Topic 1: Change Management Best Practices

Q 1.1: What is change management, and why should I care?

A: No matter how large or small your network environment, change is inevitable. Hiring new employees, adding new offices, supporting new network services, improving security, fixing bugs—all of these activities result in change, especially to your network infrastructure devices, such as routers, switches, hubs, firewalls, and so forth. Although change is almost always a good thing in the end, change can cause bad things to happen. For example, a careless typo in a firewall configuration file could have alarming security implications. So no matter how minor or beneficial a change may be, you should always approach change with a healthy dose of caution. *Change management* is a set of policies and procedures that you adopt and follow to formalize that caution into a repeatable, consistent process.

At its simplest, change management simply means keeping track of the changes you make and evaluating proposed changes for their effect before actually implementing them. In practice, change management involves some fairly well-defined tasks:

- Maintaining documentation that describes the current configuration of all network devices
- Maintaining documentation that describes the purpose and details of any changes
- Maintaining an archive of older configurations so that they can be used in an emergency
- Implementing policies that control the rate of change
- Implementing policies that control who may perform changes

Why should you bother with all of that? Primarily, to improve network uptime. Unauthorized or unplanned changes are the number one cause of network device failures and unplanned downtime for organizations. Failure to document current configurations makes it difficult, if not impossible, to recover gracefully from a failed change procedure. Failure to control the rate of change as well as who can make changes results in an inconsistent environment that is difficult to maintain long-term. Instead of thinking of change management as extra work, I like to think of change management as *saving* me work: By simply following some simple methodologies and processes, I can ensure that changes to network devices never become a nightmare. Or, at least, if they *do* become a nightmare, I can quickly recover without having to spend all night at the office!

Q 1.2: What's the best way to "do" change management with network devices?

A: The actual mechanics of change management depend on which types of devices and tools you have on your network; the ways in which you should conduct a change management program, however, are universal. There are two main steps to a change management program: planning and management.

Planning for Change

Too many change management methodologies ignore the planning phase, which is perhaps the most important. Planning allows you to identify and reduce risk, provide a means to rollback in case of disaster, and so forth. Essentially, planning requires you to

- Identify everything that could possibly go wrong as a result of a change.
- Assign a level of likelihood and severity to each potential risk.
- Identify means of mitigating risks or, at least, provide a means of recovery should the risk actually become a reality.

A solid change management planning methodology will make it easier for you to prioritize changes according to their business impact. For example, if you find yourself making several high-risk, low-benefit changes, you can implement policies to reduce such activity, for example, by adopting a policy of only making low-benefit changes during a regular update cycle, such as at the end of each month.

How you actually conduct each step of the planning process depends on your environment and your personal preferences. In the next four sections, I'll provide some examples to get you started.

Identify Risks

What might go wrong when you update the routing table on one of your routers? Many possibilities spring to mind:

- You could mistype something and corrupt the entire routing table, making the router functionally useless.
- You could enter incorrect information, preventing the change from working properly.
- You could enter incorrect information that makes existing routes stop working correctly.
- While uploading changes to a router, you could lose your network connection, resulting in a partial change to the router.
- You could upload changes to the wrong router, causing routing problems across the network.
- Any incorrect changes you make could replicate across your network through routing protocols, corrupting all of your network devices.

The objective with your risk list is to identify everything that could *possibly* go wrong, not just the things that are *likely* to go wrong. Keep in mind that changing the configuration of *any* network device, not just a router, creates a set of potential risks.

> ☞ Keep your risk lists handy! After you've developed a list of risks for a particular type of change, such as a router update or a firewall change, keep that list. You're likely to make the same type of change again in the future, so there's no reason to unnecessarily repeat the risk-identification process. You will be building your risk list into a checklist for *avoiding* risks, so the list can become part of your network's change management documentation and act as a list of procedures to be followed to help avoid unnecessary risk during network device management.

Categorize Risks

After you've got a list of everything that could go wrong, assign likelihood and a severity to each item. I prefer a simple scale of 1 to 3, where 1 represents highly unlikely risks, or risks that would be very minor if they did occur, and 3 represents risks that are likely to occur and would be very severe if they did. Working with the previously created list of potential risks, you might assign the following ratings:

- You could mistype something and corrupt the entire routing table, making the router functionally useless—likelihood is 2, severity is 3. The likelihood is high because you manually type all the router configuration information and, although you're always careful, there's no data-validation process in place.
- You could enter incorrect information, preventing the change from working properly—likelihood is 2, severity is 1. Severity is less than that of the first risk because you're simply failing to implement the change, not affecting anything else.
- You could enter incorrect information that makes existing routes stop working correctly—likelihood is 2, severity is 2. The severity is 2 for this risk because you're affecting an entire device.
- While uploading changes to the router, you could lose your network connection, resulting in a partial change to the router—likelihood is 1 because you've got backup power supplies everywhere and a very reliable network; severity is 2 because if the risk did occur, it would take the entire device offline.
- You could upload changes to the wrong router, causing routing problems across the network—likelihood is 1 because you are careful; severity is 2 because if you did make this blunder, you would ruin an entire router.
- Any incorrect changes you make could replicate across your network through routing protocols, corrupting all your network devices—likelihood is 3 because if you do make an incorrect change, it *will* propagate fairly rapidly; severity of 3 because this mistake could potentially take your entire network offline.

The purpose of this list is to help identify the risks that are in most need of specific mitigation. The risk list for a switch reconfiguration might include similar items, but the risks listed would be unique to switches; the same can be said of firewalls, managed hubs, or any other network device. One simple way to rank your risks is to add your two ratings, giving you a prioritized list of things that could go wrong:

- Any incorrect changes you make could replicate across your network through routing protocols, corrupting all of your network devices—risk: 6.
- You could mistype something and corrupt the entire routing table, making the router functionally useless—risk: 5
- You could enter incorrect information that makes existing routes stop working correctly—risk: 4
- You could enter incorrect information, preventing the change from working properly—risk: 3

- While uploading changes to the router, you could lose your network connection, resulting in a partial change to the router—risk: 3
- You could upload changes to the wrong router, causing routing problems across the network—risk: 3

With this list in hand, you're ready to start planning ways to avoid these risks and, should the worst happen, recover as quickly as possible. Again, although I'm using a router in this example, you'll want to prioritize the risks associated with changing any type of network device.

Mitigate Risks

Risk mitigation is a planning process in which you try to think of ways to prevent your identified risks from ever occurring; while at the same time coming up with a means of recovery should the risk become a reality in spite of your efforts. Add the mitigation and recovery ideas to your list to create a risk-avoidance and recovery checklist:

- Any incorrect changes you make could replicate across your network through routing protocols, corrupting all your network devices.

  Avoidance—Disable routing protocols on router until change is verified by a senior administrator.

  Recovery—Ensure that a backup of all router configurations is available before you make a change. In the event that incorrect data propagates, immediately restore device configurations from backup.

- You could mistype something and corrupt the entire routing table, making the router functionally useless.

  Avoidance—Use vendor-supplied tools to make changes rather than manually entering changes. Vendor tools provide some data validation to help prevent data-entry errors. Also, document all changes and have another administrator review and approve them for accuracy. Have the other administrator verify the accuracy of the changes after they are made.

  Recovery—Back up the device configuration before making a change. Immediately restore device configuration if changes made do not comply with the change documentation.

- You could enter incorrect information that makes existing routes stop working correctly or prevents the change from working properly.

  Avoidance—Use vendor-supplied tools to make changes rather than entering changes directly in router. Vendor tools provide some data validation to help prevent data entry errors. Also, document all changes on paper and have another administrator review and approve them for accuracy. Have the other administrator verify the accuracy of the changes after they are made.

  Recovery—Back up the device configuration before making a change. Immediately restore device configuration if changes made do not comply with the change documentation. Always have a workstation available that can make a direct serial connection to the router in case an incorrect change disables the router's ability to communicate via network interface.

- While uploading changes to the router, you could lose your network connection, resulting in a partial change to the router.

  Avoidance—Ensure that router, administrative workstation, and intermediate devices (hubs and switches) are on power backup. If possible, place an administrative workstation on same network segment as the router to be changed to eliminate the possibility of an intermediate router failure during upload.

  Recovery—Back up the device configuration before making a change. Ensure that the router being changed is accessible to a local-segment workstation on which the back up resides, allowing easier restore. Always have a workstation available that can make a direct serial connection to the router in case an incorrect change disables the router’s ability to communicate via network interface. As a last-ditch recovery method, many network devices offer a hardware reset switch that restores the device’s factory configuration. Combined with a recent configuration backup, you can use this reset function to quickly get the device up and running again.

- You could upload changes to the wrong router, causing routing problems across the network.

  Avoidance—Have another administrator confirm your changes and settings prior to upload.

  Recovery—Back up all network devices before making a change. If data is uploaded to the wrong device, restore that device’s configuration from backup. Always have a workstation available that can make a direct serial connection to the router in case an incorrect change disables the router’s ability to communicate via network interface.

Some network devices, such as managed hubs and switches, might offer simpler recovery methods. Some managed hubs, for example, can create a backup of the last-known good configuration to a built-in flash RAM module, and let you recover that configuration with a hardware reset switch. Other network devices, such as firewalls, might require more extensive planning to ensure that a fast recovery is possible.

☞ After you’ve developed a complete risk list, including mitigations, for a particular type of change, save it! This list should become a checklist for all future changes of the same type. By following the checklist each time you make that type of change, you’ll automatically mitigate the potential risks as well as have prepared recovery options in case the worst happens. If your network administration is primarily accomplished by junior administrators, these mitigation lists can become a mandatory part of the procedures the administrators follow, helping ensure that you’re sort of looking over their shoulder, even when you’re not.

Prioritize Changes

Don’t get into the habit of making every change that pops into your head. Prioritize changes based on their impact on business operations. You can use a simple 1-to-3 scale or something more complex. High-priority changes are worth more risk, of course, while lower-priority changes—especially those with a high-risk rating—should be put off until they can be made under tightly controlled circumstances. For example, I’ve worked with companies that save all low-priority changes until the end of the month. Before implementing any changes, they carefully review them all. They also back up every single network device in case something goes horribly wrong, and they put the necessary support personnel on alert. This process requires a lot of effort and isn’t something that these companies want to go through on a daily basis. For emergency changes that need to be implemented immediately, the companies have a fast-track process that requires two senior administrators to approve and implement the change; the idea being that senior administrators have enough experience to pull off the change with less risk. How you prioritize and handle changes really becomes a matter of change management policy, which I’ll discuss next.

Managing Changes

Changes can easily get out of control, and the only way to rein them in is to have in place a firm set of change management policies that all administrators are required to follow. For example, you might implement a change management policy as follows:

- All changes must be documented and approved by a senior administrator. Change documentation must include the current state of the device as well as the proposed change.
- Changes identified as high-priority require a senior administrator’s approval. All other changes require the approval of two administrators, including at least one senior administrator.
- All changes must include a detailed description of the intent of the change (for example, “To allow the Nevada office to communicate directly with the Seattle office rather than communicating through the New York hub office.”)
- All completed changes will be reviewed at a weekly meeting of administrators. This meeting will help make all staff aware of recent changes and allow an opportunity to review failed changes.
- Changes classified as emergency priority can be made only by two senior administrators working together. These changes can bypass the normal review process, but that process will be completed as soon as possible after the change is complete to ensure a complete set of documentation for the change.

The actual policies your company might adopt may differ; however, the important thing is to have some procedural guidelines in place.

Want to Know More?

No matter what you do, make sure that you have a system in place for change management. If you'd like some ideas for how to physically implement such a system, check out the University of Kentucky's Change Control FAQ, located online at http://www.uky.edu/~change/faq.html. That should give you some ideas of how a change management system works at a very high level, including change requests, tracking, and so forth. You should also check out Cisco System's excellent white paper about change management, available at http://www.cisco.com/warp/public/126/chmgmt.shtml. This white paper provides a great overview of change management and gives detailed examples of process flows. The white paper also provides examples of change management documentation, which can help kick-start a new change management process in your organization.

Q 1.3: How can I prevent overzealous administrators from making unauthorized changes to network devices?

A: First, realize that many administrators feel that they're doing users and the company a favor by performing so-called "minor" changes without following their company's sometimes complex change-management process. Some administrators are frustrated by the politics involved with making a change to a network device, and dislike the fact that they can't simply reconfigure their routers when they need to do so. Your first step is to overcome that mindset and make sure that all administrators understand the purpose and benefits of the change-management process:

- In the end, change management reduces work. Changes are less likely to cause failures and recovery is easier in the event of a problem, resulting in fewer late nights spent at the office.
- Change management shares the responsibility for making changes. A good change-management process includes several sign-offs, eliminating the need for a single person to bear the brunt of mistakes.
- A well-designed change-management process can reduce stress by eliminating the "do it now!" demands often placed on network administrators. The process can help absorb that stress by regulating change requests into a manageable stream.

In addition to changing administrators' perception of the change-management process, you can take advantage of the fact that most network devices offer a physical means of ensuring changes occur only when authorized (that is, through passwords). I've worked in environments in which utilities were used to automatically change router passwords every day. Before they could make changes, administrators had to check out the day's password, which forced the administrators to follow procedure.

Other utilities can retroactively catch unauthorized changes. Doing so makes it easier to correct or undo unauthorized changes before they cause problems and to educate the offending administrator on the proper change-management process in your organization. Unfortunately, assuming your administrators have password access to your devices, there's almost no way to prevent them from making changes without following your change-control process.

✎ In a few years, you'll see a new class of network device-management application that builds upon the solutions already available. This new class of solutions will provide a complete front end to device management, letting administrators use a friendly graphical user interface (GUI) or even an intermediate command-line command to make changes. The application will push the changes to the network devices using the devices' configuration passwords. Administrators won't actually know the configuration passwords; instead, the administrators will authenticate to the application separately, perhaps using their regular network-security credentials. The result will be a front-end application that provides business rules and processes to the change process, then pushes authorized changes to devices on the back end.

These types of applications already exist for network server management. Aelita (http://www.aelita.com), for example, makes a suite of applications that provide this type of management interface for Windows networks. Network device management hasn't caught up yet, due in part to the variety of devices in common use on large networks. But these solutions are on their way. Companies to watch for these applications include AlterPoint (http://www.alterpoint.com) and Network Mantra (http://www.networkmantra.com); both emerging leaders in network device management.

Today, you can utilize software tools such as Tripwire (http://www.tripwire.com), which periodically logs onto your network devices and compares their configuration with a known-good *baseline* configuration. Changes to the configuration generate an email alert, giving a senior administrator the opportunity to analyze the change and either accept it—making it part of the baseline—or reject it, causing the original baseline configuration to be restored to the device. Other tools, such as AlterPoint's DeviceAuthority, add detailed change-management reports, which you can use to not only review the specific changes made to your network devices but also to get a better idea of the type and volume of changes made to your devices over specific periods of time.

Q 1.4: How can I ensure uniform device configuration throughout my organization?

A: Companies with a large number of network devices often have difficulty maintaining consistent configurations across those devices. The benefits of consistency are fairly numerous:

- Consistent configurations make it easier to train new administrators and make it easier for administrators to take over each other's tasks based on workload.
- Consistency improves network reliability by using tried-and-true configurations on all devices.
- Consistency simplifies troubleshooting because the standard configuration has predictable, known behaviors. In addition, deviations from the standard can be easily detected by simple file comparisons.

Some of the highest-end network device management solutions include the ability to enforce consistent device configurations within an enterprise. A senior administrator develops configuration policies, which are enforced by the software on new configuration changes. Most packages with this capability can also review existing configurations for compliance with policies, allowing you to retrofit the software into an existing environment and clean up inconsistent configurations.

You don't necessarily need fancy software to enforce consistency, though. You can create configuration templates quite easily, making it easier for other administrators to use the same configuration settings and syntax across your organization.

Start by configuring a single device to be a model of your new, standardized configuration. Get the device's configuration into a text file either through TFTP, FTP, HTTP, or whatever other means the device supports. Then modify the text file—adding comments where necessary—into a template. Listing 1.1 shows an example for a Cisco Catalyst switch.

> Note that the exclamation marks in this sample file are comment lines and don't affect the actual configuration.

```
.
!
!! For Cat switches / firmware 3 / template v4
!
interface FastEthernet0/1
 description [officename]_local1
 duplex half
 speed 10
!
interface FastEthernet0/2
 description [officename]_local2
 duplex half
 speed 10
!
interface FastEthernet0/3
 description [officename]_local3
 duplex half
 speed 10
!
interface FastEthernet0/4
 description [officename]_local4
 duplex half
 speed 10
!
interface FastEthernet0/9
 description [officename]_backbone
 duplex full
 speed 100
!
interface FastEthernet0/10     !!! Omit speed and duplex
 description [officename]_lab  !!! for autoconfiguration
!
interface: FastEthernet0/11
 description [officename]_admin
 duplex half
 speed 10
!
interface VLAN1
 ip address [ipaddress] [subnetmask]
 no ip directed-broadcast
 no ip route-cache
!
.
end
```

Listing 1.1: An example for a Cisco Catalyst switch.

☞ Use version numbers! This sample configuration not only includes its own version number, but also includes a comment to tell administrators which type of device and which version of the device's firmware is required to use the template.

Notice in this example that several replacement variables, in [brackets], are included in the text. To use this template, simply use a text editor's search and replace feature to replace the variables. For example, replacing [officename] with newyork will result in the desired interface names, newyork_local, newyork_backbone, and so forth. Create a separate document that describes the variables in use. For example:

- [officename] = Replace with the name of the city in which the device is installed.
- [ipaddress] = Replace with the device's VLAN1 IP address (obtain the IP address from master tracking list).
- [subnetmask] = 255.255.255.0 in all field offices and 255.255.0.0 in all labs (see IP master tracking list for exceptions for certain subnets).

After you edit the file, you can load it into the new device. Most devices support a means of loading configuration files.

☞ Store your configuration templates in a version control system. Using a version control system allows you to retrieve older configuration templates in the event of a problem with a new template.

In addition to making it easier to ensure consistent device configuration, configuration templates make it easier to deploy new devices throughout your enterprise. Rather than having to follow a complex set of configuration instructions—or worse, configuring new devices from memory—you'll have an easy-to-use template that can you can quickly complete, load into the device, and place into production.

Whenever you make changes to your network devices, be sure to consider the changes for inclusion in your templates. For example, you might decide to add several RIP configuration commands to your routers to improve RIP performance; be sure to make those same changes to your templates.

Q 1.5: How can I ensure that all of the devices on my network are accounted for and under change management control?

A: Whenever I walk into a new client to talk about change control for network device management, one of the first things I ask for is a network diagram and an inventory of network devices. After all, a key in change management is ensuring that *every* device is under control, and the client's mix of devices often dictates which change management solutions I can recommend. A few clients have surprised me by having a complete inventory on hand; most have, at best, a few incomplete diagrams and an administrator or two who have a pretty good recollection of which devices are on the network. Needless to say, some research is required to generate a more accurate inventory.

Automatic Device Discovery

There are a number of products that can automatically search for managed network devices. For example, Microsoft Visio 2002 Enterprise Edition provides a fairly accurate network discovery feature. Microsoft Systems Management Server (SMS) 2.0 includes a similar feature. Some change management solutions, including AlterPoint DeviceAuthority and ReadyRouter also include device discovery capabilities.

One software package that I've had a lot of success with is Synexsys Inventory. It's actually a very full-featured network inventory package, including capabilities that extend to desktop software license management. Thus, many companies in need of asset control capabilities can benefit from it. It's also got a fantastic Simple Network Management Protocol (SNMP) discovery module, which seeks out managed network devices. Practically every device supports SNMP (unless you've disabled SNMP), so you can use Synexsys Inventory to quickly acquire an accurate report of your network device assets.

Any automated network discovery software will likely return devices that you don't care about. Network-attached printers, for example, will often respond to SNMP queries, and you might not have any desire to place them under your change management program. However, it's far better to have a complete, accurate inventory that contains too much information than to miss some critical device simply because everyone's forgotten about it!

Device Discovery and Security

Automatic device discovery can also be a great security aid. I've been to a number of clients at which, at one point or another, unauthorized network devices wreaked havoc on their networks. In one case, an unauthorized router started advertising itself and its routes. The router wasn't correctly configured for the network, so all routing operations became pretty unreliable in a short while.

If you've picked up a package such as Synexsys Inventory or a change management solution that includes an auto-discovery option such as DeviceAuthority, you can periodically repeat the auto-discovery process to see whether any new devices have popped up on your network. For example, Figure 1.1 shows DeviceAuthority's auto-discovery wizard, which queries for SNMP devices on your network.

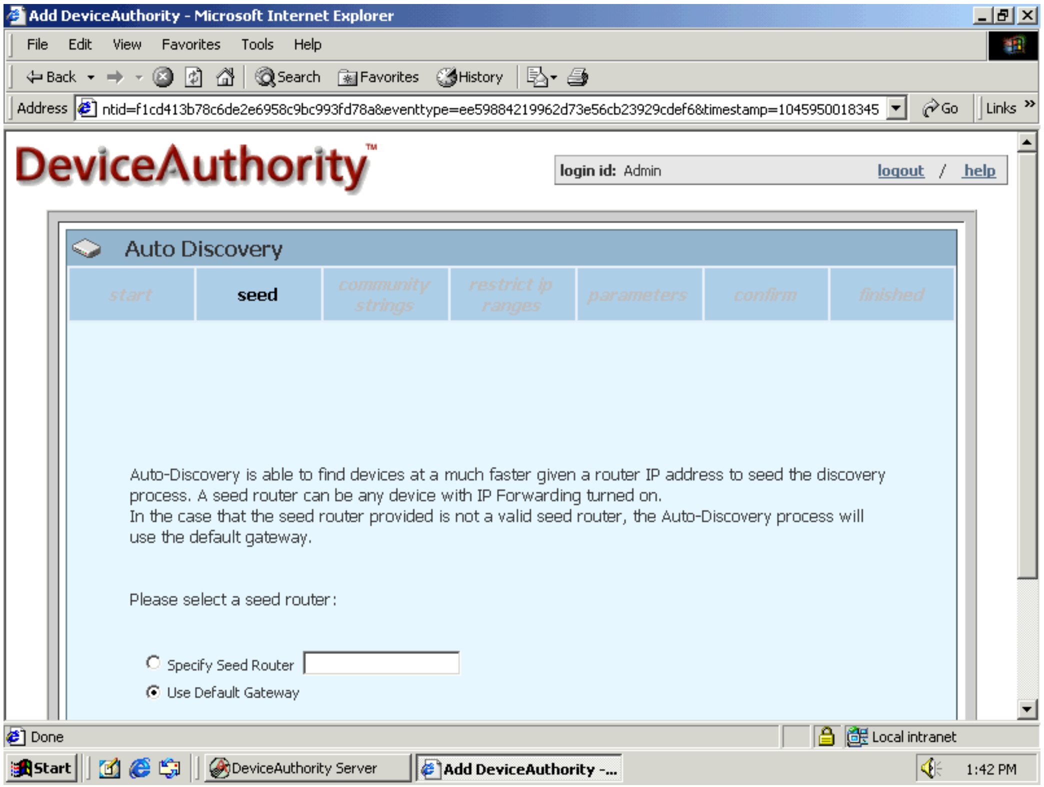
Add DeviceAuthority - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites History
Address ntid=f1cd413b78c6de2e6958c9bc993fd78a&eventtype=ee59884219962d73e56cb23929cdef6&timestamp=1045950018345
Go Links
DeviceAuthority™
login id: Admin
logout / help
Auto Discovery
start
seed
community strings
restrict ip ranges
parameters
confirm
finished
Done
Local intranet
Start
DeviceAuthority Server
Add DeviceAuthority -...
1:42 PM

Auto-Discovery is able to find devices at a much faster given a router IP address to seed the discovery process. A seed router can be any device with IP Forwarding turned on.
In the case that the seed router provided is not a valid seed router, the Auto-Discovery process will use the default gateway.

Please select a seed router:

○ Specify Seed Router

● Use Default Gateway

Figure 1.1: The auto-discovery function in DeviceAuthority can start with a seed router for faster operation

Most network device discovery routines simply query for SNMP devices on your network. Typically, that will return all of the managed devices on your network. It's possible, of course, that you have non-SNMP devices. For example, most inexpensive hubs and switches won't respond to SNMP. However, those devices are inherently unmanageable. You don't need to worry about keeping them under change control, because there's nothing in them that can be changed. Capturing a list of your SNMP devices will tell you which devices can be configured, and those are the devices that you'll want under change control.

Device Discover and Documentation

One reason I'm a fan of Visio 2002 Enterprise Edition is its network discovery feature. In addition to locating the devices on your network, it automatically constructs a complete network diagram, showing how the devices are connected and even labeling router interfaces with the appropriate IP addresses and other information. The discovery process can take a very long time to complete, and the finished diagrams usually need some cleaning up and rearranging to be more useful, but it's a great feature that can be a real timesaver. The feature works best when it's running on a very powerful workstation and when your routers use a routing protocol (such as RIPv1, RIPv2, or Open Shortest Path First—OSPF) to exchange routing information.

SNMP, Discovery, and Security

For security reasons, some organizations choose to disable or in various ways restrict the use of SNMP on their network devices. Unfortunately, doing so will prevent most auto-discovery routines from locating managed devices, because those routines generally rely almost entirely on SNMP. Thus, you'll have to manually configure each device in the change management solution rather than having a tool seek them out automatically.

If your organization has SNMP enabled on your network devices, you'll have an easier time generating a network device inventory. However, don't underestimate the potential security risk that SNMP and auto-discovery can represent. One of the biggest battles an intruder fights is finding out the structure and architecture of your network, and with SNMP enabled, those intruders can utilize auto-discover routines just as easily as you can. Auto-discovery is also very, very difficult to detect when it's in progress. You can minimize the effectiveness of an intruder's auto-discovery efforts by changing your devices' SNMP community strings to something very difficult to guess, and by taking steps to secure your physical network so that unauthorized persons don't have access to it.

> ☞ There's been a recent upswing in so-called *war driving,* the practice of driving around with a wireless laptop looking for unsecured wireless networks. Most war drivers are simply looking for free Internet access through your wireless LAN; others might have more nefarious purposes. An unsecured wireless LAN can represent an easy point of entry for intruders who want to conduct an SNMP auto-discover on your network. Changing community strings, using wireless encryption to restrict wireless access, and configuring wireless access points not to bridge SNMP packets will help prevent your wireless LAN from becoming a way for intruders to learn more about your network infrastructure. Of course, if you disable SNMP in your wireless access points, you won't be able to use automatic device discovery features as effectively.

Q 1.6: How can I reduce network device problems through change management?

A: Throughout this book I've been preaching the wonders of change management as a means of preventing problems—or at least of easily detecting them. Change management, however, is really just a form of documentation, and it's documentation that can really help avoid—or quickly determine the case of—network device issues.

Inventory

One way to prevent network device problems is by maintaining an ongoing inventory of your network devices. This inventory should include items such as:

- Hardware inventory, software versions, module descriptions, and so forth
- Port assignments, connected media type and speed, and other logical configuration values
- Routing configuration, VLAN configuration, access lists, and other security concerts
- Any out-of-band management configuration
- Cable requirements

For example, suppose you receive from your change management system a notification email that informs you that a particular router had a particular interface's media type changed. You examine the change and find that it was switched from a 10Mbps Ethernet connection to a 100Mbps connection. "No problem," you think "all the hardware is 100Mbps anyway."

Except that this particular router is connected with CAT3 cabling, which doesn't reliably support 100Mbps connections. The router starts to experience periodic failures that are tough to pin down. What you should have done is examined that configuration change in light of the router's hardware inventory, which would specify that it was using CAT3 cables. You'd know to get right over and switch to CAT5 or better cabling to prevent problems.

Your inventory should also contain a list of users or services that would be affected by the device's failure. When a network problem does occur, you can search through your documentation to match affected users or services, and quickly narrow the problem to the devices servicing those users or services.

Change Management

As I've mentioned before, change management over devices' running configurations is critical. You need to maintain—through whatever means you prefer—a running list of changes, a backup of the current and previous configuration files, and so forth. Most network device problems occur as a result of a configuration change; thus, being able to quickly spot the change and roll back to a known working configuration is your best weapon in the fight against network downtime.

Audit Trails

Configure your devices, if possible, to use Remote Authentication Dial-In User Service (RADIUS) or Terminal Access Controller Access Control System + (TACACS+) accounting. Configure them to synchronize their time with a Network Time Protocol (NTP) server so that accounting messages will include timestamps. Accounting makes it much easier to track the source of a problem in time. This ability is especially important if your device configurations have changed multiple times between configuration management pulls, because the accounting log might be your only indication of which configuration settings have changed, who changed them, and when the changes occurred.

Network Topology

Always, always, always have an up-to-date, accurate diagram of your entire network. Troubleshooting any kind of problem—especially with routers—is much easier when you have a diagram that shows how things *should* be working. Maintenance of your diagrams is especially important and should be part of a formal change management process. Although up-to-date diagrams can make troubleshooting vastly easier and more efficient, outdated diagrams can be a significant hindrance.

Q 1.7: We're an ISO9001 shop. How can we incorporate network device change management into our processes?

A: ISO9001 can seem like a difficult framework in which to run a business because you can't seem to do anything without following a flowchart. However, this requirement is a good thing and designed to make sure that your business' results are predictable and repeatable. In fact, ISO9001 is an excellent framework for change management—it forces you to use a formalized, repeatable process.

The easiest way to incorporate network device change management into your processes is to create an official ISO process for doing so. Figure 1.2 shows a sample flowchart that you can use as a starting point.

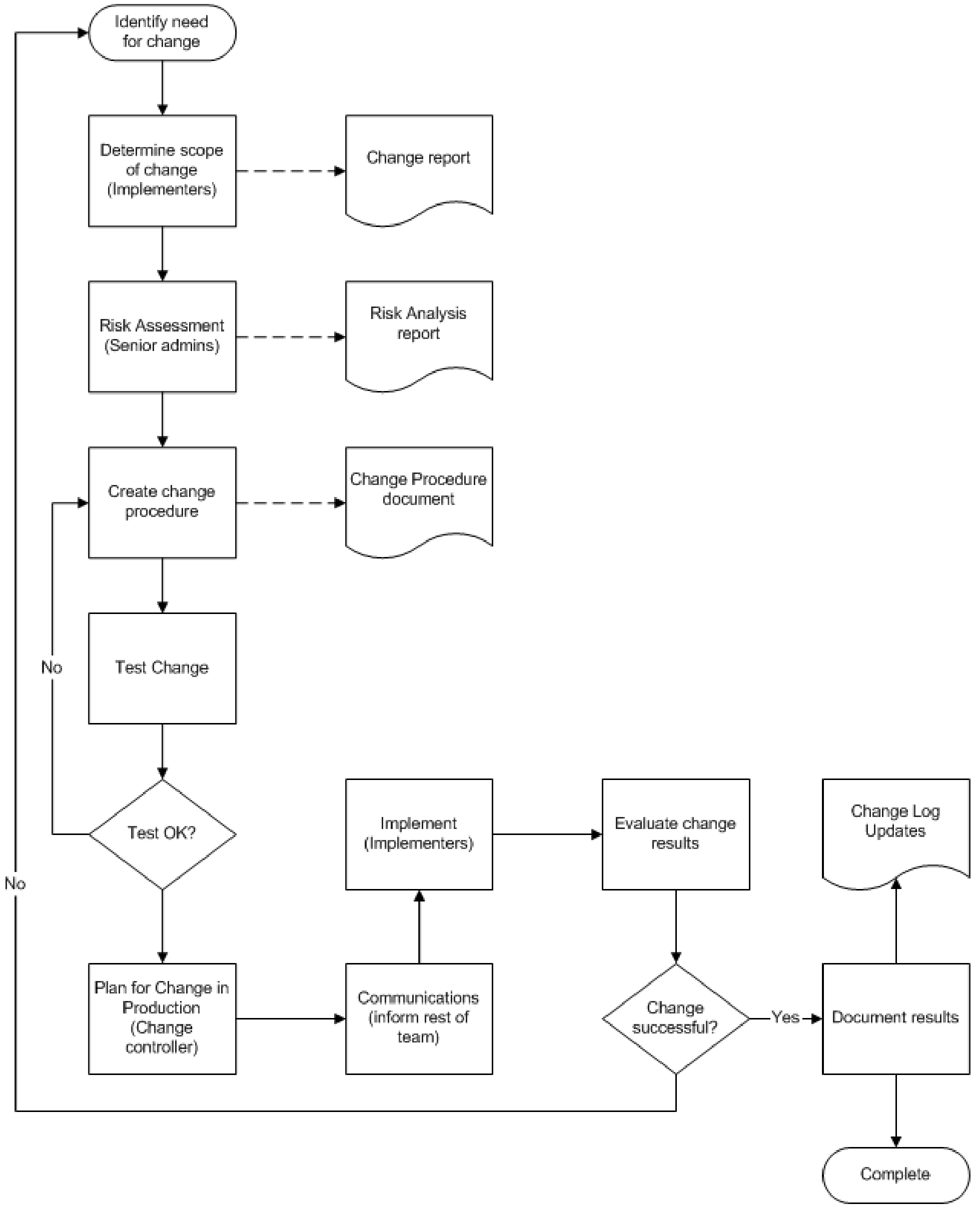

Figure 1.2: Sample change management process flowchart.

This flow incorporates many of the change management elements I've discussed in previous tips:

- Identify the specific change that needs to be made.
- A technically knowledgeable person then needs to figure out exactly what that change involves. For example, adding a new office may involve adding new static routes to a number of routers or simply involve a bit of additional routing protocol traffic with no actual configuration changes required. This scope—essentially a list of everything that will change—needs to become a formal document.
- Senior administrators should evaluate the risks of making the changes—as I've touched on in several earlier tips. Their analysis should be documented in a formal risk analysis, and the decision to proceed should be based upon that analysis.
- An official change procedure—the precise steps that will be taken to implement the requested change—must be documented and approved.
- The procedure should be tested in a lab environment if at all possible, particular if the change is deemed risky. If the test is unsuccessful, the procedure can be modified and the test repeated.
- If the test goes well, a change controller needs to plan for the change to actually be made in production. This process might involve the controller ensuring that all devices are currently backed up in case the change fails in production and a rollback is required.
- Communications involves informing anyone who might be affected by the change so that everyone involved can be alert for signs of failure and those affected don't misinterpret the effects of the change as a problem and initiate corrective action.
- The implementation should be evaluated for success. Ideally, the change plan should identify metrics for success so that the change can be definitively proclaimed a success or not.
- If the change was not successful, it needs to be rolled back and the entire process repeated to find out what went wrong. A change management solution can be especially effective at the rollback step, allowing devices' former configurations to be restored quickly.
- If the change was successful, it needs to be documented. The baseline configurations for affected devices should be updated to reflect the change, and the devices' configurations should be backed up. That way, if the device fails for other reasons, a restore will also include the recent changes.

Your company's change controller—usually a network device administrator—should own this change management process. You'll likely need to tweak this flowchart a bit to meet your specific needs, but it should give you a useful ISO-quality start.