

Cisco IT Data Center and Operations Control Center Tour



Process in the OCC: Staff, Process, and Ongoing Process Development

3. Process in the OCC: Staff, Process, and Ongoing Process Development

OCC Staff

Figure 1. OCC Staff on the Bridge



Dick: “Now we’re inside the gallery behind the operations center. In this environment we’ve got two or three people on shift at a time. During night shifts, we can typically manage with two people, but depending on the day of the week, we may have three. We generally look at it as a three-person model. Everybody on the team is cross-trained.

“The responsibility of the operations command center personnel is to ensure that when something fails, a case is immediately opened and within five minutes they send out a notice, telling people about the situation. They don’t ask questions—that’s critical to managing the OCC—they just open the case as fast as possible.

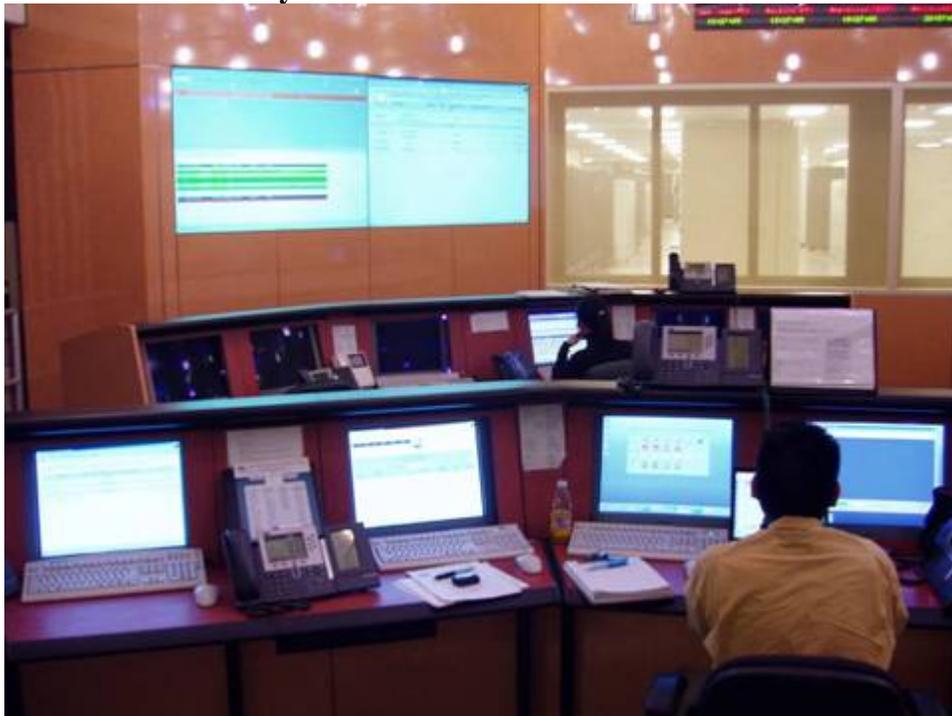
“When they open a case on a failed resource, a notice is automatically sent to all the people who are associated with that resource, so they can begin to take steps to fix the

problem as quickly as possible. It's not our job to identify these people; instead, people who are responsible for that resource, or who are interested in finding out about the health of that resource, subscribe to e-mail aliases that we use to alert them via pager. That way the OCC staff can identify the Cisco employees interested in that particular resource.

“Our OCC staff manages the case in our trouble-ticketing tool, from the time the alarm is raised until the time the immediate problem is solved. They require the on-call support engineers to phone, to get on the bridge, and not leave the bridge until after the problem is resolved or they've negotiated with the OCC staff as to when they're going to come back. That means that the support engineers can leave the bridge, but they have to tell our staff when they're coming back, and if they leave prematurely, our OCC people are going to call them right back until they get back on the bridge.

“If the on-call support engineers for the down resource don't respond within five minutes, our staff will escalate to the first tier – usually that person's manager.. The OCC staff will try the on-call duty support pager first, then the next thing they do is try to reach the on-call person's mobile phone. If they can't reach the duty engineer by pager or phone, they will escalate to that person's manager by paging or calling the manager. If that doesn't work, they page or call the director, and after the director then the VP, however high up they need to go to get somebody to respond to the outage situation. Also, when the OCC has to go to the VP level, they call me in and together we send pages to a couple of strategic people – the people in IT who can make things happen -- and that will solve the problem pretty quickly.”

Figure 2. The OCC Day Staff



Q: How do you get by with only three people on staff to handle all these resources?

Ian: “That’s been a real victory for our process, and our automation tools. Four years ago there were upwards of 20 people handling 400 or 500 incidents per year. Now we have this down to a team of 14—with 3 people per shift; 24-hour coverage seven days a week; covering nights and weekends; it works out to about 14 people. What’s amazing is that in that time we’ve gone from about 500 incidents per year to about 9000 incidents per year, with a reduced number of staff.”

Q: Why did the number of incidents grow so much?

Ian: “Part of the increase in incidents came when we went from monitoring just the five global production data centers to monitoring all 40 plus data centers; that is, when we started monitoring and doing incident management for the engineering data centers and server rooms as well. The rest of the increase took place because we started getting much more stringent on reporting and recording every single incident that took place.”

Q: So what other kinds of things do they do in the OCC?

Dick: “They manage nightly tape backups, month-end processing, and quarter-end financials.”

Q: What is month-end processing?

Dick: “Month-end processing is a large set of batch jobs that need to be run, in parallel, to close the books at the end of the month. To do that we need to stop and sum up the total revenue coming from our sales applications, and the total costs coming out of our manufacturing partner applications. We have to stop these processes long enough to take a snapshot of the financial flow within Cisco. This whole enterprise resource planning process runs on an Oracle database, which Cisco IT is migrating to Oracle 11i. In fact, our manufacturing, sales, finance, and human resources applications all run on the Oracle database platform.

“The OCC team coordinates with a product scheduling team in my group that works very closely with the month-end processing teams to ensure that we get our quarter-end financial reports run on time. At this point we’ve got the month-end processing to the point that it’s predictable and consistent, and the OCC team almost always handles it themselves.”

Q: What types of people do you hire to handle all these different activities?

Dick: “We’re looking for people who can work with people, and who can talk professionally just as well when they’re sitting around a table as when they’re on the bridge in the middle of the night, working with people who are screaming because they

just woke up to respond to a dozen pages and their baby is crying in the background and their manager is on the line yelling at them.

“We hire the OCC staff to be able to manage problem situations, to be calm in a crisis, and handle little problems or disasters by the book. I think we’ve been very successful in hiring the right people over the years.”

OCC Process: Coordination, Communication, and Documentation

Figure 3. Ian with OCC Staff



Ian: “The goal of the command center is to communicate, coordinate, and document during incidents. Our team doesn’t provide technical resolution; technical response is provided by one or more members of about 150 teams worldwide, who have P1 and P2 responsibilities for applications, systems, databases, and infrastructure such as switches and routers, for which their team is responsible.”

Q: That’s pretty catchy: Communication, Coordination, and Documentation. Could you describe that process for us?

Communication

Ian: “Communication is letting the right business or support groups know about the existence of a down or degraded resource. When a resource goes down, the first thing that happens is our team sees it on the monitor in front of them, and in front of the room. It starts with paging the person on duty responsible for supporting that resource, within 5

minutes of detecting a down or degraded service, and we expect a return phone call to the OCC within 5 minutes of paging out, no matter what the time or day.

“If we don’t get the expected response, we start our escalation process. That’s really aggressive, and requires that our team and the duty support engineers are very responsive. We ask a lot of them. And if they don’t respond, we escalate to their manager, and then their manager’s manager, to make sure we get the right engineering resources working on the problem as soon as possible.”

Coordination

Ian: “Coordination is more complex. Some failures may affect several applications or several resources at once, because so much of what is on our network is dependant on something that’s down or degraded, so we may have several duty support engineers working on a problem together from a variety of angles.

“Our team will coordinate the overall response for multiple groups. We bring them onto a phone bridge together. For example, maybe we need a database administrator, a network person, or a system administrator, all together very quickly to look at the problem to determine what really is causing the incident. We have three MeetingPlace bridges reserved specifically for this team 24 hours a day that we use, as necessary, to get people talking together.

“Our goal is to get all the duty support responders necessary to identify and resolve the problem, and to keep them talking, so they can keep the resolution process moving forward. One of the skills on this team is knowing when the engineers are working productively and when they’re not. Sometimes things get quiet on the bridge, and it’s up to us to know the difference between when things are quiet because everybody knows what the problem is and they’re working madly to fix it, or they’re quiet because nobody knows what the problem is, and nobody wants to admit it. So, one of their skills is to ask questions, like “Why has nobody sent an update?” or “Can you give me an update?” or “Did you just mention you found the problem?” or just ask gentle questions to get people talking and moving along, to find out whether we’re progressing or not.

“It’s important for people on our team to know when to prompt the support engineers for the next step. We train them to ask gentle questions, like:

- “Do we need to bring in additional resources?”
- “Has anybody let their management know about this?”
- “Have you told the business clients about this?”
- “Can we page out for that person you just talked about, the one who might know something more about this?”

“A lot of their communication skills are people skills. This team is hired for communication skills, not technical skills. Of course it helps to have a technical bent, but the technical response we expect to come from the data support groups.”

Documentation

Ian: “Everything that happens during an incident, from what we call a short term fix and all the way to a full recovery, is captured and documented by this team in real time. This keeps other technical resources and management aware of what’s being done to fix the problem, and helps the support engineers because it offloads the documentation and outbound communication from them and lets them concentrate on fixing the problem. We enter all this information as the problem resolution unfolds into an Alliance case.

“Alliance is a trouble-tracking application that we have customized for our use. For example, we added an extra tab called the Outage tab, which is specifically for capturing what resources went down, when they came back online, and what the impact areas were to the business: CEC, Cisco.com, manufacturing, ERP, and so on.. We also record which support engineers are on the bridge, how, why, and when the bridge is set up, who responded, any important decisions that were made, and any actions taken.

“Anything that’s not captured as a data value, we try to capture as a comment in the activity log. When we view the case later, we want to be able to go through the combination of the case data and the activity log to recover and piece together the sequence of events that show issue through to resolution.”

Ongoing Process Development

Process Redefinition

Figure 4. Dick Corso in the OCC



Dick: “I came to Cisco from Lockheed Martin. I think I was hired because I have a very strong process and methodology background in data center operations. The very first day I arrived, in 2000, I noticed that sometimes there would be no alarms at all, and other times there would be a lot of alarms, and it seemed to depend on who was on shift. The problem I saw really was during the times when there were no alarms, which sounds odd.

“But what was happening was part of the culture then; people helped each other. And if an engineer screwed up during a change and some resource got pulled offline that wasn’t scheduled, the people in the OCC would turn off the alarm and work with them to get it back online, rather than have a lot of people paged to try to fix it. So engineers got used to extra cooperation from the OCC, and they weren’t as motivated to use the change management process as scrupulously as they should.

“Problem tickets weren’t being created, availability was inflated, and people weren’t motivated to fix the procedural problems as long as no one was noticing it, or measuring it. In fact, at that time people were being rewarded for the reduced number of trouble tickets being created. It was seen as a sign that there were fewer problems, which was rewarded.

“Another issue was that people in the OCC weren’t always at their posts. I started to show up at different times, like in the middle of the night, and sometimes there would be nobody here. What happened was that they had gotten to know pretty well what the normal change and work schedules were, and they knew when things were likely to fail and when things were likely to be smooth. So if there was something critical going on they’d be right here, but if things were quiet they might be off having a smoke, or out to dinner, chatting, or in the closet sleeping. This also resulted in fewer trouble tickets being created, because sometimes people weren’t there to open the ticket.

“So we went through period of change when I got here. After communicating to the team how important it is to follow procedure, and how critical it is to the company to maintain control over the whole change management process, we selected a day to start out with our new process. We all know it in this group as “holding the line day,” sort of a military maneuver from my military days and from government contracting. From that day on, as soon as any resource became unavailable, for any reason, we created a problem ticket. On that day we went from almost no P1s per day to about 8 P1s per shift. We’d never seen more than 10 P1s in a day.

“People were really concerned, because a lot of bonuses in Cisco IT are tied to their resources being available, and meeting availability targets. There were a lot of managers who were unpleasantly surprised, because they’d been told that their systems were highly available and now they found out that this wasn’t true, and their bonuses might be at risk. This was all kind of a shock to people in IT.

“But after the initial shock wore off, people started understanding that Operations had become much more rigid, and much more predictable and reliable. After a while people no longer questioned it and it became part of our culture. It has resulted in P1 alarms returning to being pretty rare, and our availability statistics are back up to being very high, and we know that now it really means something. In fact, because we now track every single outage very carefully, and engineers focus on long-term root-cause analysis of problems and long-term fixes for those problems, our availability in many areas has hit all-time highs.

“My management style is a little more aggressive than most, and so I spent the first two years building relationships with people on my team and the support teams. My team knows that, in the middle of the night, any day of the week, if they need to get hold of me, they can, and they do. They could have support engineers or managers here on the bridge yelling about something that they did. And they can count on me to get on that bridge and take care of the problem, right away.”

Ian: “Another way to look at our process redefinition is to see it as our team focusing on our real core mission, and taking things that we used to do, which were not strictly incident management, and moving them out of the command center. Over the years some things had crept into our process that really didn’t belong there. For example, we used to help coordinate change requests: When someone would request a change on one part of the infrastructure, we would look up all the other parts of the infrastructure that would be affected, and contact all the other teams that were responsible for those areas to make sure that the proposed change wouldn’t affect them, or help schedule it at a time to minimize impact.

“But we decided that this was not incident management, and that the team requesting a change needs to coordinate its change request. This has not only reduced our workload but helped engineering teams to communicate with and understand all the other systems that their system affects, and all the other teams who work on them.

“Our team also used to spend time helping engineering teams through their changes through the night, or over the weekend. Well, that’s inefficient when our team is trying to jump from incident management to coordinating a change request and back, and we sometimes would end up with a conflict like, “I’ve got this P1 incidence I need to deal with, but this change request is also important, and I have to handle both at once,” which reduced our ability to manage incidents and keep our availability up. So we got our teams to stop all work that is not strictly incident management, which makes us more responsive during an incident, and allows us to do all this with as few resources as possible.

Process Automation

Figure 5. Ian in the OCC



Q: What did you automate to make it possible for fewer people to handle more than 10 times the number of incidents?

Ian: “We made a lot of changes. Part of our success came from adopting the Alliance Remedy trouble-tracking tool and adapting it to our needs. We took some of the capabilities in the previous tools we had that worked well, and we integrated them with the Remedy case-management system to make the actual process or procedure more efficient for us. We also built boiler plate incident templates to help reduce the amount of time it took to record an incident.

“And that’s what gave us the ability to take on managing incidents for the entire network and all the data centers, both production and engineering. This makes sense, because we have the monitoring systems, we have well-trained support teams, and we do this well.

“It’s far more efficient for a single team to monitor and document, to capture the information succinctly going into the case, figure out who to page for support response, capture all the values and variables and the issues and responses, and that sort of thing. It helps drive down costs to consolidate the various means by which information can come to this team so that it’s consistent. More, because we have a single standard method for

how we respond to incidents and capture the information, and a standard language we use with all the engineering teams throughout IT, it becomes more familiar and easier to use for everyone in IT.

“There are things we can’t automate, though. Even though we try to find ways to simplify the procedure, to automate wherever possible, we want to make sure we don’t lose the value that we have with this team, because the great advantage of having humans over automation is that we’re able to see the big picture and think creatively for new and unexpected situations. Now, 99 + percent of the situations that we have are pretty boilerplate; but the critical ones can be very different. So we build our procedures to be consistent, we have predictable procedures for the type of incidents that repeat most often, but we don’t try to nail every last detail to the wall. That way when something unusual comes up, people still have the latitude to manage those exceptional situations.”

Process Documentation

Q: Are these procedures documented or are they picked up by your team over time?

Ian: “They are documented. The documentation sometimes has to go through several changes until we get it right. We used to have separate procedures for different technologies, and for different environments, but found that this made our work too complex and made it hard to scale.

“Usually when we get new business capabilities and technologies that might affect our procedures, we first ask “How does this affect our current incident management process?” If we can, we try to adapt our current procedures rather than create new ones, because changing a procedure is difficult and time consuming and leads to a few initial failures and errors as people learn a new procedure. Because of that we try to make sure all our procedures are as independent as possible of any specific technical or functional area, since we know the technology will change over time.

“We also try to make our procedures independent of the applications being supported. Whether Cisco.com, CEC, or manufacturing ERP, we try as much as possible to make sure that our procedures are addressing all those areas consistently. That was one area of simplification that allowed us to really scale a lot. We went from a larger team with about 1700 cases in fiscal year 2000 to a smaller team with 8600 cases in fiscal year 2003. Simplifying our procedures helped make that happen.

“The next step for us was to engage in out-tasking, and to move to a two-location model. After we established and documented our processes in detail, we handed off responsibility for this incident management workflow to Cisco partners. We’ve tuned our processes to the point where they’re consistent, reliable, and capable, and when you can do that you’ve got something to outsource.

“You can give it to Cisco partners because Cisco is not in the business of doing incident management. We want to use our employees for our core business. We’re continually looking to the next levels of complexity, and when we have those support processes well

understood and well documented, we'll hand those off too. Moving our night shift to Bangalore India and using a two-location model for support was part of this process.”

End

In the OCC: Staff, Process, and Ongoing Process Development

You can go back to the Technology in the OCC: Telephony, Monitoring, and Backup section, move ahead to learn about how long term solutions and other changes within the data center and network are carefully managed to keep availability high, or you can go to any other part of the tour.

We hope you have enjoyed this part of the Cisco IT Data Center tour. You can contact your Cisco sales person to arrange an Executive Briefing Center visit, and request a live tour of the Cisco main production data center and operations control center.

For additional Cisco IT case studies on a variety of business solutions,
go to Cisco IT @ Work

www.cisco.com/go/ciscoitatwork

Note:

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore this disclaimer may not apply to you.



Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS
(6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International
BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on
the Cisco Website at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2004 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0406R)