OARtech
February 8, 2006

Introductions

Oarnet update
Tony Eller
Slides are on OARtech web site.

All phase 2 schools are directly connected to the TFN. Last Mile update included Kent State, Rio Grande, and Battelle as all connected directly. Outages: NEOUCOM's OC-48 card failed but they had a spare at Kent. School staff was able to install the spare card. The outage lasted about an hour. Oarnet is building a redundant link through that area and should have the link up in about a week. Oberlin has a problem increasing bandwidth on their serial
DS3 while they are waiting for a last mile solution. Springfield POP is first 10GigE circuit operational for OSC-Springfield, and the first one for the TFN Network. Partners at some K-12 sites are on or will soon be (Toledo Public,
SCOCA) and WVIZ public broadcasting station of Cleveland. Resources available for OARtech members are on statseeker.oar.net and portal.oar.net. Oarnet is currently working on a lighting ceremony for ETEC, and Windows of the Future.

Is Looking Glass coming back? As they have been turning off telnet on the routers and turning up SSH, Looking Glass has been breaking. Oarnet will be changing it to using SSH, and at that time it will work again with the new routers.

How many here have a completed Disaster plan? There were very few hands that went up.

Continuity and Disaster Recover Planning Steve Akers Slides are on OARtech web site.

Definitions:
Disaster Recovery (DR) is a part of business continuity with an objective to restore critical business processes. It focuses on data recovery with a timeframe of the first 30 days. Solution is often hot site recovery, or some other secondary site.

Business Continuity (BC) objective is to restore business back to prior state.
The focus is to return to normal with a timeframe of more than 30+ days.

DR/BC plan – the methods, processes and procedures needed to minimize the impact of a disaster. It includes the guidelines and activities required to restore the systems and operations and the business to the state prior to the disaster.
It should be well written and properly tested and allows personnel to administer recovery efforts with results in a timely restoration of services.

RTO – Recovery Time Objectives

You need to know the impact of the disaster and it helps you focus on what is most important. If you had only one person to recover things, what would you send them to first? You need to make sure it focuses on the process not the technology so that all the business pieces needed are included. The personnel responsible should be a body of people that include both technology and business people that represent the entire organization. The plan needs to be simple and should include maintenance and updates to the plan to make it a living document. You at least need to know that changes have been made. Staff needs to have a general understanding of what they need to do in a disaster situation. The plan needs to be unique to your organization so often templates are not a good idea. You need to understand the due diligence and compliancy needed. You have to get started and admit it could happen to your organization.

A DR effort should do a readiness assessment to determine how ready they are to start a DR effort. The flow of the document is VERY critical and has to seem natural in the way they do things. You need to know the critical paths and equipment and unite the technology with business processes. A non-technical manager may not know a specific server, but will understand the business processes. It is a stressful situation so you need to minimize decision making. Training has to be recurring, and must go through the entire organization. You need know what changes have been made (organization best practice is to have change management in place) so that in a disaster you know when something has to be changed.

He uses a 12 step approach to provide a proper framework to building a plan:

1. Management Commitment – you must have support from the top down.
2. Planning/steering committee – a body of people from the organization that is making the decisions. They provide guidance and should be cross departmental and help clear out the road blocks.
3. Risk evaluation – people will take advantage of the organization in a disaster. So mitigate these risks before the disaster occurs.
4. Business impact analyses – catalog the critical systems and processes and quantify the financial loss related to the outages. You need to establish a practical RTO and recovery objectives – define your pain thresholds. Illustrate your critical paths and what process might affect other processes.
5. Determine recovery strategy – was defined at a high level in the previous phases. Determine the types of continuance with personnel, technology, processes and procedures. This will help you determine what you want to use: hot-site, cold-site, second office, etc…. This may actually not be fully determined until you are done with the plan.
6. Data Collection – do an inventory and repository of all the resources, documents, procedures, vendors, personnel, contracts, and records. You need to make sure you have the contract numbers readily available.
7. Emergency operations center (EOC) – establish an EOC that is away from the recovery site that can be used for managing the recovery effort. You want 2 to 3 places.

8. Organize and write the plan – setup a skeleton/framework as to setup of the flow of the document. Allow it to be vetted by the business organization. Then fill it in to a formal documentation. Keep getting the business buy in all through the process. You want to be able to give a person their section of the document and they don’t have to worry about any other piece of the plan.
Set up work flow diagrams and color code the documents so that the various teams know what they are responsible for. Set up a disaster impact matrix with the important processes across the top and the actual pieces that are needed to run the various processes down the left column. This can tell you which of your processes are broken in various disasters. The plan is written to worst case, but in a less than worst case disaster you need to know what section of the plan needs to be activated. If a disaster affects only a piece, then the affected processes are highlighted. Business people can probably work through the plan with just the work flow and the disaster impact matrix.

9. Develop the materials – these would be used for both training and testing.
You should establish the scope, criteria and type of test (full or tabletop).

10. Awareness and Training – familiarize the people with their roles; use repetitive learning with different staff participating.

11. Testing and maintenance – expect something to fail. If you don’t have a failure, you probably don’t have people paying enough attention. Then make improvements to the plan and improve the flow and establish maintenance process.

12. Approval – get a final approval of the first draft with testing results. It is a living document so will need continual updates. The plan itself should not include the step by step detail to recover a server, but an overall that says you restore from backup. It should point to existing operations documentation, so you are using working documents for appendixes.

Failing to implement a disaster recovery plan is an indication of organization negligence. There are components of these requirements in Sarbains-Oxley and HIPPA.

If you document a specific group of people to execute the plan, some of the people may be gone in a disaster. You need to document the contingencies if people are gone. You have to do testing assuming some key personnel are dead.

When should templates be used? The actual process for building the systems would be an appendix to the plan not in the plan itself so it can be a template. However, you would not use template for the higher level, but at the lower technical level.

Comment from the floor
One school had an Xserv system broken into and then used to try and exploit windows systems. Found that the Xserv was compromised via SSH password checking, and then in /var/temp. Number one problem is generally passwords.

There is a book by Ray on OSX that talks about the crucial needs of password policies. This just points out that you shouldn't forget about the Apple systems when considering exploits.

OpenVPN
Ransel Yoho
Slides are on OARtech web site.

OpenVPN is an open source, cross platform, client/server VPN software.

VPN is an extension of a private network via an encrypted tunnel over a public network. SSL/TLS is the Secure socket layer v3 and transport laser security (SSL 3.1 = TLS 1.0). PKI is Public Key Infrastructure and allows for an arrangement which provides for 3rd party vetting of and vouching for user identities. It is usually carried out by software at central location together with other organizations.

Cryptographic objectives are to provide confidentiality, integrity, authentication, and non-repudiation. Its primitives include symmetric encryption, message digests, and asymmetric encryption. Symmetric provide confidentiality, message digests verify message integrity, and asymmetric encryption provides authentication and non-repudiation.

Typically the VPN would use addresses out of the private address space. You would want to avoid the ones used most often.

Types of VPN include different types of tunnels, user space and gateways. The types of tunnels are PPTP, L2TP (layer 2 transport protocol), IPsec (some vendor interoperability, and complex). User space include vtun, ssh, and OpenVPN. The SSL application gateways are clientless and proprietary. Some clientless SSL tunnels are using browser technologies (java) that run on the pc and thus can have problems that extend from the pc or browser. Some also negotiate improper encryptions.

Why look at OpenVPN? Is cross platform with client and server (Linux, Windows,
and Mac) and is based on the TLS/SSL openssl library. It is easy to install
and configure and can do NAT traversal (a problem for IPsec). It interoperates with other VPNs such as PPTP and IPsec.

OpenVPN PKI would have a separate certificate (or public key) and private key for the server and each client with a master Certificate Authority (CA) used to sign those certificates.

Supports bidirectional authentication based on certificates. Both server and client authenticate with other by verifying the cert. It only needs it own cert/key and will only accept clients whose certificates are signed by the master CA cert.

References:
OpenVPN 2.0 HowTo
http:/openvpn.net/howto.html

OpenVPN and the SSL VPN Revolution
http://www.sans.org/rr/whitepapers/vpns/1459.php

OpenVPN can be configured via bridging or routing.

ISSA
Brian Mueller, OSU

He is trying encouraging people to apply for the scholarship. Applicants have to be an ISSA member, but can be a student member and do not have to be in Ohio. They are giving away $10K total. Please encourage people on your campus to apply. Student memberships are minimal. Also they are interested in forming ISSA student chapters.

Lunch

ArcSight
Chris Hipskind
Dave Leslie
Contact information is on OARtech's web site.

Enterprise Security Management space is where they provide product. Their software enables enterprises to collect correlate and manage massive amounts of data from heterogeneous systems.

Anything that provides a security event management logs, can be a feed to their product – ArcSight Management. With this you get real-time analysis and correlation.

The 3 areas they work in: Help me be more proactive in hardening the perimeter; Show me who on the inside is a threat; and Compliance to regulations.

They take the logs and events from different systems and correlate events in real-time and report it back with work flow and call management. They use SmartAgents which are collectors to pull the events in, then feeds to ArcSight manager which archives and stores the event. They then provide a console to view the data.

They support 200+ devices. They also have a FlexAgent which allows you to create an agent for those devices they do not have built in agents for. They have agents out in the network collecting events that filter, aggregate, normalize and categorize the events. It uses SSL between the elements. He showed an example of taking an event from a PIX firewall and showed how it normalized and categorized the data.

You can have it setup as a single box deployment, but you really want to have the smart agents out closer to the perimeter to allow more secure transfers of the logged events, and use less bandwidth. You can have a concentrator that is closer to the perimeter with multiple agents for collecting data.

The correlation and workflow product does correlation with historical forensic investigation, threat detection, and risk-based prioritization and can do alerts. It looks at agent severity, asset criticality, severity, model confidence and relevance to set the priority scoring level. You can make some decisions on tuning how it sets the priority.

Does it look regular events in historical data? They have a piece of the product that can see patterns of events that can be mined out of the data.

Two consoles are provided. One is for Security Analysts and has full functionality, forensic investigations, visualizations, authorizing, and centralized administration.

They have over 360 standard reports. They do have a GUI report writer that
does not require SQL knowledge. The backend database is Oracle. In Oracle you setup a retention scheme to determine how much data you need right now and how much you want to keep for historical analysis.

He showed a demo of the Analysis/full console. They have several "dashboards".
Overview dashboard had top source IPs, protocols, fired rules, etc…. The dashboards are updated in real time. You can group the assets into what groups that you would like based on your own needs. Showed a Worm event dashboard which showed the attacker and who it was attacking. You can click on an IP to see the active channel to look at all events related to the specific IP. You can also open active channels that can show many different subsets of events (e.g. all medium priority events with a specific IP or group of items).

Some of the device vendors' data used in the example was Snort, Cisco, McAfee (ePolicy Orchestrator), Check Point, Tripwire, Intruvert, CMA, radius logs,

You can see the chain of events to see how a specific event was correlated to be a problem. You can drill down into the event to all the information associated with the event. You can view the rules to see why a rule fires and to create some of your own rules. You can setup a rule that can automatically moves a target from an active to hostile list, or have a rule that can determine for an compromised target and thus get a list of all compromised targets.

You have multiple layers of administration for the software (e.g. view only, or admin for only some assets) and can setup zones within a network to use them in your logic. You can setup locations that will actually look at geographical views of those locations. Reports can be done by IP, groups, assets, etc….
You can create a case by adding records to the case, and see how they are worked and the impact etc....

Is product is meant for security event management, is not meant to be a network polling station, and it can categorize any type of event data.

List of devices supported is available on the web, but they are adding new devices very rapidly so the posted list may not be up-to-date. They have connectors to Remedy type applications to do automatic alerts on a high priority events.

New business
Nominations for vice chair were opened. Notice of nominations will be sent to the list.

Meeting adjourned.