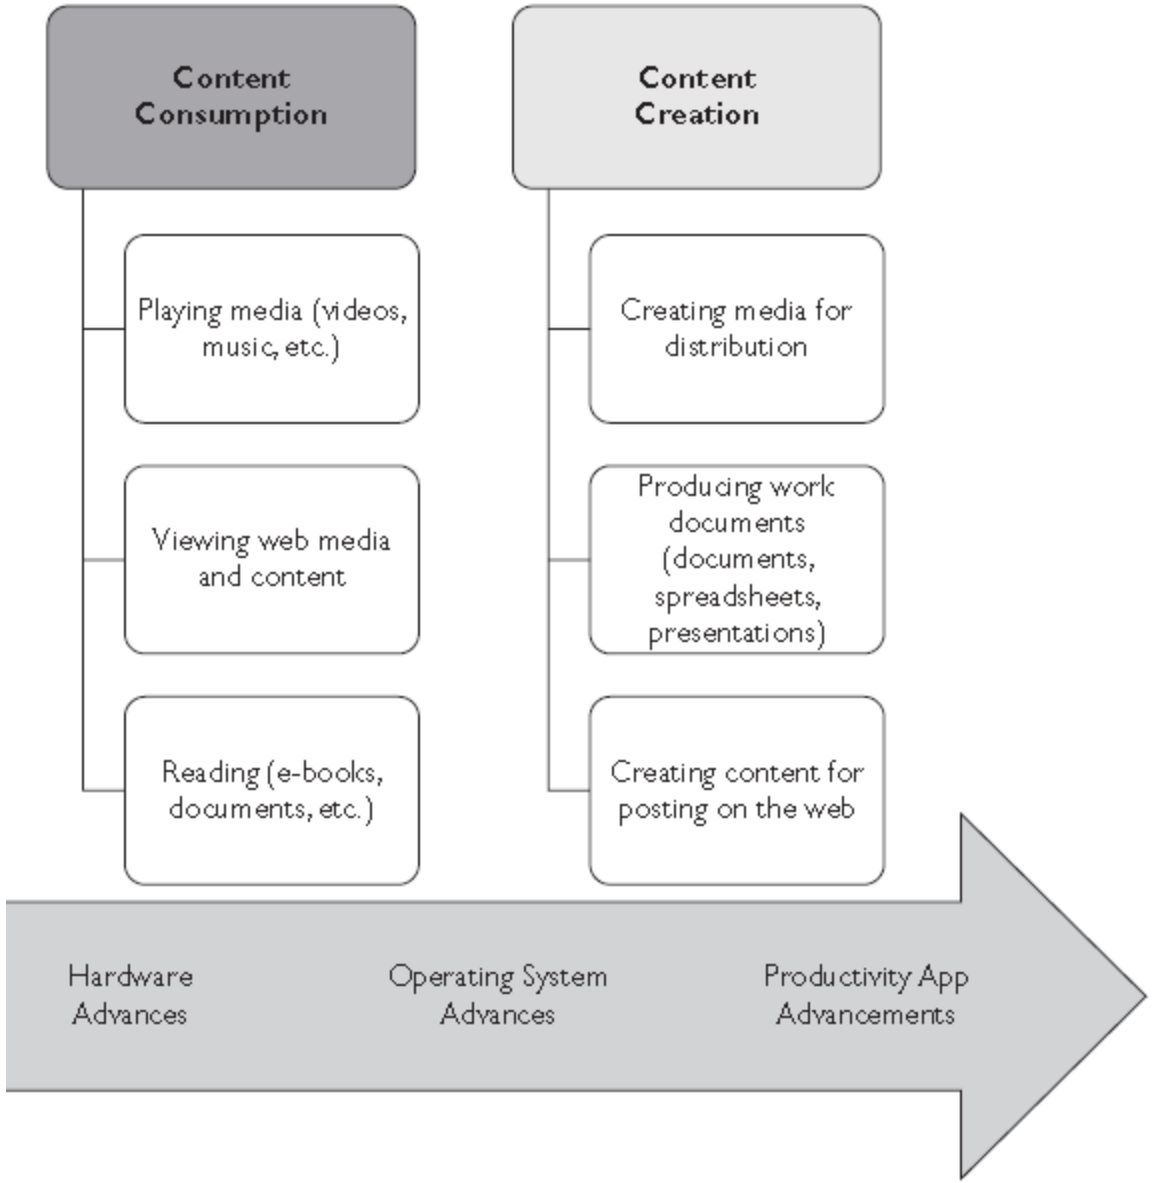

Two paradigms, and how advances in hardware, operating systems, and apps have really changed what we use mobile devices for.

One challenge, however, is that corporate data is often is stored or processed on mobile devices with little regard for security or control.

Comparing Apples to Androids: Apple

- With the past few iOS versions that Apple devices could be centrally managed in a formalized structure.
- Apple has a monolithic development model.
  - Apple solely controls the development of the operating system as well as the hardware device and also maintains strict controls on app development for its platform.
- Apple has very strict development policies and controls for developers.
- iPhone and iPad apps are almost exclusively installed and updated through the use of iTunes.
  - No third-party providers except for providers of line-of-business apps specific to a particular organization.
  - These internal development groups reside within an organization and can develop for Apple devices, but only for those that are under the organization's control.
  - They still have to undergo a type of Apple partnering and enterprise licensing approval process.

BlackBerry

- First in the enterprise with smart devices.
- First BlackBerry devices were just an extension of an organization's email…
- BlackBerry controls a significant portion of the enterprise market, particularly in the U.S. government.
  - Slowly changing as other devices gain market space.
- Like Apple, BlackBerry has a monolithic, or vertical, development and control model, in that they control the operating system and the hardware device platform, as well as apps that are developed for the platform.

Android

- Android is an open platform, based on yet another open platform, Linux.
  - Owned by Google.
- Android operating system is available for device manufacturers to alter or customize.
  - Some differences between vendors.
- The development model used by Android is more open.
  - Characterized by a wide variety of apps, as well as a wide variety of devices that run the Android operating system.
- Apps marketed and sold by a variety of sources, including the Google Play store.

Windows

- Microsoft has struggled to break into both the consumer and the enterprise smart device market spaces.
- Microsoft mobile devices lend themselves to centralized management and can become an integral part of enterprise mobile infrastructure.
- Microsoft also maintains its own app store, but there are also third-party app providers.
- Microsoft primarily controls the OS portion of its platform, but naturally has developer requirements as well, although these are not as restrictive as Apple's.

Mobile Devices

- A small sample of the many available devices from all four major platforms.

MDM Concepts
Mobile Device Management

- Mobile device management, as a unifying concept, basically means that the organization must develop a formalized structure that can account for all the different types of devices used to process, store, transmit, and receive organizational data.
- MDM can fit into the existing infrastructure along with desktops, servers, but it also addresses unique challenges of centrally managing devices that don't always stay within the corporate perimeter.
  - May include unique software specifically designed to manage mobile devices that also integrates into existing infrastructure.
- May make use of software that manages applications on these same mobile devices (called Mobile Application Management, or MAM).

MDM Challenges

- MDM attempts to manage: access control, secure identification and authentication, patch management, antivirus, and compliance with policy.
- MDM also manages things that traditional networks don't, such as provisioning a device and remotely installing and configuring software, and it can even be used to remotely wipe or lock a device.

The two major challenges

1. Mobile nature of devices
2. Extent to which personal devices are used in the infrastructure.

Mobile Aspect Of The Challenge

- Devices can be removed from the organization's physical boundaries, and sometimes its logical ones as well…
  - Makies it very difficult to manage security, patching, auditing …
- When a mobile device is connected to the corporate infrastructure, it can be managed, but when connected to an external network (such as the employee's personal network at home), it may be more difficult to manage without a direct, corporately controlled secure connection.
- Organization may, in some cases, have to simply manage the device as it connects to the corporate infrastructure.

Personal Devices in the Infrastructure

- Employees increasingly using their own personal devices to perform work functions, such as checking email, taking business calls, viewing sensitive data, and so on.
- Corporate data is being stored on personal devices.
  - A challenge because personal devices can be more difficult to control.
- Absent a policy restricting the use of personal devices in the organization, the employee has the ability to do almost anything he likes in terms of organizational data being processed on his own device.
- Attempts to control devices from the enterprise level can result in conflicts due to level of control and privacy issues.

MDM Infrastructure

- MDM encompasses the infrastructure, including servers, applications, security policies, and so forth, necessary to manage mobile devices that connect to the enterprise network.
  - Devices could be completely owned by the organization, or personally owned by the employee.
- MDM requires organization to have supporting infrastructure in place, which could include traditional network services, such as DHCP, DNS, email, security services, and so forth.
  - For the most part, these common core services will serve both traditional desktop and server devices, as well as the mobile infrastructure.
- Organization will have to implement some separate services specifically for mobile devices.
  - One such service is the mobile device management software itself.

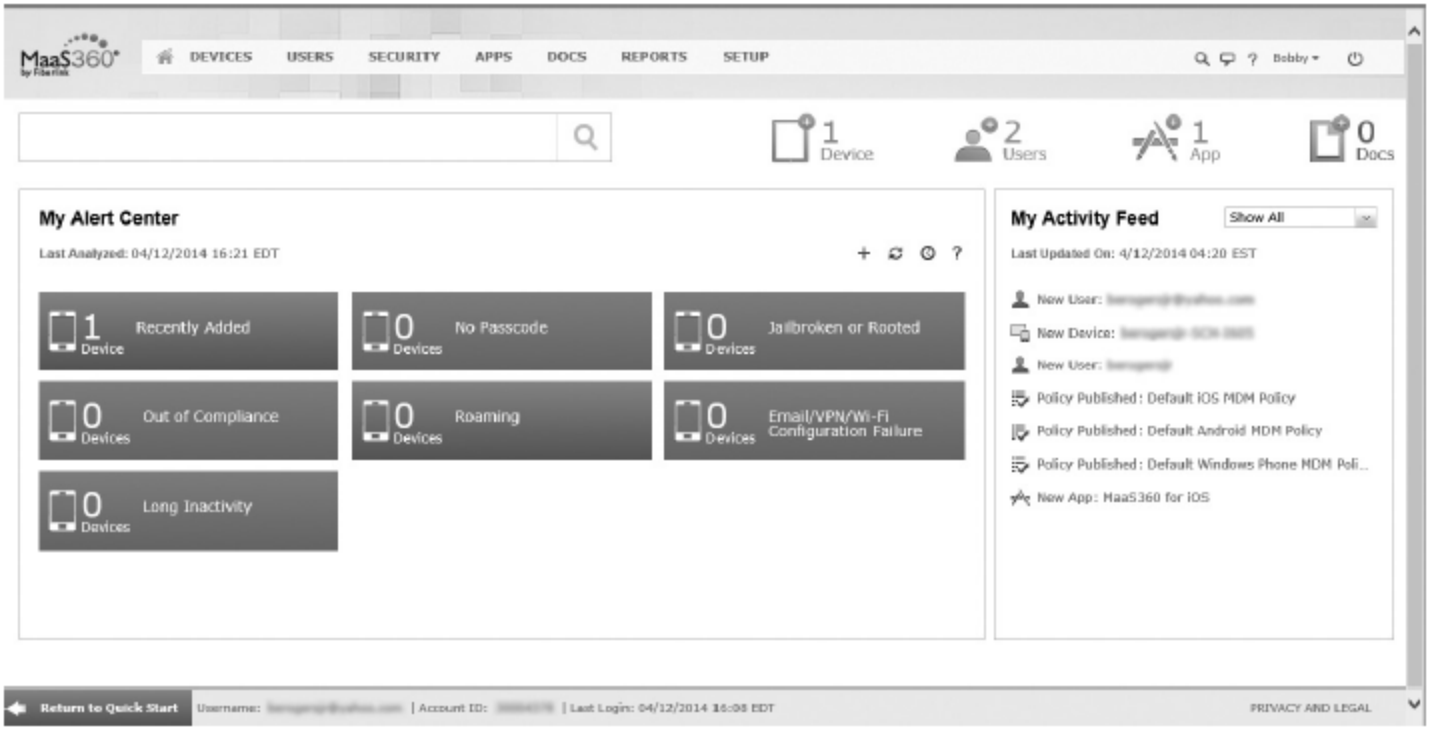

Figure 6-3 An example of a cloud-based MDM solution, MaaS360

- MDM software could be a commercial product purchased and licensed specifically for the organization.
- Also could be a cloud service.
- One example MaaS360 (above), a popular MDM cloud service that many organizations use.

MDM Policies

- Organization may have a centralized policy server as part of the MDM infrastructure that pushes security policy, as well as other types of policy, to the enterprise's mobile devices.
- Organization should use all of these different MDM components to initially provision devices, meaning that it will initialize, configure, and join the devices to the organization's infrastructure.
- Provisioning also includes setting up network services, such as email and Internet configuration, app store settings, and security settings.
- Beyond provisioning, mobile device management in the enterprise also means performing several actions on the device periodically, such as monitoring, patching the operating system, upgrading apps, and so forth.

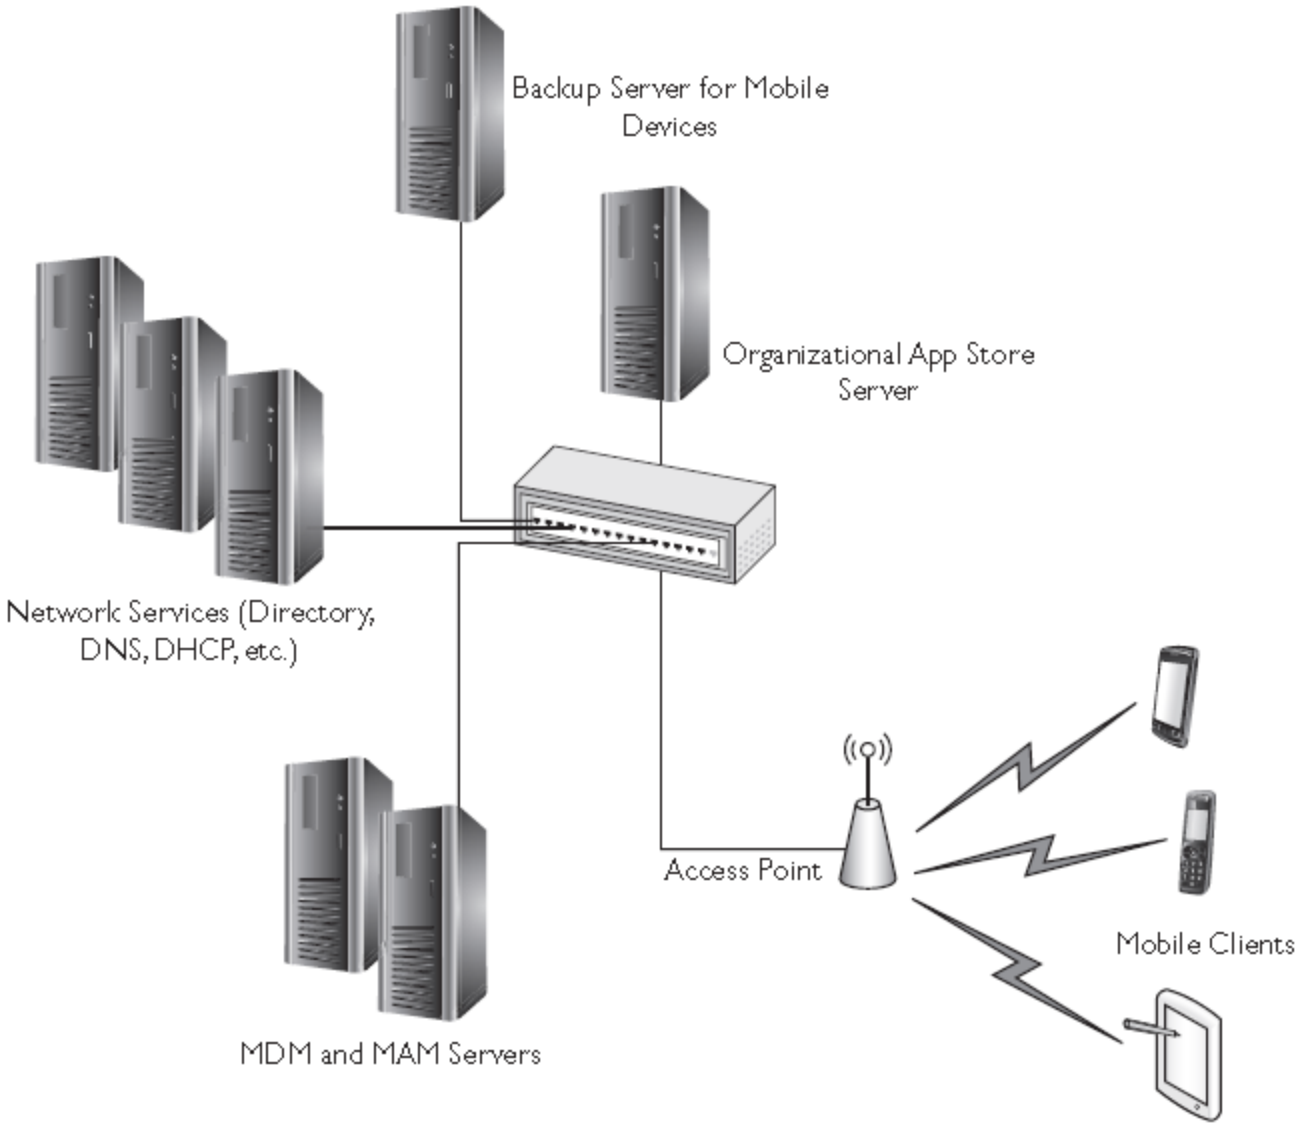

Figure 6-4 A notional MDM architecture

Bring Your Own Device (BYOD)

- BYOD war was briefly fought and lost by organizations hoping to continue the long-held tradition that IT assets belonged to (and were strictly controlled by) the company, not the individual.
- In some cases, companies may be able to enforce a policy that prohibits the use of personal devices to access corporate data and resources, particularly those in high-security environments.
- At the other end of the spectrum, some companies allow (and even encourage) personal devices, as it saves corporate IT dollars and can contribute to a much happier employee.
- Most organizations, however, probably fall into the middle of the spectrum and have a mixed environment of both corporate-owned and personally owned mobile devices.
- In some cases the organization may institute a cost-sharing program, subsidizing an employee's personally owned device by offering a monthly phone stipend or through discount agreements with mobile device and telecommunications vendors.

BYOD Challenge: Device Control

- How much control the corporation has versus the individual.
- If corporate data is processed or stored on the device, then rightfully so, the organization should have some degree of control over the device.
- On the other hand, if the device also belongs to the employee, then obviously the employee should have some control over it.
- Conflict is probably best solved by policy.

Cost and Privacy

- If the organization allows the user to use her own device for company work, does the organization help pay for the monthly bill or compensate the user for its use?
  - Best solved by defined formal policy and procedures.
- Yet another and equally important BYOD challenge is employee privacy.
- If policy allows the organization some degree of control over the device, what degree of privacy does the user maintain on her own device?
- Can the organization see private data or have the ability to remotely access or control a user's personal device and its use?

Two critical BYOD issues:

1. Personal data privacy versus protection of corporate data
2. Level of organizational control versus individual control.

Policy and the Mobile Infrastructure

- Policy governs everything in the enterprise infrastructure, to include security, equipment acquisition, provisioning devices, setting up users, and so forth.
- Policies are promulgated (top) down from the organization's senior management to be followed and implemented by everyone in the organization.
- Where policy states what is required in terms of compliance, procedures are used to supplement and support policies by explaining how something is done.
- Procedures must be developed to support any mobile device policies in place, such as requirements for encryption or authentication, for example.
- Before mobile devices are purchased and provisioned, the organization must develop a solid set of policies for them.

Organizational IT and Security Policies

- Information technology policy not only applies to mobile devices, it applies to every IT asset in the organization.
- However, because mobile devices present unique challenges and issues, special attention should be paid to developing good mobile device policy that balances the needs of the organization with available resources, device functionality and use, and, of course, security.
- Information technology policy, combined with other organization policies, such as acquisition or resource management policies, will drive what types of mobile devices are accepted and acquired for use in the organization.
- The organization may have a policy, for example, that states that only iPhones or BlackBerry devices are to be used for official company business.
- The company may allow users to bring their personal iPhones, or may purchase and issue them to its employees.
- The same set of policies may have requirements regarding who in the organization gets to use company devices, such as senior management or the sales force, for example.
- In any case, the acquisition and use policies for mobile devices should spell out how these devices will be acquired, who will be allowed to use them, what they will be allowed to be used for, and other conditions for use.

Policies

- Policies should also dictate terms concerning the extent of use of personal devices in the enterprise, either supporting or restricting a BYOD program.
- In the event the organization allows the use of personal devices to store or process company data, then the policies must also state what the responsibilities are of both the organization and the user in protecting any company data on these devices.
- Policies should dictate the level of control the organization has over the device, especially if it is a personally owned one.
- Elements stipulated in this policy may include the organization's ability to restrict network connections or applications used on the device, or even control over the data itself, including the user's personal data that may reside on the device.
- These policies should not be taken lightly and created without considerable thought, because of the ramifications that can occur when dictating corporate control over personally owned devices.
- Users must be required to agree to these policies in order to use their devices in the corporate infrastructure, so naturally they should be well briefed on the policy and the consequences of abiding by it.
- To be fair, users should also be given the option of not using a personally owned device to access corporate data, if they desire.
- Figure 6-5 gives you an idea of how policies relate requirements to implementation.

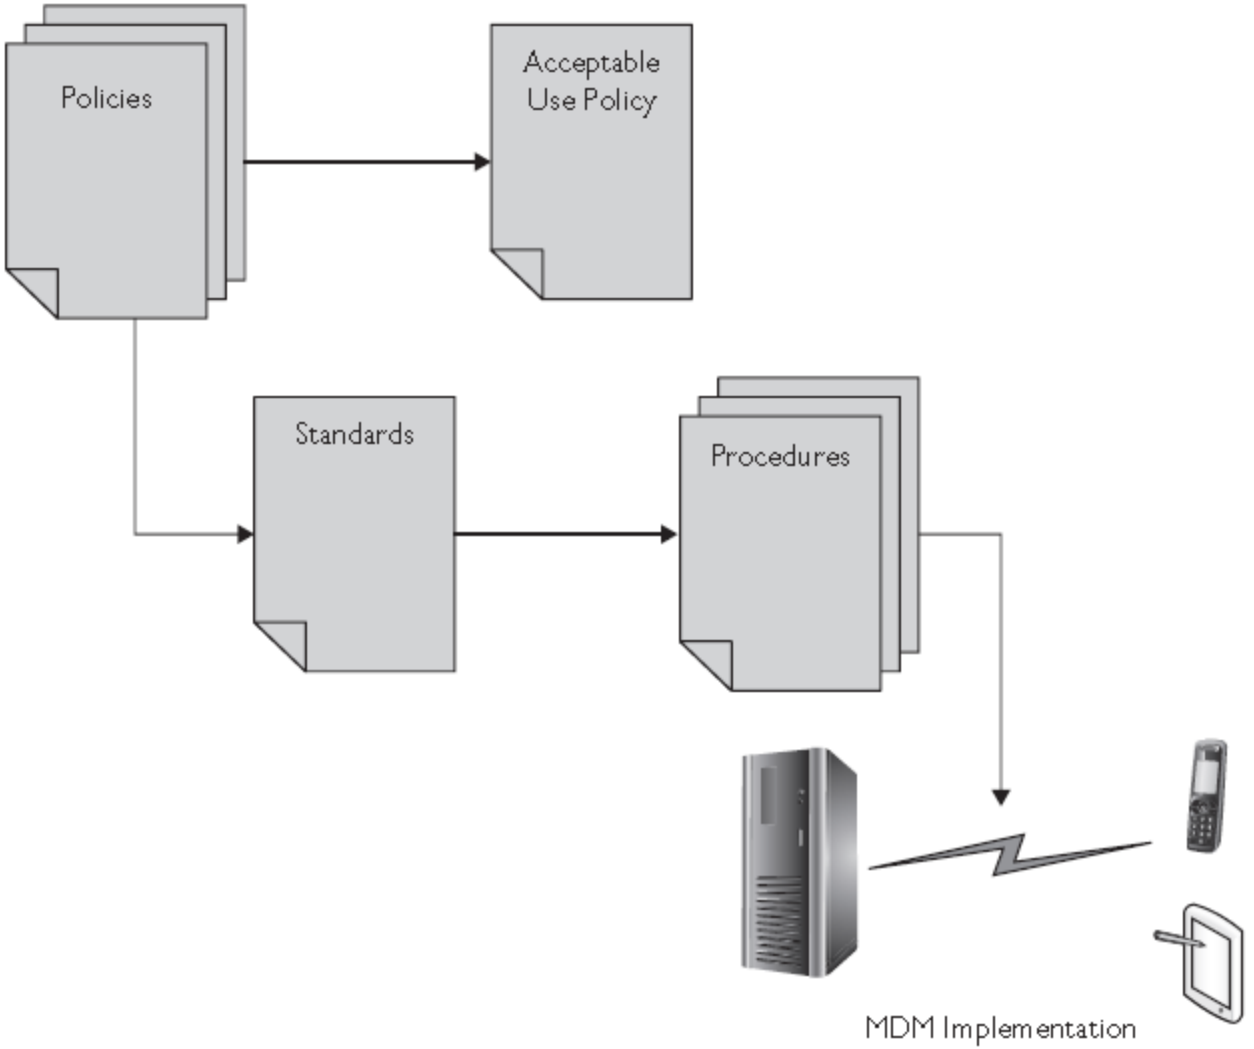

Figure 6-5 Relationship of policies (requirements) to standards, procedures, and implementation

Backup, Restore, and Recovery Policies

- Managing data backups and recovery should also be governed by organizational policy.
- Backup policies, which also should cover data restoration and recovery, dictate how often backups should be performed, to what degree they should be performed in terms of the amount of data that is backed up, and how quickly backups must be restored in the event of data loss.
- Backup policies should also dictate what types of backup are run, as well as the backup schedule itself.

Balancing Security with Usability

- The question of how much security is enough versus how much usability the user should be allowed to have depends upon the organization's policies which depends on an organizations tolerance for risk.
- With more usability comes greater risk, but with more security may come less productivity.
- The organization must take a true risk-based approach to security and functionality, not only with traditional devices, but with mobile devices as well.

Vendors, Platforms, and Telecommunications

- In planning a mobile device infrastructure, there are concerns with standardization across the enterprise.
- The organization may choose to have a uniform standard for devices and operating systems, to include only specific vendors' platforms, or only specific versions of a given operating system.
- For organizations that have a more homogenous mobile device hardware population, it still may desire standardization for each OS and version used in the enterprise.
- There are also considerations with app stores and OEM vendors, as well as telecommunications providers.
- Obviously, the more differences in operating systems, versions, providers, and so on, the more difficult it would be for administrators to centrally manage and keep up with the configuration on those devices, as well as manage challenges with billing, updates, interoperability, and supportability.
- Standardization, as much as the organization can attain, is a significant goal in mobile device management in the enterprise.
- This, again, can be affected by organizational policy, which can set standards for different devices, platforms, and providers.

- One other interesting aspect of policy that relates to telecommunications vendors and carriers is the issue of expense management.
- The organization should have a defined policy that addresses who pays for monthly carrier expenses, especially if the users are in an officially sanctioned BYOD environment.
- The company may help offset the users' expenses incurred while conducting organizational business on personal devices, for example, by subsidizing monthly bills or even paying the charges outright.
- In an environment where the organization owns the devices, this is a bit more clear-cut; likely the organization will bear all (or at least the majority) of the expenses.
- In any event, policy should be created to handle these issues.

OS Vendors

- Choice of platform vendors depends upon several factors include interoperability with existing infrastructure, the ability of the organization to support the device platform, resources available, and, of course, security.
- Each different device platform has its own advantages and disadvantages that the organization must weigh in deciding what the standard will be.
- In some cases, the organization may not settle on a particular vendor or platform as a standard, especially in an environment where there already are many different types of devices and operating systems in use.
- Even when settling on a particular hardware platform or OS vendor, the operating system vendor may have different OS versions available for a given device or platform.
- These all may be supported by the OS vendor, but for the sake of standardization, the enterprise should attempt to use the same version for all its devices to the extent possible.
- This may be problematic in some cases because there are different versions of the same device.
- For example, a given version of iOS may be the enterprise standard, but because of different versions of iPhones purchased in different fiscal years, the organization may be required to support different iOS versions until it periodically refreshes those legacy devices.
- Careful acquisition planning, as well as IT lifecycle planning, should be considered so that devices are upgraded or refreshed on a periodic basis, keeping in mind operating system version and device supportability as well.

OEM

- OEM (Original Equipment Manufacturers) vendors provide additional hardware, accessories, and even apps for mobile devices.
- Some of these may be optional, and some may be purchased by the organization out of necessity.
- For example, a given enterprise timekeeping app that comes from an OEM, or specific piece of hardware designed to plug into smartphones and scan credit cards for POS purchases, may be acquired for business purposes.
- Decisions to acquire these types of software and hardware from OEMs are also influenced by factors I discussed previously: interoperability, supportability, security, and cost.

Telecommunication Vendors

- In an infrastructure where the enterprise maintains tight control over its devices, and standardizes device hardware and operating system platform, the carrier used by the organization will likely also be standardized.
- The factors that go into selecting a carrier are based upon different elements such as pricing structure, mobile device volume supportability, and so on.
- In a less structured environment, such as a BYOD environment, or somewhere in between those two extremes, there may be several different communication vendors used by the organization to provide data services to its mobile devices.
- In addition to cost and supportability, other factors that affect standardizing with a particular telecommunications vendor include data rates, throughput, coverage areas, and so on.
- All of these elements can vary with different carriers, and should factor into the organization's decision.
- Additionally, each different telecommunications vendor can add its own unique changes to different hardware devices and operating systems.

- There may be slight differences in operating systems, apps, and even supported features found on the device, which can be attributed to the telecommunications vendor and the changes they may make to the device and its software.
- As an example, both Verizon and AT&T bundle different applications or additions to the operating system in some of their devices.
- Some of these modifications may or may not significantly affect the functionality of the device, but are typically used to personalize the device for the carriers' networks or infrastructures.
- One significant example of a carrier-specific change that did, in fact, affect functionality of a device was when AT&T previously imposed a limitation on iPhone devices that restricted the ability of the device to tether (share its Internet connection) to another device.
- As this example illustrates, the organization should completely investigate all of its carrier choices, as well as how the carrier alters devices, apps, and platforms, to ensure that they choose the right telecommunications vendor.

- Policy is a critical aspect of planning and implementing an MDM, and includes a wide variety of policies that cover security, standardization, interoperability, acceptable use, and many more.

Enterprise Mobile Device Infrastructure Requirements

- In order to implement a mobile device infrastructure on an existing corporate network, there are several requirements that should be met.
- Careful planning and involvement from all levels of management and technical personnel.
- The organization first has to decide on a policy, of course, to include privacy issues, the use of personal devices, standardization of devices, apps, and so on.
- The organization also has to have the infrastructure that can support a managed mobile device implementation.
- Includes the existing services that are provided to non-mobile clients, such as desktops.
- Also includes new infrastructure required to successfully integrate and manage mobile devices.

Security Requirements

- Security should be one of the major requirements that an organization considers when implementing a mobile device infrastructure.
- Unlike some of the big security issues the desktop operating systems suffered through their evolution, mobile devices, their operating systems, and their apps, are much more secure.
- Nonetheless, security is still an issue in the mobile device world.
- Some of the bigger security problems you see in the mobile device world relate to data loss from mobile devices because of the way they're used sometimes, or because they are not securely configured.
- When developing security requirements for implementing a mobile device infrastructure in the enterprise, several items should be considered.

Device Groupings

- Device grouping allows administrators to place devices in certain groups, based upon logical user groupings or other requirements.
- Allows administrators to apply different policy options to the different groups of devices, such as security policies, software or patching policies, or any other updates based upon the requirements of that particular group of devices.

Devices can be grouped by:

- Organizational unit, geographical location, and class of device, such as laptop, tablet, smartphone, and so forth.
- User function, such as administrators, sales personnel, and so on.

- MDM software can also be configured to group devices according to organizational specific requirements.
- For example, suppose you wanted to group devices by vendor platform or an operating system version. You may want to do this in order to push certain patches down to devices that have a certain version of operating system on them, but not to others.
- For example, you could also group them by VLAN or IP subnet if you wanted, for segmentation purposes.
- Device grouping would allow you to do that, and this should be an important management tool used in any mobile device management infrastructure.

Figure 6-6
Device groupings and different policies

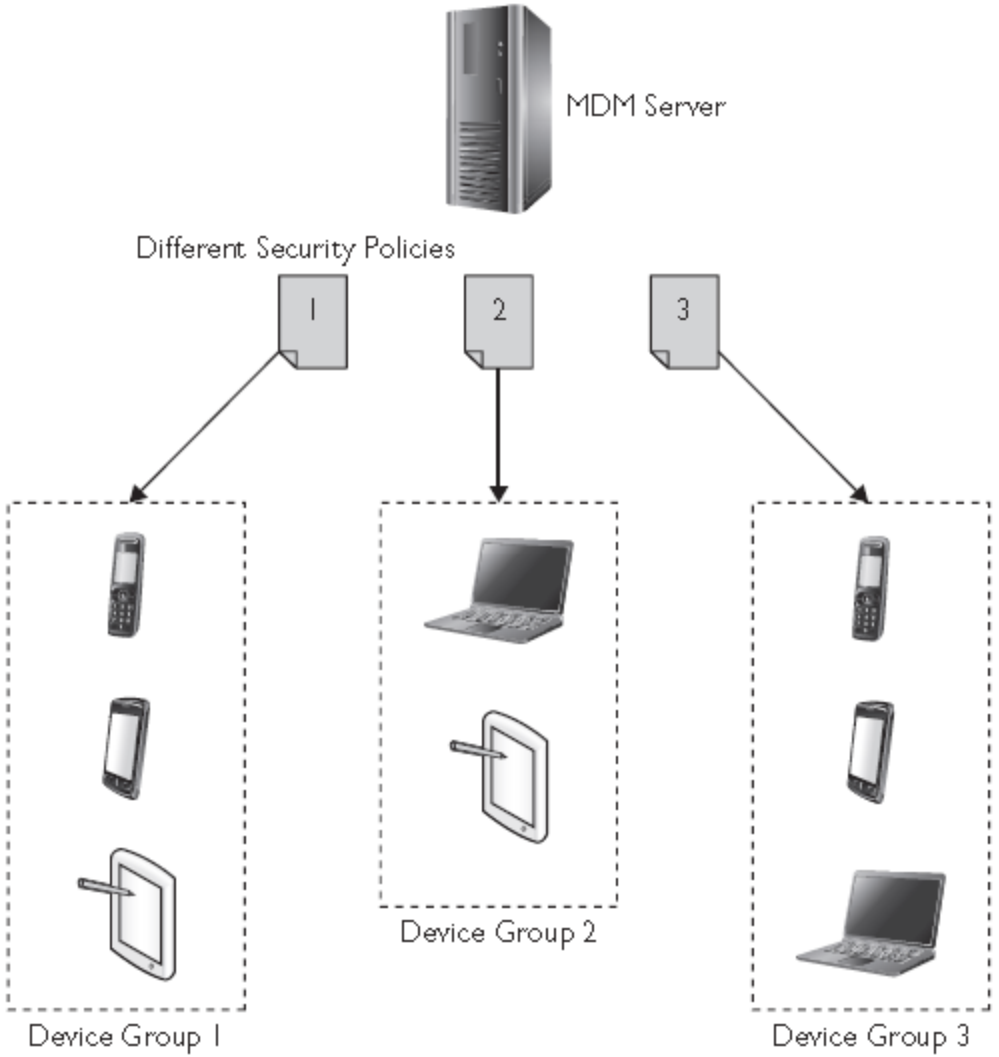

- Figure shows one example of how device groupings could help an organization from a security policy perspective.

Administrative Permissions

- In order to effectively manage mobile devices, organization has to have the elevated permissions necessary to perform privileged actions, such as installing software, patches, and upgrades; audit device use and access logs; and configure security settings on the device.
- In an environment where all devices are centrally controlled and managed by the organization, this is not difficult.
- However, in an environment that has adopted a BYOB paradigm, this may not be easily accomplished.

Password Strength

- The enterprise policy for mobile devices should closely mirror that of traditional desktop devices and other infrastructure devices.
- Only when legacy mobile devices cannot support the complex passwords required in the organization should policy be different for these types of devices.
- You should set the password strength requirements in the security policy according to the requirements of your organization.
- By and large, the same password policy used for traditional devices on the network should be used for mobile devices as well.
- Keep in mind, however, that many mobile devices may use a personal ID number (PIN) in addition to passwords for different apps.

Remote Wipe

- Remote wipe refers to the ability of the organization or the individual to access the device via wireless or cellular signal, and institute a complete erasure of the device's storage.
- Normally, a remote wipe would be used in the event the device is lost or stolen.
- A remote wipe ensures that sensitive data on the device is not compromised when the user no longer has positive control over the device.
- Most modern mobile devices, particularly smartphones, allow for some type of remote wipe capability, either through the MDM infrastructure, or using a remote wipe app provided on the device and accessed via a Web or desktop app.
- One major issue of remote wipe is that the data will be lost permanently unless it has been backed up; however, this may be preferred to losing data due to unauthorized access.

Remote Lock/Unlock

- Yet another feature related to security is the ability to remotely lock or unlock a phone for a user.
- This, of course, has to be accomplished through some type of network connection, whether it's cellular or Wi-Fi, and can be a function of an app on the device or, preferably, through the MDM infrastructure. An administrator may want to remotely lock a phone if it has been temporarily lost or is not currently under the control of the user.
- The organization may want to do this if it could reasonably believe the user will get the phone back under their control, but needs to make sure that no one can access data on the device in the interim. Remote unlock is a feature that the company may want to have in the event the user locks the phone and can't remember the passcode, for example.
- In this case, if the user inputs the wrong passcode too many times, it may actually wipe the device, if it is configured as such in policy.
- This offers the user a way to get the device unlocked without risking a device wipe.
- Again, this has to be configured on the device and in the MDM infrastructure.

Captive Portal

- Captive portal is usually an administrative or management function, possibly implemented as a web site, or through a network access control device, whereby the user is prevented from entering the network infrastructure until they are properly authenticated or verified as having sufficient need to connect to the network.
- Can also ensure that devices meet certain requirements in order to connect to the corporate infrastructure.
- Requirements may include antivirus updates or security solutions installed on the device, proper authentication and encryption mechanisms implemented, connection to a specific subnet or VLAN, or any other requirements the organization wishes to impose on a device before it connects to the network.
- This may be something that happens when the user connects to a self-service portal or when the device is first provisioned and connects to the network for the first time.
- One example of a captive portal is that of an open Wi-Fi network, for example, that requires a web site authentication from the user in order to proceed to connect to resources on the Internet or other internal assets.
- A captive portal can also be used for guest networks, for example, if the user doesn't have full authorization to access most resources on the internal network.

Monitoring and Reporting Capabilities and Features

- Monitoring and reporting allows you to ensure that:
- Devices are being used according to policy
- Security posture of the device is being maintained in terms of configuration and patching
- Actions of the user are appropriate.
- Most MDM software allows you to implement monitoring and reporting on devices, as long as the device operating system supports these functions.

Interoperability and Infrastructure Support

- Another consideration in planning an MDM implementation is the infrastructure support the organization already has or is willing to add.
- Obviously, implementing an MDM program from scratch would require a significant investment in time, money, training, and other resources.
- Leveraging existing infrastructure wherever possible is likely a good idea if it is feasible.
- For example, using existing services, such as DNS, DHCP, email, and security services, is not only feasible, but definitely recommended for integration purposes and economy of use.
- Adding additional infrastructure, such as MDM servers, cellular signal boosters, wireless LAN controllers, and so forth, should be done with seamless integration in mind. If these new pieces of infrastructure are added to an existing network that can't support them, due to legacy equipment issues
- or bandwidth issues, for example, then the existing network may have to be upgraded first before the MDM pieces are introduced.

Interoperability

- When planning a mobile device infrastructure, you should examine your existing infrastructure and compare products, protocols, services, network requirements, and application requirements, as well as user requirements, to make sure that the mobile devices you implement in the infrastructure interoperate with existing technology.
- Not unusual to have to make adjustments to the existing infrastructure, including equipment configuration, topology, addressing schemes, and so on, when installing new pieces to support mobile device management.

Self-Service Portal

- A self-service portal is the function of a mobile device management infrastructure that allows mobile device users to connect to the network, possibly to a shared folder or self-service web site, and perform many different functions for themselves without the need to interact with mobile device administrators. Self-service functions may include installing certain common software apps, such as antivirus updates, trusted certificates, proxy settings, or even application and operating system updates.
- Many different functions can be accessed through a self-service portal by the user, but some will require intervention from the administrative staff, usually those involving administrative permissions, such as significant configuration changes and major operating system upgrades.

Device Platform Support

- Organization has to determine what types of devices it will support, including from which vendors, which operating systems, and from which application stores it will allow users to download apps. Factors that affect the decisions regarding the device platform support include supportability, cost, flexibility, level of integration into the existing network, and even how deeply embedded existing devices already are in the infrastructure.
- For example, forcing all of the employees that have previously used Apple devices to switch to Android devices may affect the ease in which the organization can support those particular platforms.
- For the most part, an organization may settle on one major platform and try to steer all of its users to that platform, especially in a situation where the organization owns all the mobile devices.
- However, in a BYOD implementation, this is likely not possible.

On-Premise vs. SaaS

- Cloud services include Software as a Service (SaaS), Infrastructure as a Service (IaaS), Platform as a Service, (Paas), and, as previously mentioned, even MDM as a Service (MaaS).
- Considerations involved with both on-premise and cloud providers.
- Cost, ownership and control, security, and responsiveness.
- The organization must look at these factors from the perspective of cost savings, naturally, which would include the necessity to buy, provision, and maintain equipment, train personnel to operate the services on-site, as well as do the same for supporting infrastructure.
- However, from a security and control perspective, the organization may have limited control over a cloud provider's handling of the organization's data and services.
- The organization should do a cost-benefit analysis to determine whether on-premise support is more cost-effective and secure than contracting out services to a third-party provider.
- One way to ensure that there is a good balance of cost, security, control, and so on, is to have a well-written, well-vetted service-level agreement.

Multi-Instance

- Multi-instance refers to the use of multiple instances (meaning multiple installations) of software configured for specific groups of users or even having different configuration settings per installation. The MDM solution you choose to implement may provide the ability to have multiple instances of software or settings based upon different configuration requirements in your organization.
- Some individual software packages, such as Microsoft's SQL Server, also provide this capability built-in.
- There are a number of different reasons that an organization might want to use different versions of software.
- User groups may have different mobile hardware devices or even different versions of an operating system on the same hardware platform.

- In addition, different groups of users within an organization might have unique security requirements and require different instances of software to support their requirements.
- Multi-instance software can be managed by the MDM software itself, if the software supports multiple instances.
- In the multi-instance model, the software itself may execute on a remote server and allow the mobile device user to interface with it through the user's browser, or through a front-end app that connects to the instance.
- Cloud based services, such as those that provide SaaS, typically provide for multi-instance use,as that is part of their business model.

- Licensing would be one issue you would have to consider when deploying multiple instance software, as well as network configuration.
- Some software licenses may permit the use of multi-instance use, but others may not.
- Each separate instance may also require different ports or IP addresses.
- It also may be licensed on the basis of a certain number of users that can simultaneously access the software instance at a time.
- In the case of database access, software may require separate database stores.

Location-Based Services

- One major feature of mobile devices is the ability to track the device's location through GPS, cellular, and, when needed, Wi-Fi connections.

Geo-Location

- Because of the mobile nature of personal devices, and the built-in location-based technologies that mobile devices usually have, geo-location services are a must for MDM programs.
- In addition to helping users find out where they're at and where points of interest are near them, geo-location can be used by the organization to help keep track of devices in the event they are lost or stolen, to aid in their recovery.
- Aside from that, geo-location can also help make sure that users are compliant with policies that may prohibit the use of their mobile devices in certain locations.
- There are many features of both end-user apps and MDM that require the use of geo-location services, so it's a good idea to ensure that the service remain enabled on the device.
- This can be accomplished with configuration policies pushed to the device during its initial provisioning, or even afterwards.
- For personally owned devices, this may be a little bit more problematic, and the organization should set policy accordingly.

Geo-Fencing

- Geo-fencing is an interesting way to keep track of both mobile devices and users.
- Geo-fencing uses the geo-location capabilities of the mobile device to locate it and keep track of its location while it is within a pre-defined electronic perimeter.
- This electronic perimeter is set up using electronic sensors around a physical location, such as an office building, warehouse, or even a business campus area.
- For a few examples of how this might be useful to an MDM infrastructure, picture being able to determine which mobile devices are currently in the building, and, by extension, their users.
- Or think about being able to detect when an organizationally controlled smart device enters the building. The system could detect it and send notifications to you or even the device.
- If the device is a tablet used for inventory or sales, for example, the system could pick it up and automatically download its data from the day's work.
- Also consider how valuable this system would be to prevent mobile device theft, or even unauthorized use, if the device is not allowed to leave the property.

- you could conceivably use this capability to track user locations as well, and this might not sit too well with users. Employees might consider this a form of workplace surveillance, and in some cases may rebel somewhat at using geo-fencing for this particular use. In the most benign cases, they may simply leave the device somewhere on a desk and leave the building anyway, or, in the most serious instances, they may seek legal advice and consider bringing litigation to the organization for invasion of privacy.
- Depending upon how geo-fencing is used to track employee movement within the bounds of the employer's property, there may be legal ramifications to using it to track workers.
- It's a good idea to research the legal issues of using geo-fencing to track employees, as well as intelligently discuss the merits and pitfalls from an employee satisfaction perspective, and try to get a realistic view of what benefit you may or may not actually get from this practice.

Mobile Application Management

- Mobile application management (MAM) is a concept related to MDM, but on a different scale.
- With MDM, the idea is to reach out and control a device in its entirety through the corporate infrastructure and policy. MAM is limited to simply controlling the applications on the device itself.
- There are several different ways this is possible, including controlling individual apps, controlling the source of the apps, controlling the security features of the apps, and controlling the app's data. MAM usually isn't a solution by itself; it's typically used in conjunction with MDM to varying degrees.

- Organizations have the ability to implement MAM as part of their overall MDM infrastructure, but it isn't necessarily a requirement to implement MDM in order to have MAM.
- Organizations often choose to simply manage applications on mobile devices, particularly in a BYOD environment, rather than manage entire devices, their updates, their security, and so forth.
- Two of the reasons an organization may choose to simply manage applications on the device include the size of the infrastructure and cost; if an organization has too few users to warrant implementing a significant MDM infrastructure, but yet has a desire to remotely manage a few key apps the employee uses in relation to corporate data, then this might be a good reason to implement MAM without MDM. Another reason may be the organization's tolerance for risk; the corporate decision may be to accept any risk incurred by users managing their own devices while connected to the corporate infrastructure or accessing and storing corporate data on their devices.

Different ways to implement MAM.

- First, absent centralized management software specifically used for managing mobile apps, the organization could build management capabilities into the app itself.
- This could be accomplished by pre-configuring the app to communicate with certain corporate servers, using particular authentication and encryption methods, and restricting permissions on the app.
- This would effectively control what the app could do on the device, and how it could interact with the device and the user, as well as other apps and their data.
- The problem with this approach is that the organization would have to manage this for every single app it wanted to control. MAM is more of a centralized solution, and requires the use of specialized enterprise software either combined with an MDM solution or implemented separately.
- Another problem may be that the app, especially if it comes from the device vendor or another third-party developer, may not lend itself to be manageable to the degree the organization would like.

- Another way to manage apps on mobile devices is simply to use integrated MDM and MAM software to control the devices' operating systems, which would in turn be used to control apps on the device. Regardless of whether an organization chooses to implement MAM with or without MDM, however, there are still some characteristics of MAM that require centralized policy, planning, and management. I'll discuss these over the next few sections. Figure 6-7 illustrates the differences between using MAM as part of an MDM solution, and implementing MAM separately.

Application Store

Figure 6-7
MAM employed both independently and as part of an MDM solution

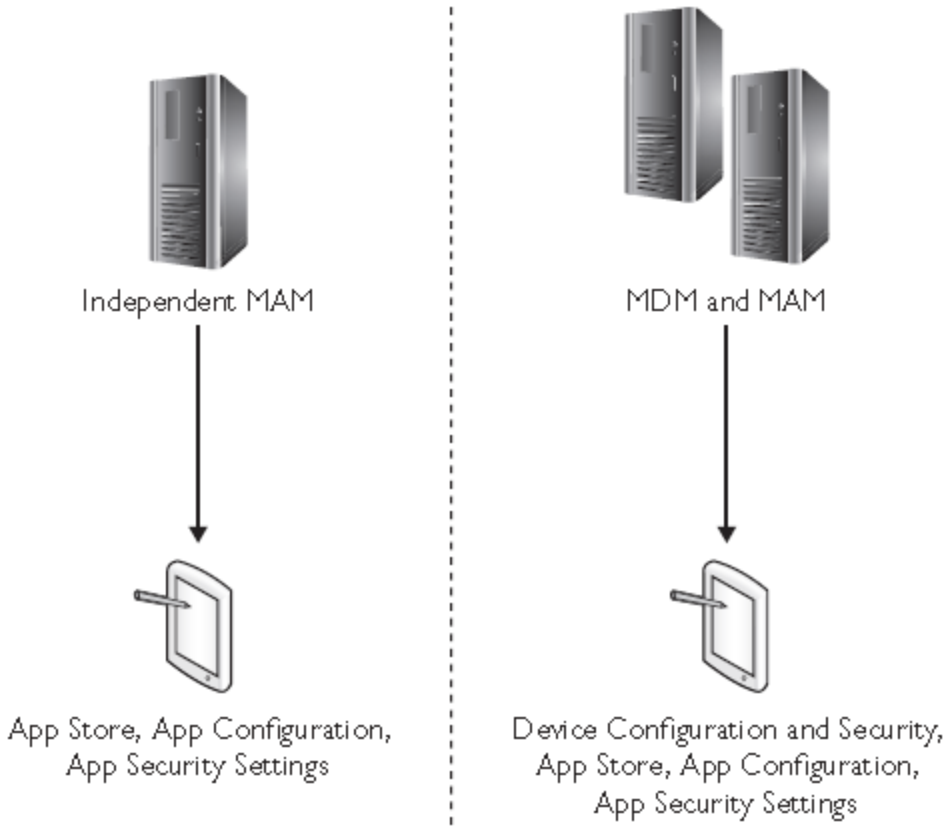

- Although the focus of this discussion is really on application management by the enterprise, it's worth it to discuss some major differences between both the vendor app stores and how mobile devices receive apps from organizational app stores

- While Apple tightly controls its app store and how apps are introduced into the Apple marketplace, for instance, Android users can install apps from other sources in addition to the Google Play store (a process popularly known as sideloading).
- These apps may come from independent app developers or enterprise-specific app stores created to develop applications specifically for the mobile users of a particular organization.
- Apple actually has several different ways to distribute iOS apps. Obviously, the most common way is through iTunes, which is Apple's own app store.
- Developers use the iOS Developer Program, where they submit an app to Apple for approval to be included in the app store.
- Apple has very stringent requirements, for both quality and security, so this can be a difficult and rigorous process for developer, particularly if the developer wishes to distribute the apps to anyone else.
- Figure 6-8 shows an example of the ubiquitous Apple iTunes storefront.

- Android app stores, such as Google Play, also have strict requirements on apps developed and marketed through them, but these requirements are not necessarily enforced in the same ways.
- There are guidelines for developing secure code, ensuring that there is no malware in the app, and so forth, but it still a lot simpler to get an app into an Android store than it is into Apple's store.
- This way of doing business in no way means that Android apps are less secure or of lesser quality, however.
- While most users of smart devices are familiar with a vendor application store, such as iTunes, or Google Play, for example, there are third-party app stores, depending upon the platform used, such as Android or Microsoft devices, from which users can acquire apps.
- The focus of our concern, however, is app stores that are provisioned and managed by the organization themselves.
- An organization might develop its own app store to deploy apps specifically for its enterprise users. Given the distribution and approval model used by stores such as iTunes and Google Play, an organizational app store is often a better alternative for several reasons.

- First, the organization can control the developmental process and ensure that quality apps are produced for its users and their devices. Second, the organization can make sure that apps include the level of security controls necessary for use in the enterprise.
- Control over which apps are installed and used on the device is yet another reason for employing an organizational app store.
- Another reason is cost.
- For larger organizations that have in-house developers, it may be more cost-effective to simply develop an app that is specific to the enterprise, versus the cost of buying a volume license for apps that may not fully meet the organization's requirements.
- In any event, establishing an organization app store and managing it through the MDM and MAM infrastructures may be an integral part of the centralized management of mobile devices in the enterprise.
- Keep in mind that at its most basic level, an app store is really nothing more than a repository for software, so mobile users can simply click on a link or go to a network share to find and install apps.
- In order for an organization to distribute Apple iOS apps to its employee users, the organization has to participate in the iOS Developer Enterprise Program, enabling them to develop and distribute apps specific to their own needs.

- These apps are digitally signed by the company, ensuring that they are legitimate and have been designed around Apple's rigorous standards.
- They can, however, only be distributed to users associated with the organization. One interesting fact about apps developed by an organization using this model is that they have to use what is called a provisioning profile.
- The provisioning profile contains metadata about the app itself, as well as information about its developer. The provisioning profile is pushed to an Apple device via the device's cellular or wireless connection.
- Provisioning profiles can also be included in the app itself to provide for ease of installation by the user.
- Because Android employs a more open method of getting apps to users, enterprise app stores don't have to go through any of the hoops that they would for iOS apps.

- EXAM TIP Understand the ways that the different app store vendors develop and distribute apps for their respective devices.

Default Applications and the Enterprise

- For ordinary users with their own devices, the default vendor applications that come with the devices may sometimes be enough.
- Most people, however, usually shop through a vendor's app store specific to their device in order to find more apps.
- In a corporate environment, the default vendor applications may not be what the organization wants.
- Second, as security is a much more important issue in the enterprise space, default applications may not be configured securely or may not offer security features the organization needs.
- There also may be some default applications that are not approved for use on devices owned or managed by the organization, such as those used for social media, for example.
- Policy and standardization will typically determine what apps are used on centrally managed devices, and they may need to be located in the enterprise app store for users to download and install.

Pushing Content

- Most mobile app management is done through a push model.
- In a push model, applications centrally managed by the enterprise app store are typically automatically installed on the device, as are updates to both the device's operating system and the different apps that reside on the device.
- Occasionally, the MDM infrastructure may poll the device to determine what version of apps and patch levels the device is currently using, and may install an upgraded version as needed.

Devices can be notified of updated or new apps via push notifications.

- Push notifications are messages sent to the device from the app store, informing the device about a new version of an app or, in some cases, about new configuration settings.
- One important item about pushing content from the application store is that, regardless of whether or not the organization has its own enterprise app store, in the case of iOS or Android apps, these push notifications don't actually come from the enterprise app store at all; they actually come from the OS vendor.
- For example, the Apple Push Notification Service requires that the enterprise register its app with Apple and will send notifications through Apple servers for forwarding to the device.
- Similarly, Android devices work the same way, using the Google Cloud Messaging (GCM) service.
- The main difference between Apple and Android notifications is that GCM notifications can hold more data, but both similarly interact with individual apps on their respective devices.

Disaster Recovery Principles and the Mobile Infrastructure

- Two main focuses of disaster recovery.
- First business continuity concerned with keeping the business going after a major event or incident.
- Business continuity requires careful planning and actions, such that the business can successfully recover from an event and continue to operate, fulfilling its mission and goals.

Disaster Recovery

- Is concerned with all the activities immediately following an incident or disaster that are designed to help the organization and its personnel recover to the state that business can continue.
- Disaster recovery could be considered a subset of activities in the overall business continuity effort. However, disaster recovery is focused on a different set of planning activities and actions.
- For our purposes here, we're going to look at the different general disaster recovery and business continuity principles that your organization should adhere to in order to keep its business going and recover from any type of incident.

Business Continuity Principles

- Business continuity concerns itself with the organization's ability to maintain its mission after a significant negative event has taken place.
- Whether this event is man-made or natural, the business must be restored to some level of operation or it risks shutting down completely and not surviving the event.
- Business continuity planning is all of the careful planning and activities that ensure a business can survive an event and continue in its mission.
- While disaster recovery typically comes before the business can resume operations, understand that the business continuity planning function is a long-term, on-going activity in the organization, and should be happening long before a disaster actually occurs.

- One of the first steps in business continuity planning is to identify the organization's critical processes and business operations.
- These processes and operations, if lost, would significantly affect the organization's ability to conduct its business.
- The organization should determine criticality for each of these processes, and how much protection each of them requires in order to maintain them after a negative incident or event.

- The second step the organization needs to take is to inventory all of its critical assets, such as equipment, data, and yes, even people, to determine how these assets relate to these critical processes and operations.
- The organization must determine how much protection each of these assets should be afforded to protect them in the event of a negative event or disaster.
- The organization should also prioritize these different processes and assets for protection and recovery after an event.
- These first two steps in the process are commonly called a business impact analysis, or BIA.
- In addition to identifying and prioritizing the different processes and assets, this part of the planning should also show what the impact would be if that particular article process was lost or that asset destroyed, for example.

- Table 6-1 shows an example of how an organization might list and prioritize these different processes and assets for business continuity planning.
- Keep in mind that this is a very simplistic example, of course, but could be used as a starting point in your business process analysis.
- Take a look at the National Institute of Standards and Technology (NIST) Special Publication 800-34, Contingency Planning Guide for Federal Information Systems, which offers some really great information on planning and conducting a BIA, to include useable templates.

Process or Asset	Value	Replacement Cost	Organizational Priority (1–10)	Impact If Lost (High, Med, Low)
Order intake process	$10,000 in sales per day	$20,000	10	High
File server	$5,000	$3,500	6	Med
Administrative Computer	$2,500	$2,500	3	Low

Table 6-1 A Sample View of a Business Impact Analysis

- Once the business impact analysis is complete, the organization should commit resources and planning to protect those processes, and its critical assets.
- Business continuity planning should also include a step-by-step process for bringing these critical processes back online after a disaster or other negative event.
- Plan should include an analysis of how much data the organization can afford to lose or how much time the organization can afford to be out of commission before bringing operations back
- up to an acceptable level.

- Careful planning has to go into disaster response and recovery, to include establishing a response team
- and chain of command, assessing the situation, and ensuring personnel and critical equipment are protected during a disaster.
- Disaster recovery also includes restoring equipment and data, and supporting infrastructure so that business continuity activities can proceed and the organization can resume operations.
- The most important part of disaster recovery planning is the recovery plan itself.
-

- The disaster recovery plan, or DRP, is designed to provide the organization a set of concrete, step-by-step actions and activities the organization will take in the event of a disaster, in order to preserve lives,
- prevent injuries, and save critical equipment and data.
- The DRP should address the different range of possible events that could happen to the organization, on a reasonable basis, of course, and plan the organization's response to each of those possible events.
- For example, part of the DRP should cover how the organization will respond to a serious weather condition, such as a tornado, for example.
- May cover alerting personnel, shutting down equipment, securing the facility, and evacuation if necessary.

BCP/Dr

- Beyond protection of personnel, planning should cover protecting facilities, equipment, and data.
- During a disaster that immediately threatens human lives, such as serious weather conditions like tornadoes, flooding and hurricanes, or fires, priority should be given to saving lives and there may not be time to secure equipment or data.
- However, if the threat is not immediate, or if there is some time leading up to the event (such as, for example, a few days' notice before a severe storm or hurricane), the DRP should cover the orderly backup, shut down, and security of systems and data, as well as the facilities if at all possible.
- Activating the disaster response and business continuity plans in advance of an event, wherever possible, is also a good idea to give the organization additional preparation.

- Other normal business operations, such as backup processes, for example, serve to supplement the disaster recovery and business continuity planning, and should be utilized with that in mind.
- For example, routine backups of the servers should be performed with a quick restoration and minimal data loss as the goal, and possibly should be located off-site in the event a disaster occurs.
- Other processes, such as alternate work locations and so forth should also be created with disaster recovery planning in mind.

Disaster Recovery Locations

- Often, a disaster will completely render the primary business location unusable.
- This is definitely a possibility in the case of fire, tornado, hurricanes, or flooding.
- In this event, the business may not be able to restore operations back in its primary site, and should plan in advance on an alternative location for restoring operations.
- There are different ways that this could be accomplished.
- Owning or leasing a separate facility located some geographic distance away from the primary site is one way of ensuring an alternate processing location.

- Cloud services provided by third parties could be another solution, especially if most of the organization's business processes occur online or via the Internet.
- It's still likely, however, that there will need to be an alternate physical location for employees to work at and restore business operations to.

Hot, Cold, and Warm Sites

- There are three types of alternate locations that provide differing levels of readiness or support for restoring business operations after a disaster.
- Each of the sites has different advantages and disadvantages directly related to how fast the business needs to restore operations, as well as how much of an investment the business can put into the alternate location.
- The first type of location is the cold site.
- The cold site is usually nothing more than a bare facility with empty space for equipment and offices. There's usually some limited level of utilities turned on for the site, such as heat, water, and electricity, but nothing more than that, to include Internet or communications access.

- In the event of a disaster, the organization would have to take a great effort in physically relocating people, equipment, supplies, and so forth to the cold site.
- In a large scale natural disaster, this might prove to be difficult due to road conditions, lack of vehicles, disrupted public utilities, and so forth.
- Disaster recovery planning, when using a cold site, should take all of this into account, as well as the time and effort required to set up the cold site for business operations.
- One key advantage to using a cold site is that there is probably very little investment required. Basically, the organization only has to lease or own an empty building, for example.
- The key disadvantage to using a cold site is that the time required to restore operations could be excessive, and it's possible that an organization would not have the manpower or tools to successfully relocate personnel or equipment over to the cold site very effectively.

Warm Site

- A warm site provides the next level of readiness for business operations restoration above a cold site; in addition to workspace and basic utilities, it also could provide communications links, such as Internet and phone service.
- A warm site also usually has some basic level of equipment already installed at the site. This could include simple office furniture, such as desks and tables, but also likely includes equipment such as servers, workstations, and other types of equipment used to bring the business back into operation.
- This equipment could be turned off and simply waiting for someone to flip the switch and turn the power on, and it could also require a quick data restoration from organizational backups to ensure that it is using the most current data available from the organization's business transactions. In some warm sites, organizations often restore data from backup on frequent basis, so that the business can come back to operations much more quickly.

Warm Site

- The key advantage to a warm site is that the time to recover business operations is much less than it would be when using a cold site.
- The key disadvantage is that a warm site requires much more of an investment of time, money, and resources.
- The organization may find, however, that the extra expense is worth it if the amount of money and business lost in the event of a disaster would be much more than what it invested in the warm site over time.
- So, business criticality and the need to restore operations much faster would be a deciding factor in maintaining a warm site.

Hot Site

- As you can imagine, the hot site is capable of providing business operations much faster than a cold or warm site.
- In a hot site, not only do you have workspace, utilities, communications services, and equipment, but this equipment is usually maintained in a high state of readiness, powered on and ready to process data. Backups from the main processing site may be transferred to the hot site and restored very often, even on a daily basis.
- This would ensure the hot site has the ability to pick up processing quickly with minimal loss of data in the event of a disaster.
- Obviously, a hot site requires much more expense of resources on the part of the organization in order to maintain this high state of readiness.
- A hot site would be used in business operations that are extremely critical, and when anything beyond a very small amount of downtime in business processing is intolerable to the organization.
- In maintaining a hot site, not only would an organization have to maintain a completely separate facility in a high state of readiness, but also all of the redundant equipment and supplies used in the facility would be almost a duplicate of anything in the primary facility.

Network Device and Server Backups

- I've mentioned backups a few times throughout this discussion over the past few sections, so it's probably very appropriate that I discuss it in depth at this point.
- This discussion assumes that the organization already has a routine backup strategy in place, and follows best practices by having different levels of backups, to include full, incremental, or differential backups, or even organizationally customized backups for specific data sets, such as transactional or specific data backups, to restore different levels of data at different required restoration points.
- This discussion also assumes that backups follow the best practice of being stored at an off-site facility, in order to protect them from disasters that could happen at the primary site.
- These are important best practices that most organizations cognizant of business continuity will follow, but they are worth mentioning here as well.

- Network device and server backups are important because they are the first line of protecting the organization's entire infrastructure data.
- This can include all the business process data, financial transactions, employee data, and any other relevant organizational data, on a massive scale.
- Server backups include not only data related to the business, but also operating system data as well. Without backing up this type of data, it would be almost impossible to restore the organization to a functioning capability in the event of a disaster or serious incident.
- Network device backups, on the other hand, usually consist of backing up the operating system (if applicable), and any configuration files or settings, so the device can be restored to an operational state very quickly.
- While some data processed by network devices may not necessarily be backed up (i.e., real-time traffic passing over a device) because it will never be restored, per se, other peripheral data created by the device; logs and other audit trails, for example, must be backed up for a variety of other reasons, including security auditing, regulatory compliance, and so on.

Directory Services

- There are several critical business and information technology processes for an organization, but few are as critical as directory services.
- he reason for this is that modern IT infrastructures rely so heavily on directory services for security, resource location, and many other critical services.
- For example, imagine what the workday would be like if first thing in the morning, when most of the users came into work and started logging into their workstations, they could not authenticate to the network because the directory services were down. In addition to network authentication, many network services, such as shared folders and DNS, for example, rely on directory services to do their job.

- Because of this criticality, directory services must be maintained in a high availability state for users in the network. In the case of Microsoft Windows Active Directory (AD) services, the distributed data stores that make up AD are often spread across several servers for redundancy purposes, in the event that one server fails or becomes unavailable, for example. In addition to this load-balancing type of set up, the Active Directory database is typically backed up several times a day in order to ensure that a current copy of the directory is maintained.
- Like directory services, several other infrastructure type services are also critical and should be protected to the use of backups and redundancy.
- Some examples of these services may include the domain name service, the DHCP database, different business critical Web services and related databases, financial and accounting transaction databases, and even key data from users' workstations.

Frequency of Backups

- How often and to what extent an organization backs up its data depends upon several factors.
- The first factor is how critical the data is.
- When doing business continuity planning, the organization should take a good hard look at its data assets and determine how critical they are.
- In some cases the organization has to prioritize different types of data to determine how critical they are and how much protection they should be afforded.
- Obviously, more critical data should be afforded better protection, which could cost the company more in terms of equipment and money.
- For example, maintaining backups on critical company servers to process customer financial data is probably a higher priority than maintaining equipment and the necessary backups to provide for data protection for an administrative assistant's workstation.

- Another factor, aside from data criticality, that affects both frequency and degree of backups, is how quickly data would have to be restored, given a catastrophic event that destroys the primary data source on a server, for example.
- Certain types of backups take longer to execute than others, and certain types of backups restore more quickly than others.
- Essentially, the amount of data backed up is the determining factor in the trade-off between a quick backup and a quick restore, with some consideration, of course to the speed and efficiency of the backup equipment used.
- With that said, it's probably a good opportunity to talk about the three primary types of backups that most organizations perform.

Backups

- Three basic types of backups:

1. full,
2. incremental, and
3. differential.

- Obviously, each of these different types of backups is chosen based on the amount of data the organization wants to back up, as well as the resources required to execute each type, such as media capacity and cost, and efficiency of the backup equipment.
- However, each type also varies in both its backup time and the time required to restore the data from the backup. Let's discuss each of these in a bit more detail to clarify.

Full Backup

- Completely backs up the entire hard drive or an entire data store or whatever else the administrator specifies as a complete backup of an asset.
  - Could also mean the operating system and applications on the server's hard drives.
- That way, if the server's drives fail, the complete contents of the hard drives, including the operating system, applications, and data, could be fully restored back to the point that existed before the failure.
- Because a full backup can get every piece of data an organization needs to completely restore a server, it would seem that this would be the preferred type of backup every time a backup is executed.
- Full backup can take quite a while to back up as well as restore.
- If the server crashes in the middle of the day and has to be restored from a full backup, it may take more time than the organization can afford to spend in downtime.
- Most organizations perform a full backup on an asset on a regular basis; once per week is a fairly normal schedule for executing a full backup, but it really depends on the organization.

Incremental Backup

- An incremental backup is a type of backup that will back up any data that has been added or changed since the last full backup or the last incremental backup.
- The incremental backup is able to determine what data has changed through the use of the archive bits set on data files.
- Most file systems have the capability to set different bits on files that determine whether or not they have been backed up recently, or have changed.
- If the file has changed, the archive bit is turned on, and this lets backup software know that the file needs to be backed up. Once the incremental backup has occurred, the archive bit is turned off.

- So if a full backup occurs, and then a file changes, an incremental backup will back up that file and then turn the archive bit off.
- Assuming the file doesn't change any more, the next incremental backup will not back up that file.
- It will only back up files that have the archive bit set to on, meaning they have changed since the last full or incremental backup.
- An incremental backup does not take very long to back up the data because it backs up only a subset of the files on the media.
- However, an incremental backup can take a long time to restore, simply because the full backup must be restored first, and then every single incremental backup performed since the full backup was executed must be restored, in order, until all the data has been successfully restored.
- For example, let's say that a full backup of a server was performed on a Sunday night.
- On Monday, Tuesday, Wednesday, and Thursday, an incremental backup was performed on each of those nights. If the server's hard drive crashed on Friday, then the full backup from Sunday night would have to be restored first, and then, in sequence, the incremental backups from each of the succeeding nights would have to be restored in order to make sure that all of the data changes were accounted for. Depending upon the amount of data involved, this may take a while, and for a critical server, this amount of downtime required to restore data from a backup may be unacceptable.

- The third type of backup is the differential backup. A differential backup is similar to an incremental one in that it does not back up every piece of data from the media.
- Like incremental backups, a differential backup also relies on the archive bit on a file to determine what data has changed since the last full backup. If a full backup is performed on a Sunday night, for example, and then data changes on Monday during the day, a differential backup will back up that data because the archive bit has been turned on for the file.
- However, unlike an incremental backup, a differential backup does not reset the bit and turn it off. The archive bit remains turned on.
- So, the next time a differential backup is run, it backs up all the data with the archive bits turned on, including all data that has changed since the last full backup. So its backups of data are inclusive of all changes since the last full backup.
- Because of this, the first differential backup in a series doesn't take very long to perform. However, as time passes and more data changes, each subsequent differential backup takes a bit longer.

- The advantage of this is that because backups are inclusive of all data changes since the last full backup, restoring data requires less time.
- For example, let's say that a full backup was performed on the server on a Sunday night. Files changed on Monday during the business day, and a differential backup was performed on the server Monday night. It would back up all files that had changed since the last full backup because the archive bit on those files would be turned on for those files.
- Because the differential backup does not turn off the archive bit, the next night the differential backup is run (Tuesday night), it backs up data that changed on Tuesday as well as the same data that changed on Monday, because the archive bits are still turned on.
- It does this every single day, increasing the amount of data that it backs up, as well as the time it takes to perform that backup.
- On Friday, the server hard drive crashes, and the full backup is restored from Sunday night, but the only differential backup that is required to be restored, in addition, is the last differential backup performed on Thursday night.
- This is because it contains all of the changes to all of the data since the last full backup. So while differential backups may take progressively longer to perform, they are usually faster to restore.

Note

- Back up and recovery have been previously covered.